On the role and challenges of CFD in the aerospace industry

P. R. Spalart

Boeing Commercial Airplanes
Seattle
USA

V. Venkatakrishnan

Boeing Commercial Airplanes
Seattle
USA
CD-adapco
Bellevue
USA

ABSTRACT

This article examines the increasingly crucial role played by Computational Fluid Dynamics (CFD) in the analysis, design, certification, and support of aerospace products. The status of CFD is described, and we identify opportunities for CFD to have a more substantial impact. The challenges facing CFD are also discussed, primarily in terms of numerical solution, computing power, and physical modelling. We believe the community must find a balance between enthusiasm and rigor. Besides becoming faster and more affordable by exploiting higher computing power, CFD needs to become more reliable, more reproducible across users, and better understood and integrated with other disciplines and engineering processes. Uncertainty quantification is universally considered as a major goal, but will be slow to take hold. The prospects are good for steady problems with Reynolds-Averaged Navier-Stokes (RANS) turbulence modelling to be solved accurately and without user intervention within a decade – even for very complex geometries, provided technologies, such as solution adaptation are matured for large three-dimensional problems. On the other hand, current projections for supercomputers show a future rate of growth only half of the rate enjoyed from the 1990s to 2013; true exaflop performance is not close. This will delay pure Large-Eddy Simulation (LES) for aerospace applications with their high Reynolds numbers, but hybrid RANS-LES approaches have great potential. Our expectations for a breakthrough in turbulence, whether within traditional modelling or LES, are low and as a result off-design flow physics including separation will continue to pose a substantial challenge, as will laminar-turbulent transition. We also advocate for much improved user interfaces, providing instant access to rich numerical and physical information as well as warnings over solution quality, and thus naturally training the user.

Keywords: CFD; aerodynamics; numerical methods; turbulence modeling

This is an invited paper to mark the 150th anniversary of the founding of the Royal Aeronautical Society in January 1866.

NOMENCLATURE

APU   auxiliary power unit
CFD   computational fluid dynamics
CRM   common research model
DES   detached-eddy simulation
DOF   degree of freedom
DNS   direct numerical simulation
ECS   environmental control systems
FMC   flight management computer
GPGPU   General-Purpose Graphics Processing Unit
IBL   integral boundary layer
LES   large-eddy simulation
MLA   manoeuvre load alleviation
ODE   ordinary differential equation
PDE   partial differential equation
RANS   Reynolds-averaged Navier-Stokes
S&C   stability and control
UQ   Uncertainty Quantification
WMLES   Wall-modeled LES

1.0 INTRODUCTION

Computational Fluid Dynamics (CFD) has become instrumental in the design and analysis of products in the aerospace industry as well as in surface transportation industries including automobiles, trucks, and boats. This paper is written from the perspective of (external) aerodynamics, although many of the issues identified carry over to the other application areas. Turbomachinery is covered elsewhere in this special issue of the journal. It is fair to write that CFD has had a tremendous although gradual impact on both commercial and military aircraft. The use of CFD in commercial aircraft is well documented[1,2], with particular success in the design of the high-speed wing (cruise shape) and its close integration with the engine, dating back to the Boeing 737 Classic of the late 1980s. The extensive use of CFD in the latest aircraft from Boeing is illustrated in the 'walk-around' chart depicted in Fig. 1; Airbus has presented a similar chart. It is seen that in certain areas the use of CFD is at a mature state but there are also emerging areas where CFD is expected to affect aircraft design in significant ways only in the future. The role played by CFD in helicopter design is documented in a 2007 paper by Strawn et al[3]; we have failed to find more recent overviews, although there is vibrant activity on subproblems such as a rotor in hover. Since the helicopter is even more complex than the high-lift configuration of an aircraft, it is no surprise that comprehensive CFD treatments are not in routine use. In the area of fighter aircraft and the many other military systems, there is understandably less public documentation. Nevertheless, the aerodynamic problems are similar if with different emphases. In the area of automobile design, CFD plays a crucial role in determining optimal aerodynamic shapes[4], prediction of aerodynamic forces[5], as well as in understanding the sources of noise generated by protuberances. CFD also contributes to the design of wind turbines and wind farms[6], while its role in ships is discussed in detail in Ref. 7.

    CFD is increasingly being used in multi-disciplinary design and analysis of aerospace products. Examples of these include high-speed aerodynamic design taking into account the flexibility of wings (aeroelastics), icing models, far-field noise propagation models and
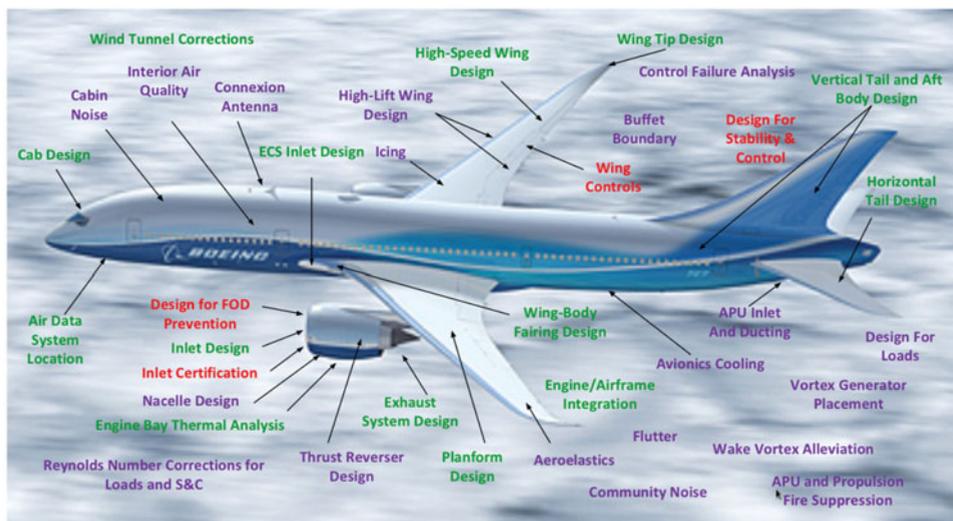
Figure 1. Impact of CFD at Boeing. Green areas have strong CFD penetration;
blue areas have some penetration; red areas present future opportunities.

conjugate heat transfer. An attractive development is the use of a range of CFD tools to calculate the benefits of formation flight for large aircrafts[8]. Increasingly, CFD results are compared directly with flight test, rather than wind tunnel, and the status of the two sources of information in the engineering process and company culture is slowly shifting, with enlightened organisations drawing on both to good effect. It is important to transition from wind tunnel to CFD for the right reasons, such as wall effects or Reynolds number and aeroelastics, whereas doing so only for speed and cost advantages has its dangers. We believe there is a tendency towards overconfidence in CFD in some circles, even to the extent of ignoring well-known sources of error, which creates a risk of backlash, were CFD to be blamed for costly mistakes.

CFD still faces several challenges that need to be addressed. The turnaround time associated with CFD is one of the factors limiting the use of CFD in the design and creation of databases and also in multi-disciplinary applications. Another limiting factor is the level of skills required of the user of CFD. CFD practiced in industry is vastly different from the CFD theory taught in universities, especially in the late 20th century. A long lead time can be required for a user of CFD to become proficient in all the various phases of CFD (geometry preparation, gridding, solution set-up, post-processing). Other limitations include various uncertainties in CFD related to numerics, physical modelling (especially transition and turbulence), and the time involved in preparing geometries for carrying out grid generation and aerodynamic analyses. The latter two tasks are still highly manual and in many instances dominate in terms of effort, compared to the solution of the fluid-dynamic equations.

This paper lays out the present and future roles we see for CFD in the design and analysis of commercial aircraft as well as the detailed prediction of their performance, and at some point their certification (in the case of military aircraft, the word certification is replaced by qualification). Boeing and its competitors are very conservative companies, first of all because of their passion for safety, but also because of the extreme industrial consequences of any design mistakes. Flaws uncovered during assembly or flight test of a new model cause considerable disruptions for the entry into service. The corresponding financial impacts are

very large, and the possibility that the new aircraft model would be impossible to certify short of, say, a complete redesign of the wing would be a nightmare. As a result, the penetration of CFD is gradual, often involving agreement amongst large communities, from engineers to top managers to company pilots, and acceptance by government agencies such as the Federal Aviation Administration (FAA).

This paper may be viewed as attempting to provide a wider update to the vision spelled out in the paper by Spalart and Bogue for off-design studies[9] and a somewhat narrower, more technical, and at times more critical view of the future than is spelled out in the well-researched Vision 2030 document[10]. Our thesis is laid out in Sections 2, 3, and 4. In Section 2, we examine the role played by CFD in various aspects of the aircraft. These include high-speed (cruise) wing design, low-speed (high-lift system) analysis, internal flows, stability and control (S&C), vibration, and noise prediction. We also make projections regarding the role we envision for CFD in the 2040 time frame under realistic assumptions. Section 3 is devoted to examining the principal challenges that we see in improving the quality and reliability of CFD. In this section, we discuss issues related to geometric modelling, computing-power limitations, numerical accuracy, physical modelling, and interaction with other disciplines. We argue that there is much room for improvement, and we categorically state that CFD is far from being a 'solved problem' or even one that would be resolved by unlimited computing power. Section 4 lays out our ideas for user awareness and education as they relate to CFD. We believe these ideas will produce better-informed and more careful CFD practitioners. They will also enable better communication of results from CFD, help non-CFD experts relate more easily to them, and more importantly trust the results to the degree warranted. We finally offer our projections and prospects for CFD in the Conclusions section.

2.0 PRESENT AND FUTURE ROLES OF CFD FOR ANALYSIS, DESIGN, AND CERTIFICATION

There are two primary ways in which CFD is used in the aerospace industry. The predominant use is in the analysis phase. Given a geometry definition, flow conditions, and appropriate boundary conditions, the task is to compute the flow field, with sufficient accuracy in the region close to the aircraft (wake-vortex applications require special provisions to maintain the accuracy much farther downstream). An appropriate physical model is used. This can and does span the gamut from lifting line and vortex lattice methods to panel methods (potential flow) and Euler equations coupled with Integral Boundary-Layer (IBL) formulations, Reynolds-Averaged Navier-Stokes (RANS) formulations which require turbulence to be modelled, to Direct Numerical Simulation (DNS) methods in which all the scales of turbulent motion are captured. It is essential for users to choose the most appropriate level of sophistication and cost for each application. Except for panel methods that are classified as boundary integral methods, all the other formulations require a grid to be generated that fills the space occupied by the fluid. And herein there are multiple choices as well: single-structured grid for simple topologies, multi-block structured, overset and unstructured. We believe automatic grid adaptation, or 'self-gridding,' is a very powerful ingredient of CFD; however, it has proven very difficult, and even the talent in government, industry, and academia and the competition amongst CFD code suppliers have had only modest levels of success.

The Partial Differential Equations (PDEs) governing fluid flow are then discretised using any of a variety of methods: finite volume, finite elements (continuous and discontinuous) with many choices available for numerical flux approximations. That such a variety of methods

exists in CFD, with no clear winners and losers, after so many years of research is surprising indeed; this is not the case in other fields such as structural mechanics. At this point, one also settles on the order of accuracy desired, second order being the most common, although methods possessing higher-order accuracy are becoming more practical. After discretising the time derivative, the last aspect is to solve the non-linear system of equations to obtain a flow field at every time step or the steady flow field at all the discrete locations (grid points). Once the flow field is computed, one is able to extract global quantities of interest, such as forces and moments as well as local flow-field characteristics such as skin friction, velocity and temperature profiles, surface pressures, entropy, and total pressure. A flow field, even steady, computed with CFD is rich in information. Typically, however, much of it is ignored in favor of near-field quantities such as surface pressures, and force and moment coefficients.

Regardless of the level of approximation chosen for the flow model, when a reliable solver capability has been achieved, it is possible to automate many of the phases of CFD. Much of the internal work in industry consists of building and validating tools built around solvers from various sources. This is particularly true with respect to geometry and grid generation if geometric complexity is restricted (e.g., cruise wing, body, nacelle, horizontal and vertical tail in the case of commercial aircraft). With sufficient speed and automation and some established level of confidence in CFD, it is then possible to generate an entire aerodynamic database, for certification and possibly to drive flight simulators. This is already being done to a degree with full-potential and Euler equations.

At present, CFD and wind tunnel are used in a complementary fashion. The initial cost of CFD can be substantially lower than the initial cost of a wind-tunnel test (model fabrication, installation, and so on) but the cost comparison switches in favor of the wind tunnel when hundreds of conditions are needed such as drag polars over a Mach-number range (and possibly yaw angles, control-surface positions, and the like).This is separate from the accuracy issues, which we discuss at length in Section 3 for CFD. In the wind tunnel, Reynolds number limitations (requiring scaling to flight conditions), tunnel-wall effects, and the inability to include aeroelastic effects very accurately are dominant (although knowing the exact shape for CFD purposes is not easy either). We believe with sufficient investment, many of the shortcomings associated with CFD can be addressed providing sufficient confidence in CFD to enable entire aerodynamic databases to be generated. Venturing further, it may be possible to 'fly the Navier-Stokes equations' someday, if computers and algorithms get powerful enough to accurately calculate the flow over the aircraft in real time even during manoeuvres far from the design point[9].

Once a reliable analysis capability is in hand, it becomes possible to use it in the context of product design. This provides a crucial advantage over wind-tunnel design work, for which even a subtle twist or aerofoil change requires the fabrication of another wing. For commercial applications, at the cruise condition, a 1% difference in lift/drag ratio is very significant. As a first step away from straightforward analysis, initially CFD was used to produce the shapes that would achieve target pressure distributions on wings; this is referred to as inverse design. However, this approach requires that favourable pressure distributions be known in advance, with no assurance that this could be achieved through changes in design parameters, especially in transonic flow where the Hugoniot relations are unavoidable (since smooth shapes are needed, so that the shocks are normal to the wall). A more modern and better approach is to cast the design problem as an optimisation problem with constraints. One defines an objective function (e.g., drag coefficient) and sets about minimising this objective function with respect to the design variables (shapes, freestream Mach number, angle-of-attack etc.) subject to the flow equations being satisfied as well as other geometric and flow constraints (lift coefficient).

Also, effective optimisation approaches use 'multi-point' design to avoid the pitfall of creating a point design, which would perform very well in only one condition, and develop problems such as strong shocks and separation at other conditions[11]. The gradient-based optimisation algorithm requires the formation of the gradient (dI/dD where I is the objective function and D is a design variable). There are multiple ways of computing the gradient, for example, via finite differences, direct method (which requires the inversion of the large Jacobian matrix dR/du where R is the residual and u is the vector of unknowns), and the adjoint method. The latter is particularly efficient with the effort independent of the number of design variables. The gradient-based approach only computes a local minimum. Global optimisation techniques typically do not use gradient information. Instead, they form a response surface from the seeds and selectively refine this surface in regions of interest. Global optimisation techniques are limited to handling only a small number of variables because the number of analysis runs required varies as $O(N_D^2)$ where N_D is the number of design variables. There are some approaches that combine elements of the global and local optimisation techniques.

Potential new areas for CFD to contribute are in the certification of various phases of an aircraft development, and in airline customer support. In addition to communicating to authorities the abilities of CFD to produce results that can be trusted for particular needs, it is imperative that CFD processes become traceable and repeatable. Efforts are already underway for certification using CFD in specific instances, such as small changes in configuration. In fact, 'certification by analysis' refers to establishing information via theory, comparisons amongst aircraft models, and ground tests, and not only by CFD, with a view to making some flight tests unnecessary. The principal motivations for this are cost, schedule, and avoidance of danger during flight testing. Concerted efforts are needed if much of the database in the flight simulator is to be populated using CFD. The level of uncertainty associated with CFD needs to be quantified and deemed to be acceptable by the technical experts in the particular application area before approaching the certification authorities, or even the company test pilots. This uncertainty will be compared with the inherent noise in flight-test measurements. Wind tunnels have their errors, but they have been very stable (only the few cryogenic facilities that can achieve full-scale Reynolds numbers for large commercial transports are relatively new), so that the extrapolation from wind tunnel to flight is built on a considerable knowledge base. This is hardly the case for a CFD capability, which is rapidly evolving from year to year.

We now discuss different areas of CFD applications in more detail.

2.1 High-speed flight conditions

CFD is relied upon heavily in the design of the high-speed wing (cruise). It is a particular challenge to design the wing in the presence of nacelle and engine, and the Boeing 737 in its several generations has greatly benefited from this capability. The ever-higher by-pass ratios and fan sizes in recent models and derivatives also benefit from it. CFD has evolved to a point that wings can be designed for optimal performance taking into account the interaction with the nacelle and the engine, and with acceptable off-design behaviour through multi-point design. It is now the case that wings are designed using CFD, and confirmed in the wind tunnel, and that very few wings are tested. In the case of Boeing, one of the primary design tools is the TRANAIR optimisation tool[12], which features a full-potential formulation coupled with an Integral Boundary-Layer (IBL) method, solution-adaptive gridding, and shape changes effected via changes in transpiration velocities. The latter feature eliminates the need to generate body-fitted grids around changing geometries, which is a major shortcoming when using RANS codes. The level of confidence in CFD is very high in and around cruise
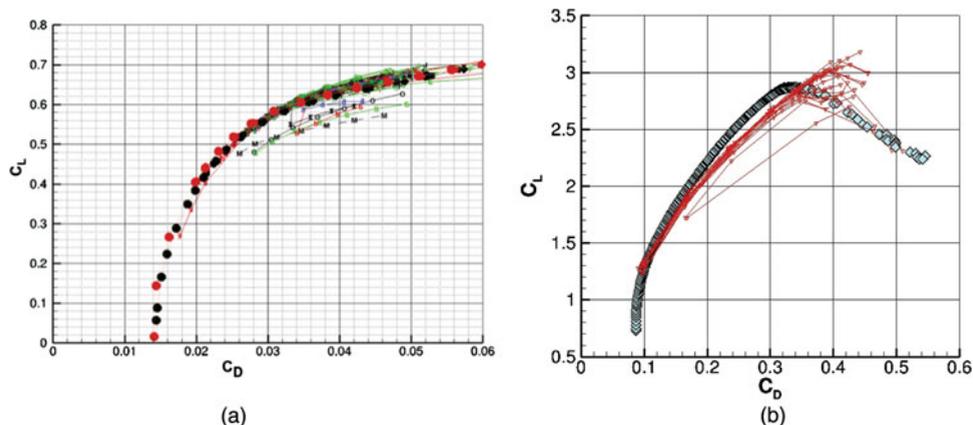
Figures 2 (a) and 2(b). Lift-drag polars at two CFD workshops. Left: high-speed, clean-wing configuration, workshop DPW4[13]. Bullets: experiments in two NASA wind tunnels; lines: CFD results from various authors. Right: high-lift configuration, workshop HiLift2[14]. Diamonds: experiment; triangles: CFD results from various authors.

conditions. We acknowledge that transport aircraft configurations near cruise conditions are special with thin boundary layers and essentially irrotational flow elsewhere (outside of shocks and wakes), which permit the full-potential formulation, coupled with IBL, to do surprisingly well. Other components, such as fairings and aft bodies, are designed using Navier-Stokes solvers. In the full-flight regime outside of the cruise range, there is less confidence in CFD. For general configurations with more complex flow physics or for off-design conditions and especially buffet, at least the Euler with IBL or RANS level of modelling is needed, of course without any claim of perfection. Still, Fig. 2(a) presented at a high-speed workshop for a representative commercial aircraft wing suggests that a high level of confidence has been established, although still not fine enough for drag prediction to rely only on CFD. Recall the importance of a 1% change in drag in the commercial aircraft business. The figure shows that many codes give very similar and accurate answers, and a few are 'outliers.' We elected to preserve those in the figure and thus render the full state of the art, knowing that in industrial use, these codes would not be trusted.

2.2 High-lift configurations

The level of confidence in CFD when dealing with flow past complex configurations such as high lift (with leading-edge and trailing-edge devices deployed) is considerably less compared to the high-speed clean-wing area. This is illustrated in Fig. 2(b). The figure again has a cluster of similar results, and a minority of outliers. However, even for the cluster of more-accurate curves, the CFD drag is somewhat too high in the intermediate range, and conversely the maximum lift is too high and delayed to too high an angle-of-attack. Experimental comparisons demand more care at high-lift coefficients, for instance 3 instead of 0.7, in particular because of stronger wall interference. We find it unfortunate that the general practice is to run CFD 'in free air' at an angle-of-attack corrected from the wind-tunnel value. It would be best to run CFD with solid walls, even if treated as inviscid, without corrections. We recognise that this is more difficult for transonic tunnels with slotted walls. All three of lift, drag, and pitching moment would be more representative of the true predictions of CFD. This presentation also ignores the presence of multiple solutions for this very geometry, discussed

in Section 3.3. In addition to the obvious difference in complexity of the geometry, a strong contributing factor is that for the clean wing, the viscous wake is thin and (in normal attitudes) does not interact with other airframe parts. In potential-flow solutions, there is little penalty in using a flat vortex sheet and neglecting its roll-up. In high lift, there is separation from the slat lip and flap edges, and thick viscous layers and vortices directly interact with the lifting surfaces so that their correct response to inviscid and viscous (turbulent) effects becomes crucial. The increase in grid count needed to model these features well is then much more than the order of magnitude, which would appear justified by the geometric complexity in itself.

As a result, in contrast with the trend we mentioned for clean wings in steady flight, in our industry even now the wind-tunnel effort devoted to the high-lift system, S&C, loads, failure conditions, icing, and airframe noise is considerable. Altogether, such testing takes of the order of one-year round-the-clock occupancy for a new wing, including its high-lift and control systems. This reflects the countless combinations of configurations and attitudes, of course, but it also reveals the relative weakness of CFD when dealing with the full complexity of an aircraft. This weakness includes accuracy concerns and turnaround time. The flow physics are considerably more complex with viscous effects being dominant (interactions amongst multiple boundary layers, shear layers, wakes, and sometimes shocks). As a rule, the flow is almost always separated in some regions and most nettlesome of all, the flows are characterised by smooth-body separation. The geometric complexity is high requiring upwards of 50-100 m grid points for a fixed grid in current practice. Even such grid counts are far from allowing precise resolution of the many shear layers and vortices that dominate the flow field and whose positions are not known in advance. As of today, both lack of grid resolution and physical modelling errors appear to be limiting factors, of comparable impact. Still, CFD is used in a limited fashion in high lift to weed out configurations, obtain force increments due to configuration changes by carefully controlling parameters such as grid topology, and to thin out wind-tunnel test points by anchoring a few CFD solutions to wind-tunnel data. Current practice is to use steady RANS on affordable fixed grids and apply engineering judgment, which consists of acceptable convergence of force coefficients and residuals as well as the use of flow visualisation techniques to determine whether the solution is 'trustworthy.' Solution adaptation is clearly needed to capture more of the flow physics and still keep the problem sizes manageable. The issues associated with physical modelling are covered in the Accuracy of Physical Modelling section, while the particular numerical issues germane to high lift are covered in the Accuracy of Numerical Solution section.

Another rapidly growing and similar area of application is active flow control, made very challenging numerically by the very small length and time scales of the actuators, compared with those of the full airframe, and physically by the intense three-dimensional turbulence involved, for which RANS modelling is questionable. CFD is of course coordinated with wind-tunnel and flight tests.

2.3 Internal flows

CFD is heavily used in internal flows, which are not unique to the aerospace field. The geometries here tend to be very complex (e.g., thrust reversers) made up of hundreds or even thousands of surfaces. Grid generation of such complex geometries is a tedious task, and the chances that every important region is well resolved are slim. Internal flows also tend to feature special boundary conditions, such as mass flux matching, bleed boundary conditions, radiation boundary conditions etc. Multiple chemically reacting species are also modelled. Other uses of CFD are in Environmental Control Systems (ECS) where cabins are modelled

for improving passenger comfort in terms of air temperature, draft, and noise, and to study the effects of pollutants. The ECS duct systems are very complex, and CFD is only slowly moving towards a complete treatment of them. The priorities include a good distribution of the air flow, and low noise. Another use of CFD is in the area of fire suppression where fuel tanks and spaces such as engine nacelles and APU housings are modelled. Many of these applications also deal with multi-phase flows.

2.4 S&C, loads, and load alleviation

S&C applications add a moderate amount of difficulty to the accurate simulation of the flow, at least for normal control-surface deflections and flight attitudes, but they add a considerable amount of volume to the database needed. Even approximating the flow as steady is not always correct. In some cases, coupling with control systems is necessary. A classic and benign example is the yaw damper, but modern Flight Management Computers (FMC) are far more complex, and operate on shorter time scales. This is the case for commercial aircraft, and of course even much more for fighter planes with relaxed static stability. Modern, light airframes are also visibly more flexible than older ones, which brings their natural frequencies close to the frequencies of the FMC and human pilot, of the order of a few hertz, and creates new possibilities for detrimental interactions such as aircraft-pilot coupling. Some models have Maneuver Load Alleviation (MLA), which is a prime example of coupling between aerodynamics and systems with crucial implications to the integrity of the aircraft. A comprehensive simulation approach to MLA based on unsteady CFD tightly coupled with control laws is not available, and we would not expect it to be until at least 2020. Gust load alleviation is also of value for passenger comfort. Outboard aileron reversal is an old problem, but as present as ever, and differences on lateral control systems (aileron types) by different companies suggest that the optimal solution may not have been found even for completely new wings.

The calculation of loads on every part of the aircraft in every configuration and flight condition is a very large task, and there is great value in high-accuracy predictions early in a program, in order to size the structure and anticipate the exact mass of the aircraft. In recent programs, CFD has by no means been the sole source of that information, partly for reasons of confidence but primarily because of the size of the database. In this domain, the competitive pressures leave little room for conservatism, but errors are very damaging. This would be true even if the error concerned the sizing of an aileron actuator, for example. As a result, the industry is increasing its reliance on CFD prudently, collecting all the possible lessons from each new program.

A fascinating example of S&C progress is the provision of an 'electronic tail skid' to prevent tail strikes on commercial aircrafts during takeoff or landing. Just like the flare manoeuvre mentioned above, it involves rapidly changing attitudes, ground effect, high-lift systems, control surfaces, and the FMC in a quite complex manner. Again, we expect CFD to be an integral part of its prediction, but not to be the sole source of data for years.

2.5 Noise and vibrations

The contributions of CFD to noise prediction, whether community or cabin noise, lag far behind its contributions to aerodynamics, and at best this is a nascent field. However, there is deep potential, especially as we enter the LES era. Numerous methods have been proposed to use Reynolds-averaged turbulence quantities such as turbulent kinetic energy k and dissipation rate ε to build noise-source models, but in our opinion they have not delivered practical

validated tools. Instead, we support an evolution towards the use of first principles, based on unsteady simulations. These have been validated, particularly for jet noise, but almost exclusively on simple geometries. Performing LES with high-order low-dissipation numerics around the complex geometry of an installed jet engine with a supersonic stream remains very difficult, even ignoring the sources of noise linked to the moving blades of the fan and turbine. The non-uniform surrounding flow adds to the difficulty when it comes to predicting far-field noise radiation. However, the research studies on simple geometries proved that LES, combined with noise-radiation post-processing of the Ffowcs-Williams and Hawkings type, accurately captures the effects of dual streams, temperature and Mach-number variations, and shock cells. Airframe noise predictions, for instance from landing gear, are not as well validated, partly because of the difficulties in closing the integration surfaces downstream. Controversies over low-Mach-number and similar approximations linger. However, both for landing gear and high-lift devices, LES is quite successful at reproducing the turbulence itself and in particular, the wall pressure fluctuations[15]. Its potential is certain.

An intermediate application is to shock cells, which are important to cabin noise; in that case, accurate enough CFD of the installed engine is a tool to obtain the shock-cell pattern and strength, which can be used qualitatively or empirically to predict the noise. Still, the accurate prediction of a long train of shock cells in cruise flight is very challenging in terms of grid quality and solver performance, and in addition, the RANS turbulence models disturb the flow field and are not well understood in this arena as of now. Another area of partial success is in using CFD to reduce flow separation, which is known to create noise. However, the prediction of the noise effect itself is still qualitative as of today. For instance, since 2004 the Boeing 737 has carried pairs of small vortex generators, easy to identify on its nose while waiting at the airport, intended to reduce separation at the base of the windshield. They were designed by CFD and no aerodynamic tests were conducted, but the noise reduction at the pilot's ear was measured in flight[16].

Vibrations and sonic fatigue are also obvious candidates for unsteady simulations. There are frequent applications to poorly streamlined components such as temperature probes, drains, or windshield wipers, and also to appendages such as turrets on military derivatives of commercial aircrafts, or blisters over antennas on commercial planes. Cavities needed for landing gear or ordnance have been a widespread and often successful application. The simulations may be of unsteady RANS type, in which case there is not firm guarantee that the simulation will correctly predict the onset of unsteadiness, particularly if the appendage is somewhat streamlined, or is immersed in the boundary layer. Simulations of detached-eddy simulation (DES) type are always unsteady, and have been highly successful for instance around tandem cylinders for research and around many aircraft components we cannot describe here; however, they are more expensive and demanding of user skill. Industrial use will grow with the rise in computing power and the development of best practices.

3.0 PRINCIPAL CHALLENGES TO QUALITY AND RELIABILITY IN CFD

As engineers, we have knowledge of the progression of CFD at The Boeing Company and a sense of its bright future. As developers on the flow physics and the numerical sides, we are keenly aware of its imperfections and heavily invested in limiting the dangers of overconfidence and simplifications, as well as planning and applying research and development efforts in the most effective manner.

3.1 Accuracy of geometry

Geometry is obviously one of the essential inputs to CFD and is probably one of the most difficult aspects in the entire CFD process. Defining geometry for manufacturing is hard enough (e.g., using Computer Aided Three-Dimensional Interactive Application (CATIA)), but in CFD what is desired is as simple as possible a definition of the geometry that can be analysed aerodynamically. When the geometry is imported from external sources, the number of surfaces could be in the thousands. The geometry may contain gaps, multiple definitions, intersecting surfaces etc. that have to be resolved unambiguously and uniquely. In the context of fixed grids, it is possible to make such decisions on the fly during grid generation. There are programs available that try to automate much of this but still require tolerance specifications that may have to be varied depending on the geometries. The problem of clean geometry specification is more pressing and delicate when adaptive gridding techniques are employed. Gaps that were resolved during grid generation are unacceptable, as are trimmed curves that possess different parametric representations on two intersecting surfaces. This problem is currently being addressed by using surrogate geometries (e.g., piecewise linear geometry definition) in the vicinities of gaps and intersections.

3.2 Surface and volume gridding

When a clean geometry is finally created, surface gridding (sometimes also called paneling) is the next task. Surface gridding takes into account the curvature of the surfaces using measures such as chord-height and chord-height/chord, smooth variations of the surface grid, desired spacings in regions of great solution variation (wing trailing edges, leading edges, nacelle lips etc.). Every surface grid generator uses its own set of rules. Clustering of surface grid points near singularities (e.g., trailing edges) is considered important but there is no accepted set of rules to accomplish it. Surface gridding in such areas also has a profound impact on the volume grid. In principle, the rules for generating surface grids should not matter because as part of grid refinement, the surface will also get refined. In practice, however, highly irregular grids present challenges to many of the flow solvers. So, much time is spent in creating smooth, graded meshes. It should be noted, however, that grid refinement is seldom done. And even if done, uniform grid refinement is inadequate. Therefore, we believe surface grid generation is another source of error in CFD simulations that should be better controlled. For most CFD codes possessing second-order accuracy, linear or bilinear representation of surface grid will suffice. However, when higher-order methods are used for discretising the flow equations, the surface grids also need to have similar higher-order properties to approximate general curved geometries.

Once a surface grid has been created, a number of grid-generation techniques can be used for the volume grid. Some of the choices are multi-block structured grids employing elliptic or hyperbolic techniques, overset methods that generate simple grids around components with overset regions, unstructured grids using advancing front and/or Delaunay tetrahedralisation techniques, oct-tree grids, or a combination of all these techniques. Typically, for turbulent simulations, one can estimate the initial normal spacing based on the spacing in viscous wall units y^+, which can be estimated based on the flow conditions. Geometric stretching is employed in the normal direction and is a key parameter (typically 1.1-1.3). In particular, contrary to common expectations, reducing the first y^+ and holding the stretching ratio the same does not refine the grid away from the wall by the same proportion; far from it. Few users of CFD understand this. Wakes are resolved to a degree by choosing the 'right' grid topology (e.g., C-mesh) or in some instances wake sheets are prescribed as additional 'surfaces' to

concentrate the grid points. It should be mentioned that most of the flow features, such as wakes, shear layers, and shocks are seldom known in advance, which again underscores the value of adaptation. Grid adaptation for turbulence-resolving approaches such as DES brings additional challenges, starting with the choice between static adaptation and time-dependent adaptation. This is already the case for URANS studies; for instance, in a buffet simulation, the shock wave has a range of positions during a cycle. The expectation is that with grid refinement, such features will get better resolved along with using smaller geometric stretching factor in the normal direction. However, non-adaptive refinement is very expensive, especially in three dimensions.

3.3 Accuracy of numerical solution

There are many sources of error associated with numerics in CFD even after an appropriate physical model is chosen. Issues related to improving the robustness of CFD codes are discussed in depth in Refs 17 and 18. Here, we merely touch upon a couple of major issues that we see as most important. Ideally when numerical results are presented they should be a close approximation to the solution of the set of continuous PDEs, at least in the near field. Limiting ourselves to steady problems, the sources of error are iteration or convergence error and discretisation error (associated with the use of a finite grid). Considering that problems of ever-increasing complexity are being solved, it is lamentable that an often-asked question is simply, 'Is the solution converged?' The answer varies widely across codes and flow conditions, and is seldom satisfactory. Much effort is spent on deriving criteria for declaring convergence of CFD codes (e.g., five orders of reduction in residual norm for steady-state calculations and/or 0.1% change in force coefficients). However, even such arbitrary criteria are seldom met, especially at difficult flow conditions. Turbulence model equations are notoriously difficult to solve because of stiff source terms. Their convergence usually leaves a lot to be desired. There are no consistent measures for measuring and gauging convergence across CFD codes either. For unsteady flows, when implicit methods are used, a source of error is also the degree to which the non-linear system is solved at each time step. Based on work done in stiff-ODE (Ordinary Differential Equation) solvers, it may be possible to bound such errors.

The other main source of error is related to the use of a finite grid. Grid resolution studies should be and have been performed to reduce the discretisation errors, obtain grid-converged results, and increase confidence in CFD solutions. Such studies are seldom done except for benchmark test cases (AIAA DPW 1–5[13], high-lift prediction workshops 1 and 2[8]). Many of these studies are still inconclusive in some respects, partly because they have amounted to uniform refinement starting with the initial grid, which may not be the best route. Mavriplis[19] shows that despite obtaining what could be considered grid convergence, the answers changed substantially with a different sequence of uniformly refined grids. A recent paper by Diskin et al[20] concludes that even in two dimensions, the only practical way to establish grid convergence for meaningful geometries is through the use of solution-adaptive grids. Obvious reasons for this are the presence of shock waves and geometry singularities such as trailing edges. We agree with this sobering conclusion and hope that grid-converged solutions using adaptive gridding will be demonstrated for drag prediction, high-lift and propulsion workshop test cases in the near future. Adaptive gridding technology has been developed for full-potential flow (TRANAIR[21]), Euler equations (CART3D[22]), and two-dimensional RANS[18,23]. The codes cited above use solution-adaptive gridding as a matter of course to compute CFD solutions and can be considered to represent the state of the art. For three-dimensional RANS, adaptive gridding has not become routine by any standard and
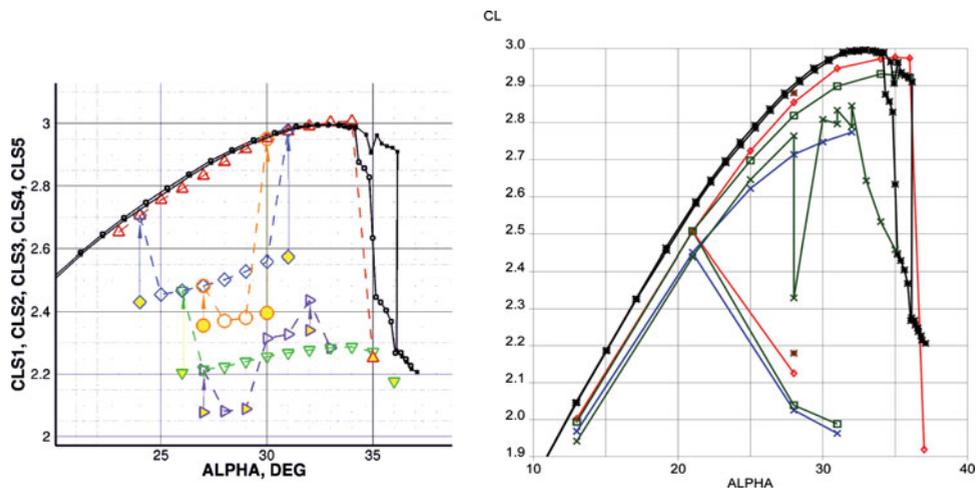
Figure 3. Lift coefficient of multiple solutions obtained for the trapezoidal wing. Black lines are experiments. Left, 3(a): results from single code on the same grid; dashed lines are CFD. Right, 3(b): colour-coded results from four different codes, one of them with two turbulence models.
(Courtesy of D. Young.)

is an active research topic. Until adaptive gridding does become routine, the authors feel that providing results on two grids, one of which is a uniform refinement (say by a factor of 1.5-2 in each direction) of the other grid would go a long way in increasing confidence in CFD results, by weeding out poor solutions.

3.4 Multiple solutions

Another serious issue that has been uncovered recently has to do with multiple steady solutions in CFD[24]. Given the flow conditions, several residual-converged (to machine zero) solutions could be obtained depending on the path taken to obtain convergence; for instance, a start from freestream conditions versus a start from a solution at a lower angle-of-attack (note that this procedure is unphysical, and does not amount to reality in the wind tunnel). This is exhibited in Fig. 3(a), in which a single grid of intermediate density, namely 11m nodes, was used. Finer grids were not found to have as many solution branches, but it is not known whether that is a definite trend, or simply due to the exploration of the solution space not being as extensive on finer and therefore more expensive grids. A recent and important finding shown in Fig. 3(b) is that at least the upper and one lower branch have now been confirmed by other codes with very different grid systems (including structured and unstructured) and algorithms, including the fact that the width of the branches is much larger than the width of wind-tunnel hysteresis loops obtained with slow increases and decreases of the angle-of-attack (this figure includes unpublished work by J Bussoletti, D Williams, and D Kamenetskiy at Boeing). Multiple solutions have been encountered in a range of situations, from simple extruded wings to complex high-lift configurations, and seem to be almost always associated with smooth-body separation (although transonic multiple solutions are not unheard of e.g. Ref. 11).

Mathematically speaking, the existence of multiple solutions for a non-linear system of equations is well known. Additional causes for this include the far-field boundary conditions, which are approximate and provide an unlimited supply of energy, and the presence of turbulence models. These models are 'creations of the mind' aimed at working solutions for

the system of Reynolds-averaged equations, which itself is incomplete and physically can be viewed as quite artificial[25]. There are few if any known mathematical facts about uniqueness (for instance, entropy principles), or even stability; it is easy to create a plausible model that has immediate Hadamard singularities, even by only altering the value of a constant. The models in common use have demonstrated their robustness in practice, and the full agreement between different CFD codes using them is reassuring even as it has been secured only for a small class of 'academic' cases. As a result, we do not view the turbulence models as the primary cause of non-uniqueness, and venture that laminar problems could encounter similar phenomena.

In reality, hysteresis is a physical phenomenon that features multiple possible states typically in a wind tunnel, for the same controlled conditions. The multiple solutions observed in Ref. 24, however, occur over too wide a range of angle-of-attack to correspond to the exper-imental observations to a convincing degree, and we believe some of them are purely numer-ical. The difficulty that confronts CFD is somewhat existential: Even if one obtains a residual-converged solution, to what extent does it reflect reality? Efforts are under way to uncover the various solutions by systematic means such as deflation[26]. It is also possible to analyse the stability of the obtained solutions under time evolution. A silver lining may be that with grid refinement (perhaps solution adaptive) most of the spurious solutions will cease to exist, leaving only one or two solutions that are physical, or 'more physical' in the sense of being affected only by turbulence-modelling errors. This is hinted at in the results shown in Ref. 24.

3.5 Slowing growth of computing power

CFD has benefited tremendously from the explosive growth in computing power over the last 30 years or so. Simulations that required supercomputers such as a Cray YMP in the 1980s can now routinely be carried out on laptops at a fraction of the cost. CFD has also benefited considerably from algorithmic improvements, such as algebraic multi-grid and pre-conditioned Krylov methods. Most of the algorithms are able to exploit parallelism available in present-day distributed-memory parallel computers through domain decomposition. MPI and to a lesser degree OpenMP are programming models that are used to exploit such architectures, with some success. The challenge for CFD is how to adapt to newly emerging architectures such as coprocessors and GPGPUs[27]. It is fair to say that apart from a few simple algorithms that have been implemented on these more difficult architectures, much of the CFD community is adopting a wait-and-see attitude. The accepted rate of growth for a single chip has long been Moore's law, or roughly a factor of 2 every two years (if not 18 months). The rate of growth of the fastest supercomputer has been much faster, of the order of a factor 3.8 every two years, due to the increase in the number of chips, and therefore cost and electrical power. CFD has benefited from this to a considerable degree. Past projections for the treatment of turbulence by DNS and LES have relied on Moore's law[28]. However, Moore's law encountered an apparent ceiling for chips around 2012. In terms of supercomputers, the Tianhe-2 machine, which has held the worldwide top position since 2013, is rated at 30 petaflops, and the recent U.S. initiative aims at the exaflop speed in 2025. The ratio between the two represents a factor of only 1.8 every 2 years, which is 57% slower than the factor of 3.8. In addition, the full success of the initiative is not granted, considering the issues of cost and electrical power, and by some estimates, the future machine may be used only at an efficiency of a few percentage points. In other words, it may achieve exaflop performance for LINPACK cases, but be far below that in CFD practice. This is barring a breakthrough, of course, originating in quantum computing or another concept. The marked

weakening of the computing-power rate of rise is likely to be durable, and to disappoint the common expectations of 'free' progress in CFD.

3.6 Accuracy of physical modelling

The ubiquitous challenge in transportation CFD is turbulence, with transition equally challenging but in question only in relatively small regions and in special applications. The treatment of turbulence has been vexing for a century and will not be a 'solved problem' for decades. Progress will be dependent on new ideas and their correct implementation, and on computing power. To illustrate the wide range of turbulence problems, as of 2010 DNS, which is a treatment purely from first principles of a non-rotating golf ball, was possible with about 1 Bn grid points and more than 100 grid points per dimple diameter, and gave 'reasonable agreement with measurements' in these authors' own prudent assessment[29]; note that the experimental measurements had a 20% spread. We suspect DNS is now definitely accurate even with rotation, the companies engaged in it not making their results public, and that DNS will actually lead to better dimple designs. However, if we stay with sports, DNS of even a tennis ball is not possible, if only because of the hair on it, which could not be resolved numerically. Treating it requires a model of rough-wall turbulence, and therefore empiricism. At the other end of the scale of flows of interest, there would be great value in predicting hurricanes days in advance, and the Reynolds numbers there far exceed even those in aerospace systems. Again, these simulations are dependent on modelling physics empirically on a relatively very small scale; for instance, the interaction of wind and water waves (P G Sullivan, personal communication 2015).

Progress in RANS modelling has not been rapid, and in some cases a rigorous validation has been hampered by a proliferation of versions of each model. This issue has been addressed, although not yet for a large enough number of models at the time of writing, by the excellent NASA turbulence-modelling resource[30]. The site provides both definitive formulations and a nomenclature for all model versions, and grid-converged solutions for a variety of canonical flows. The range of cases is gradually expanding.

To set the stage in our industry, we may consider the problem of calculating the flare and landing maneuver of an airliner, therefore a configuration with high-lift devices, landing gear, spoilers, moving control surfaces, ground effect, thrust reversers, and unsteadiness lasting many seconds. Curiously, with computing power gradually rising from current levels to infinity, the criticality of turbulence modelling within CFD is likely to rise, and ultimately fall. The computing power is, further, likely to dictate the type of physical modelling that is optimal at a particular time. More specifically, as of today a solution for this landing maneuver that is accurate to the degree needed in our industry is out of reach even with the least costly type of turbulence modelling, namely, RANS. Therefore, it is arguable that physical modelling is not the dominant source of error. With power in the exaflop range and beyond, and effective grid-adaptation systems, the numerical errors will subside, so that physical modelling will stand out as the source of error; we do not expect this to be the case until the mid-2020s. Predictions beyond that time are not as easy, but are not impossible either.

It is conceivable that computing power will someday make DNS in aeronautics possible, so that modelling proper would disappear, and the turbulence considerations would be reduced to ensuring that the grid and time resolution are adequate. In 2000, one of us boldly anticipated this to happen around 2080[28], but by now we are not confident of this for the 21st century, or even that it will ever happen. The reason is that, without being experts in computer hardware, as mentioned above we believe Moore's law is clearly weakening, and fail to see where the
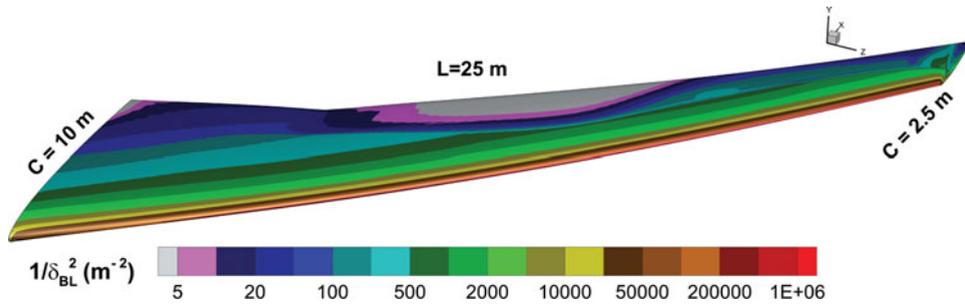
Figure 4. Contours of the quantity $1/\delta^2$ on a wing at flight Reynolds number, with fully turbulent boundary layer. (Courtesy of M. Strelets.)

other 20 or so orders of magnitude (doubling every two years from now until 2080, as had to be assumed) will come from. The expression 'post Moore's law era' is in use in high-performance computing circles. Naturally, the present considerations ignore the chance of a major breakthrough; for instance, through quantum, biological, or other radically different type of computing.

Turning to the type of turbulence treatment, of course RANS and DNS are not the only options, and LES has great potential, although methods which soundly combine elements of RANS and of LES have even more and earlier potential. The world of hybrid RANS-LES methods calls for a fairly complex taxonomy. 'Pure RANS' and 'pure LES' are quite clear in terms of the equations that are solved. This is although pure RANS can produce unsteady and even chaotic solutions for massive separation, which is generally favorable for accuracy, but difficult to control. Pure LES we define as a method that requires no modelling akin to RANS, and in particular does not depend on specifying the value of the Karman constant κ; this is also called wall-resolved LES, or quasi-DNS, and it will not become possible long before DNS does. Much more practical is Wall-Modelled LES (WMLES), in which the lower part of the boundary layer contains modelling of very numerous eddies such as streaks, implicitly contained in each grid cell. This modelling acquires a RANS nature, and implies a value for κ, which we find to be a simple test of whether 'turbulence modelling' is involved. This gives access to arbitrary Reynolds numbers and is bearable, considering that κ and the associated log law are the best-established facts of turbulence. In WMLES, the upper part of the boundary layer relies on LES, with of course adequate grid resolution. Our prediction in 2000 that LES would prevail in the 2045 time frame assumed wall modelling, and a few other generous assumptions[28]. Therefore, we have no reason to make any more optimistic predictions, especially in this post-Moore's law era, and progress in RANS modelling remains a high priority.

Figure 4 is key to the realities of WMLES, which are worth reiterating, because wishful thinking is all too common in this area. In such a method, with fully successful wall modelling, the number of grid points per cube of boundary layer, of size δ, is nearly independent of Reynolds number. As a result, the cost per unit area is proportional to $1/\delta^2$, and this quantity is shown on a wing at flight Reynolds numbers (work of Dr M Strelets, 2015). It is striking how the (red) band that straddles the attachment line of the swept wing dominates the cost of the entire WMLES. If we count roughly 32^3 points per boundary-layer cube, each square centimetre in this band is costing close to 100,000 points. For both wings the integral of $1/\delta^2$, in other words the number of cubes needed to fill the boundary layer, is roughly
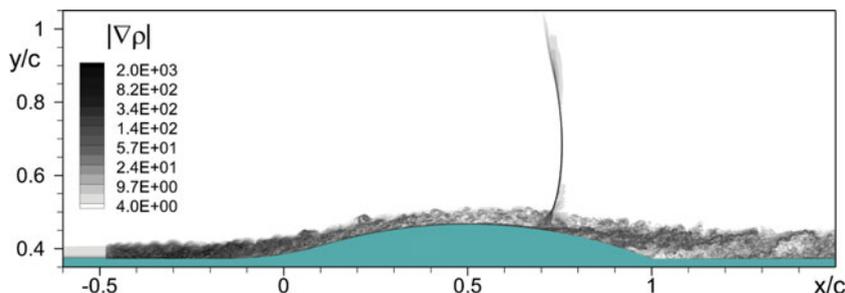
Figure 5. RANS followed by wall-modelled LES of a shock-boundary-layer interaction. Numerical schlieren visualisation. The switch from RANS to LES occurs at *x/c* = −0.48. (Courtesy of M. Strelets.)

$N_{\text{cubes}} = 4{\times}10^6$, leading to 10^{11} points as already pointed out in 1997[31]; for a high-lift configuration, the numbers would roughly triple. The number of time steps from a 'cold start' would be in the millions[31]; pure LES does not enjoy any acceleration to steady state, the way steady solvers do. This is why RANS is the only option in that band. However, what is also striking is how rapidly the quantity $1/\delta^2$ falls away from the attachment line (notice the exponential spacing of the contour levels). The consequence is that switching from RANS to WMLES in the boundary layer becomes manageable fairly rapidly, and this is why we contend that this will be the turbulence treatment of choice in the foreseeable future. If we rule out a major breakthrough in accuracy in RANS modelling for sensitive regions such as shock-boundary-layer interactions[25] the recourse to LES is necessary not only for massive separation, as was argued when creating DES[28], but also for boundary layers in severe pressure gradients and other non-trivial perturbations. Fortunately, the region with very high $1/\delta^2$ is also a region with favourable pressure gradient, which is precisely what keeps the boundary layer so thin, and such regions are easy to predict for RANS. This suggests a zonal treatment. Now the switch from RANS to WMLES is not trivial, technically. It requires a large reduction in the grid spacing from 'RANS spacing' to 'LES spacing,' and the artificial generation of three-dimensional unsteady 'LES content.' This is a very active research field, for good reasons. The engineering CFD method of this type will proceed through preliminary RANS solutions, followed by the automatic selection of RANS and LES regions, appropriate grid generation with a strong dependence on δ, and arrangement of the synthetic turbulence generator. These generators are not perfect, and the simulation will have a narrow adjustment band in which the skin friction will be inaccurate, but not enough to have a global impact, especially if the pressure gradient is still favourable. As computing power rises, the RANS-LES boundary will move towards the attachment line. The substitution of LES for RANS will consist in a boundary between the two approaches moving forward, rather than being sudden.

Figure 5 displays the beauty of such a hybrid simulation in a boundary layer, in a research case. The attached boundary layer has been seeded with LES content and encounters a shock, which causes separation with a λ pattern. This type of simulation has proven to be far more accurate and reliable than any RANS model. However, this did not happen with resolutions such as 32^3 in each boundary-layer cube; the one in Fig. 4 used 2 Bn points total, and the number of points per cube just ahead of the shock was about 10^6 (which is 30 times larger than 32^3).

We now summarise our predictions for turbulence treatment at the Reynolds numbers of interest. DNS and wall-resolved LES will not be used. Pure RANS cannot be fully eliminated

(and is far superior to LES with a skeletal resolution such as 3^3 per δ cube), but is not to be trusted after massive separation, and ultimately not even in boundary layers in strong adverse pressure gradients. The switch from RANS to WMLES will not happen globally, but instead, hybrid simulations will see the boundary move forward to gradually shrink the RANS region, reducing it to the thinnest areas of boundary layers, which are the least difficult to predict but cannot be ignored. The laminar regions will be treated with 'RANS grid spacing' and a choice of transition-prediction approach; this is another very active research area, which we do not delve into here. The complete turbulence treatment will be quite complex with several steps and fully integrated with the grid design; this complexity is unfortunate and aesthetically unappealing, but this approach is the only one with the ultimate potential to face the hurdle of turbulence in our industry.

3.7 Integration with other disciplines

Structural mechanics is a major discipline that interacts with CFD because we are dealing with flexible structures such as aircraft wings, control surfaces, and fan blades. Recall that geometry is an essential input to CFD. Thus, when dealing with flexible structures, there is no recourse other than to couple CFD with a structural model. The fidelity of the structural model could range from simple approximations such as beam model (early in the design phase of an aircraft) to a sophisticated finite element model. Fluid-Structure Interaction (FSI) is fast becoming an important area with many applications. Typically, for static situations, aerostructural coupling is done in a loose fashion and converges in just a few iterations. Dynamic load predictions (including flutter boundary calculations) also call for time-dependent or time-harmonic CFD analysis to be coupled with a structural module[32]. CFD can interact with 6-DOF models to determine trajectories (e.g., of missiles or aircrafts in stall/spin). CFD can also be used to compute dynamic derivatives[33]; it then becomes possible to couple with control laws and carry out aircraft manoeuvres.

At a deeper level, CFD and structures need to be integrated because the aerodynamic forces size the structure and influence the mass of the aircraft, and therefore its cost and its value. In addition, the critical conditions are at the edge of the flight envelope, and therefore probably more difficult to predict physically. A prime example of this difficulty is the slope of the lift curve of each section of a transonic wing at higher angles of attack and Mach numbers.

Unsteady CFD, using DES or LES approaches, is tentatively used to predict vibrations and noise. This is tentative, first, because the detection of unsteadiness via CFD is not error-proof. For instance, buffet remains a challenging field, but one with great impact, whether it is high-speed wing buffet or low-speed flap buffet. These are phenomena with a rather narrow frequency range. CFD has also helped with the noise of small components such as temperature probes. Broadband noise adds the cost of resolving, in the time domain, a very large number of cycles. An example is jet-flap interaction. The propagation of noise adds to the difficulty. Propagation to the fuselage demands a large volume of fine grid, followed by careful coupling with structural equations, ultimately leading to cabin noise. Community-noise prediction is perforce done by post-processing the unsteady simulation with a far-field propagation approach; these approaches remain very delicate, especially in non-uniform flow fields with boundaries, and marked by controversies particularly over the role of quadrupoles in the acoustic analogy.

Another major discipline that interacts with CFD is icing. Given the flow-field information from CFD, sophisticated icing models (e.g., Ref. 34) are used to generate ice shapes and the new geometries are analysed in CFD. These models account for surface roughness, phase

change, runback, and other factors affecting ice accretion. Grid generation of configurations with ice is a non-trivial task because the ice shapes are complex, and separation off the ice horn is likely. For some types of icing, a treatment as a rough surface, entering the turbulence model, is preferable.

As of now, most multi-disciplinary analyses are carried out by cycling through the solvers of the various disciplines involved in a loosely coupled fashion. This is a sequential algorithm with parallelism (as it exists) exploited only within each discipline. The approach can be characterised as a non-linear Gauss-Seidel algorithm and is not guaranteed to converge. The loosely coupled approach still has some advantages. It is relatively easy to build the infrastructure, it requires minimal exchange of information at each iteration, and leaves the discipline-specific solvers unaltered. On the other hand, a strongly coupled approach, based on a damped Newton method, is guaranteed to converge. This approach also exploits parallelism across the various disciplines permitting them to be solved concurrently. The strongly coupled approach is daunting in terms of problem formulation and creation of an analysis framework; it also requires far greater computer resources.

4.0 USER AWARENESS AND EDUCATION

4.1 Current practices and concerns

The quality of CFD answers depends at least on three factors: the code, the available computer resources (which limit the grid resolution and number of iterations), and the user. It is unfortunately easy to misuse a code and computer (one of us was recently shown a flow field past a complete aircraft in an engineering context, but the user had taken the silhouette of the aircraft, and run a 2D solver!). The obvious possibilities include accepting poor iteration convergence (in another incident, a presenter stated that the residuals had levelled off and that it 'indicated convergence'). This flaw in a solution is obvious since all codes provide residual histories, but insidious under the usual work pressures; today, a very small proportion of solutions in industrial work achieves machine-zero convergence. Independently, except in the simpler cases, there is often poor grid convergence, which is more difficult to detect, and surprisingly difficult to test for in a direct and simple manner with the common grid generators and gridding strategies. The geometry may also have been simplified too much, say by omitting slat supports, or not have been adjusted for aeroelastic effects carefully enough; this issue is just as present in the wind tunnel, of course. Other possibilities include inadequate treatment of transition and turbulence; in this case, adequate treatments may not even be available, for reasons discussed at length above, and the correct response from the user is to treat the results with scepticism, and to seek validation on similar cases. Excellent workshops have been held, but at times they illustrate how small a region of the flight envelope we actually master.

We recognise the time and cost pressures engineers work under, but we believe the software could be much better at boosting the 'situational awareness' (an expression we borrow from pilot evaluations) the users have of their solutions, and at improving their skills through the practice of running CFD. Both the engineering companies and the CFD providers, whether private or government, offer training but most users appear to spend only a few hours a year in formal training.

A prime example of poor practice is accepting a series of iterations to steady state, in which the residuals dropped by only a few orders of magnitude. It is common then to take the limit cycle of the forces and moments, and average them. This is surprising to us and not justified,

primarily because the iteration steps towards steady state are non-physical, and not closely related to time integration; in particular, they do not conserve vorticity in the inviscid regions. In some cases, the limit cycle may concern only a small region of space, where the grid may be especially poor, but if the forces and moments vary noticeably, it is likely that large regions are affected. The entire solution is therefore tainted. A logical measure to resolve this is to continue the run in time-accurate mode, but when this is suggested the users usually object that it is not affordable (in addition, spurious vorticity created during failed iterations towards steady state could be distributed over a large region, making it difficult to evacuate). This suggests effort would be well spent on the performance of codes in time-accurate mode, which is very uneven in our experience. There is no guarantee that the limit cycle of unsteady RANS is perfect, but at least the inviscid physics are correct, and the turbulent physics addressed in the best way possible inside the RANS framework. In that respect, we note that no code should be dependent on only one turbulence model, and that all users should be familiar with more than one model, and test the sensitivity of solutions to the model when there is a sense that an inaccurate prediction may be caused by modelling flaws.

Regarding grid convergence, essentially all codes should use a sequence of grids. The user should be automatically exposed to the results from the two finest of these grids, for forces and pressure distributions. If they differ noticeably, the warning is clear. We have not seen this done, although it is easy to arrange. The most common danger in Navier-Stokes work is that the grid is not adequate to capture the turbulent regions, including boundary layers, shear layers, and vortices. In particular, the sufficient grid depends on the angle-of-attack, which is rarely achieved except with automatic adaptation. Presently, almost all users run an entire polar on the same grid, a practice which would be adequate only with an 'overkill' of grid density.

Another support for awareness of physics and a tool for design will be available if and when a rigorous definition of induced drag, wave drag, and parasitic drag from viscous flow fields is established. These concepts are constantly used by designers, but only for clean wings. The extension of lifting-line theory is a fascinating and stubborn problem (even once a theory is created, it could be defeated by grid coarsening in the wake). Success in this domain could temper the erosion of classical aerodynamic knowledge in the younger generations.

4.2 Flow visualisation and other interface possibilities

Flow visualisation is also very helpful in some situations, and again should be provided automatically. Pressure, skin friction, turbulence index, first wall spacing in wall units y^+, and surface streamlines giving good coverage without the user setting starting points should be available with no effort. Views of the vortical turbulent regions in the field should be easy to obtain, marked by vorticity, the Q criterion, or even eddy viscosity. We know that providing this is made more difficult by massively parallel computer architectures. However, the frontier for the expanded and accurate use of CFD is precisely in cases for which viscous effects become pivotal, and the user's judgment of the physical and numerical soundness of the solution should be nurtured with intent. When the message from visualisation and other indicators is negative, the appropriate response will range from simply declaring that CFD is not ready for this particular flow, to refining the grid, to switching from a steady to an unsteady RANS formulation, to using a hybrid RANS-LES approach such as DES. DES would be helpful for massively separated cases, or any which require detailed unsteady information such as noise and vibration.

In the longer term, virtual reality and similar systems beyond the simple two-dimensional screen will be considered. Without turning CFD into too much of a source of entertainment, feeling the surface with one's fingers, with its pressure fluctuations and vibrations, the flow speed and direction and the turbulence, will be a resource when solving design problems such as separation, vortex breakdown, and buffet. The user will select a component such as flap by voice, and reach through the other ones to touch it; he/she will hold the component, say a landing-gear door, by hand and feel the (unsteady) force and moment on it. This idea is not ours, since Boeing's history of the B-47 wind-tunnel tests of the 1940s and the now-common underslung nacelle concept mentions that 'the concept was tested in the Boeing wind tunnel by mounting model engine nacelles on the end of a pole (the 'broomstick' test) and moving the nacelles around the wing until the optimal position was discovered – forward and below the wing.' Even moving a point by hand with a 'three-dimensional mouse' when exploring the flow field would empower the user when looking at a screen; the equivalent of a smoke wand would be provided. In propulsion applications, the local temperature is very valuable; for noise studies, the user will have instant aural access to noise anywhere in the field. We understand that the storage implications of this type of communication are very large. However, the current displays of isosurfaces or contours remind us of Plato's cave, and in addition they are supplied in slow motion.

The 'experience' will, further, direct attention to regions of weakened accuracy, either due to 'bad cells,' or marginal grid density possibly detected by comparing the best two grids, or of degraded confidence in the transition/turbulence treatment. Colour will show regions with marginal separation, or values of the *n* factor (used in transition prediction) close to critical. Again, the idea is to educate the users continuously about numerics and physics and to illustrate how, when used at the frontiers, CFD is not a black box. As well as showing rich information contained in the solution, the interface will issue gentle reminders of the 'ethics of CFD.'

In a more abstract exercise, the user will 'travel' the flight envelope for each flap setting, with a choice of independent variables including speed, altitude, mass and centre of mass, g factor, and so on. The buffet or stall boundary may be revealed by a buzzer or 'stick shaker.' Naturally, one day CFD will be directly driving real-time flight simulators, and real stick shakers will obviously be involved. The boundary of the envelope of high-confidence predictions will also be indicated, and in particular, conditions that permit multiple solutions to the equations will be clearly indicated. The existence of multiple solutions has been strongly established and usually blamed on the sensitivity of smooth-body separation[24], but reliably extracting them will be a very delicate problem.

5.0 CONCLUSIONS

The widely expected substitution of CFD for the vast majority of ground and flight testing in the aerospace and similar industries, although announced in the 1970s, will take decades from today to complete, gradually expanding from the center to the edges of the operational envelope, from isolation to complete collaboration with other disciplines, and from innocuous to safety-critical decisions. We believe that the recent marked weakening in the rate of increase of computer power is durable and linked to the laws of quantum physics, but cannot exclude the possibility of a revolution in hardware design. This substitution by CFD will contribute somewhat to addressing the perennial concern that aircraft and similar programs take much too long and cost too much. It can also reduce industrial and schedule risk. Another benefit of

CFD is the increasing power to approach true global optima in aerospace design, breaking down the stubborn barriers between disciplines. The challenges in physical modelling of transition and turbulence will not be truly overcome in this century, if ever. The solver and grid-adaptation issues are still resilient in spite of the talent and effort applied worldwide, but we see no option other than addressing them squarely. A substantial advantage of CFD will be that, with higher accuracy standards and more flexible airframes, wind-tunnel testing will greatly suffer from the impossibility to reproduce the shape of the loaded wing or blade over enough conditions.

We perceive a danger of overconfidence and under-competence in CFD. There are many ways to produce bad CFD, from careless attitudes, to serious inattention to numerics issues and convergence, to being unaware of the mysterious failings and fallacies of turbulence modelling. These are not highlighted by the providers of CFD capabilities, at least not in the packaging of the product, although some of these entities have discussed failures and partial successes openly in scientific papers. We do not doubt that mistakes are made and will be in the future, and we do not blame our non-CFD colleagues for their prudence in adopting CFD, whether on the industry or regulatory side. Such mistakes occur also in physical testing, of course. Another danger is the erosion of physical and engineering judgment in aerodynamics, not to mention ignorance of the very equations the codes are solving and their connection to physics. The training of users deserves more attention, and we have made proposals in that direction. We can hope the power of CFD will not only improve the economics of conventional aircraft, but also empower us to bring new concepts to the air, for example a supersonic transport with acceptable fuel-burn and sonic-boom penalties. Widespread acceptance of certification by analysis, primarily by CFD, will be a remarkable achievement and is for us an inspiring mission.

ACKNOWLEDGEMENTS

REFERENCES

1. JOHNSON, F.T., TINOCO, E.T. and YU, N.J. Thirty years of development and application of CFD at Boeing Commercial Airplanes, Seattle, USA, *Comput Fluids*, December 2005, **34**, (110), pp 1115-1151.
2. ABBAS-BAYOUMI, A. and BECKER, K. An industrial view on numerical simulation for aircraft aerodynamic design, *J Math Indy*, 2011, pp 1-10.
3. STRAWN, R., NYGAARD, T., BHAGWAT, M., DIMANLIG, A., ORMISTON, H.S.R. and POTSDAM, M. Integrated computational fluid and structural dynamics analyses for comprehensive rotorcraft analysis, *AIAA Atmospheric Flight Mechanics Conference and Exhibit*, 2007, 10.2514/6.2007-6575.
4. DUMAS, L. CFD-based Aerodynamic Optimization for Automotive Aerodynamics, *Optimization and Computational Fluid Dynamics*, THEVENIN, D. and JANIGA, G. (Eds), 2008, Springer, pp 191-215.
5. ISLAM, M., DECKER, F., VILLERS, E.D., JACKSON, A., GINES, J., GITT-GEHRKE, T.G.A. and FONT, J.C.I. Application of Detached-Eddy Simulation for Automotive Aerodynamics Development, 2009, *SAE Int*, paper 2009-01-0333.
6. SANDERSE, B., VAN DER PIJL, S. and KOREN, B. Review of computational fluid dynamics for wind turbine wake aerodynamics, *Wind Energy*, 2011, **14**, (17), pp 799-819.
7. BERTRAM, V. *Practical Ship Hydrodynamics*, 2nd ed, 2011. Butterworth-Heinemann.

8.   Slotnick, J., Clark, R., Friedman, D., Yadlin, Y., Yeh, D., Carr, J., Czech, M. and Bieniawski, S. Computational aerodynamic analysis of the formation flight for aerodynamic benefit program, 2014, *AIAA-2014-1458, 52nd Aerospace Sciences Meeting*, 2014.

9.   Spalart, P.R. and Bogue, D.R. The role of CFD in aerodynamics – off design, *Aero J*, 2003, **107**, pp 323-330.

10. Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W. and Laurie, D. CFD Vision 2030 Study: A path to Revolutionary Computational Aerosciences, January 2014, Technical Report*NASA/CR-2014-218178*.

11. Le Doux, S.T., Vassberg, J.C., Young, D.P., Fugal, S., Kamenetskiy, D.S., Huffman, W.P. and Melvin, R.G. Study based on the AIAA aerodynamic design optimization discussion group test cases, *AIAA J*, 2015, **53**, (17), pp 1910-1935.

12. Melvin, R.G., Huffman, W., Young, D., Johnson, F., Hilmes, C. and Bieterman, M.B. Recent progress in aerodynamic optimization, *Int J Numer Methods Fluids*, 1999, **30**, (12), pp 205-216.

13. Vassberg, J.C., Tinoco, E.N., Mani, M., Zickuhr, T., Levy, D.W., Brodersen, O.P., Eisfel, B., Crippa, S., Wahls, R.A., Morrison, J.H., Mavriplis, D.J. and Murayama, M. Summary of the fourth AIAA computational fluid dynamics drag prediction workshop, *J Aircraft*, 2014, **51**, (14), pp 1070-1089.

14. Rumsey, C.L. and Slotnick, J.P. Overview and summary of the second AIAA high lift prediction workshop (Invited), 2014, *AIAA SciTech, 52nd Aerospace Sciences Meeting*, National Harbor, MD.

15. Spalart, P. and Wetzel, D. Rudimentary landing gear results at the 2012 BANC-II airframe noise workshop, *Int J Aeroacoust*, 2015, **14**, (11–2), pp 193-216.

16. Anderson, B., Shur, M., Spalart, P., Strelets, M. and Travin, A. Reduction of aerodynamic noise in a flight deck by use of vortex generators, 2005, *AIAA-2005-426*.

17. Allmaras, S.R., Bussoletti, J.E., Hilmes, C.L., Johnson, T., Melvin, R.G., Tinoco, E.N., Venkatakrishnan, V., Wigton, L.B. and Young, D.P. Algorithm issues and challenges associated with the development of robust CFD codes, *Variational Analysis and Aerospace Engineering*, Buttazzo, G. and Frediani, A. (Eds), 2009, Springer, pp 1-20.

18. Johnson, F.T., Kamenetskiy, D.S., Melvin, R.G., Venkatakrishnan, V., Wigton, L.B., Young, D.P., Allmaras, S.R., Bussoletti, J.E. and Hilmes, C.L. Observations regarding algorithms required for robust CFD codes, *Math Model Nat Pheno*, 2011, **6**, (103), pp 2-27.

19. Mavriplis, D.J. Results from the 3rd drag prediction workshop using the NSU3D unstructured mesh, *J Aircraft*, 2008, **45**, (13), pp 750-761.

20. Diskin, B., Thomas, J.L., Rumsey, C.L. and Schwoppe, A. Grid convergence for turbulent flows, 2015, *AIAA-2015-1746, SciTech-2015*.

21. Bieterman, M.B., Bussoletti, J.E., Hilmes, C.L., Johnson, F.T., Melvin, R.G. and Young, D.P. An adaptive grid method for analysis of 3D aircraft configurations, *Comput Methods Appl Mech Eng*, 1992, **101**, pp 225-249.

22. Aftosmis, M.J. and Nemec, M. Cart3d simulations for the first AIAA sonic boom prediction workshop, January 2014, *AIAA 2014-0558*.

23. Fidkowski, K. and Darmofal, D. Review of output-based error estimation and mesh adaptation in computational fluid dynamics, *AIAA J*, 2011, **49**, (14), pp 673-694.

24. Kamenetskiy, D.S., Bussoletti, J.E., Hilmes, C.L., Venkatakrishnan, V., Wigton, L.B. and Johnson, F.T. Numerical evidence of multiple solutions for the Reynolds-averaged Navier–Stokes equations, *AIAA J*, 2014, **52**, (18), pp 1686-1698.

25. Spalart, P. Philosophies and fallacies in turbulence modelling, *Prog Aerosp Sci*, 2015, **74**, pp 1-15.

26. Farrell, P., Birkisson, A. and Funke, S.W. Deflation techniques for finding distinct solutions of nonlinear partial differential equations, *SIAM J Sci Comput*, 2015, **37**, (4), pp A2026-A2045.

27. Witherden, F.D., Vermeire, B.C. and Vincent, P.E. Heterogeneous computing on mixed unstructured grids with PyFR, *Comput Fluids*, 2015, **120**, pp 173-186.

28. Spalart, P, Strategies for turbulence modelling and simulations, *Int J Heat Fluid Flow*, 2000, **21**, pp 252-263.

29. Smith, C., Beratlis, N., Balaras, E., Squires, K. and Tsunoda, M. Numerical investigation of the flow over a golf ball in the subcritical and supercritical regimes, *Int J Heat Fluid Flow*, 2010, **31**, (13), pp 262-273.

30. Rumsey, C.L., Smith, B. and Huang, G. http://turbmodels.larc.nasa.gov, (online).

31. SPALART, P., W.-H. Jou, STRELETS, M. and ALLMARAS, S.R. Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach, *First AFOSR International Conference on DNS/LES, Advances in DNS/LES,* LIU, C. and LIU, Z. Z. (Eds), 1997, Greyden Press, Columbus, OH.
32. SCHUSTER, D.M., CHWALOWSKI, P., HEEG, J. and WEISMANN, C. Summary of data and findings from the first aeroelastic prediction workshop, July 2012, *ICCFD7*.
33. RONCH, A.D., VALLESPIN, D., GHOREYSHI, M. and BADCOCK, K.J. Evaluation of dynamic derivatives using computational fluid dynamics, *AIAA J*, 2011, **50**, (12), pp 470-484.
34. BIDWELL, C. Icing Analysis of the ANSA S3 Icing Aircraft Using LEWICE3D Version 2, 2007, SAR Technical Paper 2007-01-3324.