To CFD or not to CFD that is the question?

– Part II

David Hollom looks at where CFD fits into the design spectrum

P NIETO/DPPI

Williams F1's previous half-scale wind tunnel. Williams moved to a $50 million full-scale tunnel in 2005 but teething troubles cost the team much of that season's competitiveness

Twenty years ago hull design programs, operating on computers using the ubiquitous Motorola 68000 chip, took around 20 seconds to calculate and update the on-screen view of the lines plan, after the movement of one of the vertices that control the hull shape. And that was at quite a low resolution (ie a course surface mesh). At higher resolutions, giving greater accuracy, you might have to wait at least a minute to see the change and then a further 20 seconds or so to view the hydrostatics and the rating. Quick, if you had been used to drawing by hand but deadly slow by today's standards.

With modern computers a set of hull lines can be produced in minutes at the highest resolution and thus accuracy that you could desire (not the final design but a good starting point). And in not too many more minutes it can be just about finished and be accurate to within a few decimal places for every target hydrostatic parameter a designer could ever wish to achieve, together with a spot-on target rating (of course a designer will usually then tweak away for varying amounts of time). And there is no wait-time for an on-screen display of the hydrostatics or rating.

With our own program, HullDes, it is possible to have windows open, simultaneously, for hydrostatics and rating and see them alter as you move the vertices to inform the hull shape. Although we have not yet done so, it would seem to be entirely feasible to have another window open, displaying results from the velocity prediction program (VPP), although there would have to be some considerable interaction within the program, altering, in a predetermined and automatic way, a number of factors, such as appendage size, sail area, hull weights and bulb size and so on as displacement, righting moment, hull size and rating etc changed.

As always the computer will not guarantee a fast boat – that depends upon the operator. But it does allow a lot more options to be examined in a given time.

To digress a little, while I am 100 per cent behind the use of computers in yacht design I do find it a little strange that modern courses in yacht design do not require the student to draw a set of lines by hand. Everything from the outset is now carried out by computer and, while I would never willingly go back to producing a set of lines by hand, I do feel that you should have done it a few times; I am certain that after sweating over a set of lines for a few days, striving to achieve the desired shape and struggling to make them fair, it produces a better understanding of shape and how water might flow around that shape. Or perhaps I'm just turning into a grumpy old man?

Turning again to the Navier-Stokes equations it is, as already mentioned, only this huge increase in computing power that has made their partial solution possible. To differentiate it from other mathematical methods such as panel or potential flow, that are much more economical of computer power, this branch of theoretical fluid dynamics is known as computational fluid dynamics (CFD). But how accurate is the end result and to what extent has it made tunnels and tanks and flumes redundant?

It all depends upon what you mean by accurate. Every time you go sailing, or fly your glider, or race your car, conditions are different. The amount and type of turbulence in the ocean or in the air is probably never quite the same from one day to the next and, as we have seen in Part 1 (last month), the degree and type of turbulence have a large bearing on performance, particularly where the maintenance of laminar flow is important. Partly, but not entirely because of their effect on Reynolds number (Re), which affects the boundary layer and thus the viscous drag, changes in temperature and particularly in the case of air, density, also affect performance; a result that is accurate for one set of conditions on one day and in one place may in fact be inaccurate for another set of conditions on another day or even on the same day in another place.

Assuming that by accuracy we mean the ability to repeat the results of a physical experiment in a tunnel, tank or flume, at a particular level and type of turbulence, density, temperature, Froude or Mach number and Reynolds number, there are several reasons for inaccuracies. First, there are the simplifying assumptions and approximations necessary to solve the equations. Second, there are rounding errors.

Because the number of decimal places used in the calculations must be finite they must be rounded up or down. While this would have little effect on a small number of calculations, over a calculation having, say, ▷

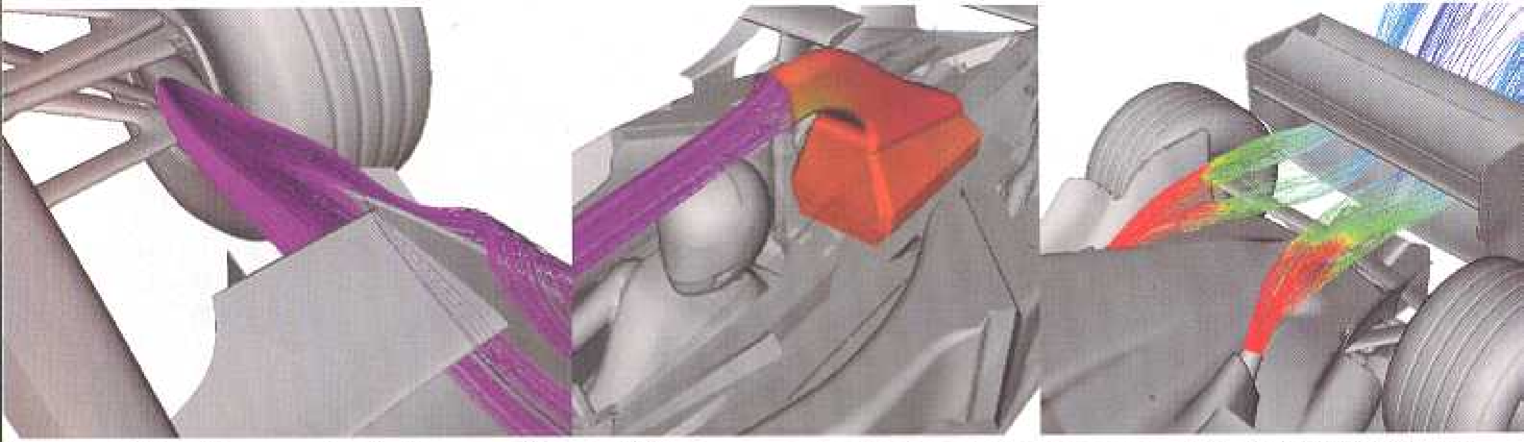

Williams F1 are one of only a tiny group of teams with the supercomputer power to model an entire car as a single unit in CFD – largely as a result of their support from computer manufacturer Hewlett Packard. Nevertheless the bulk of Williams' CFD work is confined to studies of individual elements (*above*) with the wind tunnel remaining the ultimate arbiter – particularly when it comes to global car behaviour

WILLIAMS F1

a million steps these can mount up cumulatively, although one might expect that the roundings up and down might average out. However, in manipulating small and large numbers this can become a factor.

Finally, until very recently it was necessary to run the programs using either a fully turbulent or a fully laminar boundary layer. As the computation of viscous drag depends very much on the type of flow, laminar or turbulent, and the amount of each, this would seem to be a major source of inaccuracy. However, remember that in tank testing it is normal to fix transition near the bow, so that traditional methods of tank testing suffer from the same problem when comparing results with reality.

More recently, developments of the programs known as Large Eddy Simulation (LES) are being evolved where the mesh is small enough to completely accommodate the largest predicted eddy. According to Kolmogorov's theory of self-similarity, large eddies of the flow are dependent on the flow geometry while smaller eddies are self-similar and thus have an almost universal character. Thus, in LES the large-scale motions of the fluid are calculated while the effect of the smaller universal eddies on the larger ones are modelled using a sub grid scale (SGS) model.

The program is then able to predict transition, and thus the amount of laminar and turbulent flow and also, one assumes, the position and extent of any bubble and its effect on the flow, both of which would appear to be a great improvement. Before you get too excited, though, computationally, it is far more expensive than Reynolds-averaged Navier-Stokes methods (RANS), particularly in the region of walls (the vehicle skin) so that simulating such flows usually exceeds the limits of even today's supercomputers.

For this reason zonal approaches are often utilised, with RANS or other empirically based models replacing LES near the walls; but this would seem to somewhat defeat the purpose, as transition could not then be determined. There is, also, the intriguing thought that, while the flow within a laminar boundary layer is fairly predictable, the flow within a turbulent boundary layer is, by definition, random and chaotic. That flow will, thus, never repeat itself so any solution, no matter how accurate, is only true at that moment. At any other time it will be different.

Comparisons are often made between designing a Formula One car and designing an America's Cup boat and thus, because they, the car fraternity, make extensive use of CFD we, the sailing community, should do so too. However, while there are many similarities in methodology the actual design problems are very different.

The car designer is very much interested in downforce. Yes, he is interested in drag, but on some slow circuits he will willingly increase downforce for an equal drag penalty, ie he would consider the extra downforce worthwhile down to a lift/drag ratio of unity. The reason of course is that the car carries the extra speed that the extra downforce provides in the corner, onto the next straight and some way down it. Yes, the lower-drag car will eventually exceed the speed of the high-downforce car, in a straight line, but if the straights are not long enough for this advantage to materialise and there are many corners, the high-downforce car will be faster around that track. Obviously there is a trade and on circuits with longer straights and fewer corners, a better lift/drag ratio is used for deciding whether the extra downforce is worthwhile.

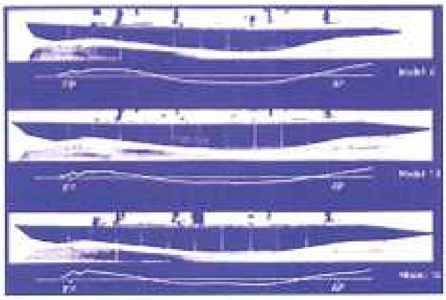

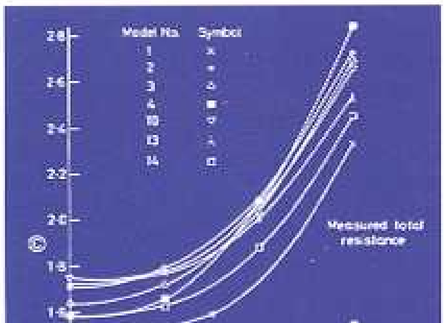

Top: Fig 1 – calculated versus actual (in the tank) wave systems for 12-Metre designs. *Above*: Fig 2 – computed wave drag with total measured drag (wave plus viscous)

We in sailing, on the other hand, are, or should be, vitally interested in drag not lift. As mentioned in previous articles, the rig forces predetermine hydrodynamic lift (which is easily calculated or, alternatively, can be obtained from a breakdown of forces from within the VPP), so that all that is required to sail faster to windward is to produce the required lift force for less drag (assuming the same stability). That way you will either point higher, by having a better hydrodynamic lift/drag ratio, or sail faster, which in itself will reduce induced drag which will, yet again, allow you to sail either faster or higher. As always, a good helmsman or a VPP will find which solution is optimum, or at least very near optimum. The lower drag, obviously, also pays downwind.

Navier-Stokes programs are very good at predicting lift or, in the case of a racing car, downforce but are, as we have seen, not always as good at predicting drag. For that reason they are more useful in situations where lift or downforce dominates rather than drag, as with an F1 car. Also the Formula One car has many downforce producing features, which mutually affect each other. It is therefore desirable to model a chunk of air around the car to catch this interaction, which necessitates the use of either Navier-Stokes or their non-viscous counterpart, Euler programs.

Cars, I suspect, also work in a turbulent

WILLIAMS F1

environment, so that the necessity to choose either a fully turbulent or a fully laminar boundary layer does not create a problem.

Boats, on the other hand, have a mixed boundary layer, the flow starting as laminar but then somewhere a transition to turbulent flow occurs. Also, we only really have the interaction of the fin and the rudder to contend with so that it is not strictly necessary to model a block of fluid around the boat. Thus, simpler panel or potential flow methods, which only model the flow at the surface of the body, are probably adequate for our needs. Because they normally utilise an integral boundary layer method, such programs also have the advantage of perhaps predicting the state of the boundary layer and thus the viscous drag better, and therefore may well give results that are as good or even better than Navier-Stokes or Euler programs... but achieved very much more quickly and cheaply.

However, to reduce drag it is necessary to have a good understanding of the flow and how it can be improved, and it is here that Navier-Stokes programs really seem to score. The flow visualisation is excellent and, assuming that the model of the flow is accurate, there are many areas of interest where improvements to the flow offer large reductions in drag.

Interestingly, during the last British America's Cup Challenge in 12-Metres we used two potential flow programs (not strictly CFD) that were then very much in their infancy. Through one of our sponsors, British Aerospace, we used a potential flow program called SPARV, developed by Brian Maskew at BAE in Brough. Brian Maskew left BAE and joined Analytical Methods in Seattle and subsequently wrote VSAERO, which was then used by Dennis Conner's successful design team.

SPARV did not have a free surface facility and so could not model wave drag whereas VSAERO did. However, to partly overcome this deficiency, we used a separate potential flow wave-making prediction program developed by George Gadd at BMT (previously the UK National Physical Laboratory) where we were tank testing (for part of the time in the tank where Barnes Wallis tested his bouncing bomb).

As a simplifying assumption and as their name suggests, potential flow programs ignore the effects of viscosity and as clearly viscosity affects, to some extent, wave making, there must be some inaccuracy. However, these errors appear to be small. Although perhaps not 100 per cent accurate in a quantitative sense we did get the order of preference correct and we made great use of it in reducing the wave drag of the UK's final, radical 12-Metre *Crusader 2*. Fig 1, *opposite*, shows a photograph of model 4 (*Australia II*), model 13 and model 14 (*Crusader 2*) in the tank at 9kt. Under each photograph is the calculated wave system, which can be compared with the actual wave system and which appears to agree quite well. Fig 2 shows the computed wave drags with the total measured drags (wave plus viscous) in coefficient form.

So in today's world does CFD make physical testing redundant? The answer, at the moment, seems to be an emphatic no. CFD is just another design tool to be used where it is most appropriate and cost effective. In fact, in 2006 the bulk of Formula One teams only really use CFD as a screen to choose the limited number of ideas that they can test in the wind tunnel; I guess the same is true in the America's Cup.

Even with near-limitless access to the highest levels of the latest CFD, it is significant that Toyota Grand Prix are presently completing another new wind tunnel capable of handling one full-sized car or, at 60 per cent scale, two cars in-line to examine the interaction of one car with the other. Indeed, so important is physical testing that nearly every F1 team presently runs its tunnels on a 24-hour basis. Understanding the flow, however, is paramount in improving performance and it is here that CFD, at present, is perhaps most useful.

One final quote, from Dietrich Kuchemann, formerly head of aerodynamics at the UK Royal Aircraft Establishment, in his definitive work, *The Aerodynamic Design of Aircraft*, seems to sum up the situation quite nicely. After describing the many approximations and simplifying assumptions necessary to achieve an answer, the writer goes on to say:

'These very simple examples will have demonstrated the very many steps we are prepared to take in order to get near a solution. In view of this it is again and again a matter of wondrous surprise when we find that the answers we obtain in this way bear such a close resemblance to what we observe and that our thinking was not misguided after all.' ❑

T DELAUNAY/DPPI

It's easy to appreciate why so much effort is going into trying to replace wind tunnels with CFD tools when you measure the size – and cost – of Ferrari's current half-scale F1 tunnel. Most F1 tunnels run 24 hours a day and Williams' tunnel staff now number around 100