

# Reynolds-Averaged Navier-Stokes in an Integrated Design Environment

Richard Korpus<sup>1</sup>

## Introduction

Computational Fluid Dynamics (CFD) has long played an important role in America's Cup Class yacht design. As the class matures, however, it becomes increasingly difficult to improve on performance. Boat speed differences smaller than 1% now often separate competitors, and designers have to search ever harder for finding even small gains. More accurate tools are needed to identify these small differences, and the burden often falls on Reynolds-Averaged Navier-Stokes (RANS) technology.

While RANS is still considered state-of-the-art CFD, it's essentially the only practical choice. Simpler flow modeling techniques, such as panel codes, are incapable of resolving many of the flow features that affect small performance differences (e.g. form drag, separation, vortex wake details). Model tests are time-consuming and often quite expensive. More complex CFD models (such as Large Eddy Simulations) are not yet sufficiently practical for yacht Reynolds numbers or time-critical design environments. It's important to realize, however, that RANS is only a supplement to existing design techniques, not a replacement. There's no need to apply RANS when a panel code will do, and model tests are still needed to keep designs grounded in reality.

But while faster computers have made RANS analyses *possible*, most applications to date fall short of being *practical*. If an America's Cup designer is to improve boat speed, he or she must analyze hundreds of design alternatives -- not the few isolated samples usually associated with RANS [2, 9, 14]<sup>2</sup>. And even when a large number of flow analyses is available, the measures-of-merit required to rank designs are not obvious RANS outputs like flow detail or drive force. RANS codes have to be coupled to more traditional design decision tools such as Velocity Prediction Programs (VPP's).

Three crucial challenges remain before RANS can be applied in the practical design arena. First, how can a level of throughput be achieved to rank hundreds of design options? Second, how can it be reliably insured that RANS will accurately resolve small design differences? And third, how can RANS be seamlessly integrated with traditional design tools?

Meeting these challenges requires a specific blend of RANS technology. High throughput, for example, requires more than a fast computer. It requires simple yet general grid generation capabilities. The flow solvers must exhibit high degrees of robustness so that time is not wasted debugging unstable runs. A system is required to automate geometry handling, grid generation, and RANS analysis. Only by "removing man from the loop" can high throughput be achieved. Similarly, resolution of small differences requires a high level of grid repeatability, discretization accuracy and turbulence model sophistication.

A RANS package meeting these goals was developed for Oracle/BMW's 2003 America's Cup challenge. The package includes automatic grid generation utilities (both above and below the water), overset grid processors, RANS code interfaces, run management utilities, and automated post-processors. The system is integrated using a series of UNIX shell scripts, and installed on a 64-processor HP/Compaq alpha cluster. Production levels up to 400 three-dimensional (3 million point grid) RANS analyses per month were obtained.

Oracle/BMW utilized the resulting system for a number of functions within their design program. Specifically, RANS was used to:

- Rank design options based on performance;
- Analyze a greater number of options than could be tested or built;
- Verify performance gains too small to test;
- Guide a design process by linking performance to flow, and flow to design feature;
- Guide the allocation of R&D design, build, and test resources.

Performance differences as small as 0.1 seconds per mile (0.03%) were resolved, and 11 seconds per mile (3%) cumulative improvements to boat speed were identified

This paper describes how these levels of productivity and accuracy were achieved. It documents the detailed technologies brought to play, and the steps taken to implement them in the design process. Specific demonstrations are given in each area where RANS played a key role: and include examples from sail, mast, appendage, and hull design. In each case, descriptions are given as to how flow data was used to improve boat performance. The paper closes with a discussion of the expanding role RANS is expected to play in the future of design.

1. Principal Scientist, Applied Fluid Technologies, Inc.  
Annapolis, MD USA [rkorpus@appliedfluidtech.com](mailto:rkorpus@appliedfluidtech.com)  
2. Numbers in parenthesis refer to references.

## RANS Computational Methods

The motion of fluid at high Reynolds number can be simulated using the three-dimensional, unsteady, incompressible Reynolds-averaged Navier-Stokes (RANS) equations. In vector form the dimensional versions are:

$$\nabla \cdot \mathbf{V} = 0$$

$$\frac{\partial \mathbf{V}}{\partial t} + \mathbf{V} \cdot \nabla \mathbf{V} + \frac{1}{\rho} \nabla p - \nabla \cdot \bar{\boldsymbol{\tau}} + g \mathbf{k} = \mathbf{0}$$

The Cartesian vector  $\mathbf{V}$  herein represents flow velocity, and  $p$  the static pressure. Gravity,  $g$ , and its direction,  $\mathbf{k}$ , are included so that free surfaces can be modeled. Unsteady terms are retained for generality [15,17,18], but are only utilized herein for marching to steady state.

The shear stress tensor  $\boldsymbol{\tau}$  includes both mean and turbulent components, and therefore requires turbulence models for closure. The usual assumptions for Newtonian fluids give:

$$\tau_{ij} = -\frac{2}{3} k \delta_{ij} + \nu S_{ij} + \overline{u'_i u'_j}, \quad S_{ij} = \frac{\partial u_i}{\partial x^j} + \frac{\partial u_j}{\partial x^i}$$

Here,  $S_{ij}$  represents the strain of mean flow,  $\nu$  the viscosity,  $k$  the isotropic turbulent kinetic energy, and overbarred terms the turbulent Reynolds stress components.

Reynolds stresses are provided using eddy viscosities computed from one length scale and one time scale. All results presented herein are made with either both scales calculated from additional convection/diffusion equations (two-equation models); or one convection/diffusion equation and one algebraic equation (one-equation models). All two-equation approaches take the form of  $k\varepsilon$  models similar to:

$$\frac{\partial k}{\partial t} + \mathbf{V} \cdot \nabla k - \left( \frac{1}{\text{Re}} + \nu_t \right) \nabla^2 k - P + \varepsilon = 0$$

$$\frac{\partial \varepsilon}{\partial t} + \mathbf{V} \cdot \nabla \varepsilon - \left( \frac{1}{\text{Re}} + \frac{\nu_t}{1.3} \right) \nabla^2 \varepsilon - \frac{\varepsilon}{k} (c_{\varepsilon 1} P_{sol} + c_{\varepsilon 3} P_{irr}) + c_{\varepsilon 2} \frac{\varepsilon^2}{k} = 0$$

$k$  is now the turbulent kinetic energy,  $\varepsilon$  its dissipation rate, and  $P$  its production from mean strain. The Spalart-Allmaras model [20] is used for all one-equation closures, and employs a turbulent energy equation similar to that above.

Although simpler and more efficient models exist, none are known to be universally reliable for the types of flows studied here. Zero-equation models (such as Baldwin-Lomax) are robust and easy to use, but not particularly accurate in adverse gradients of weak separation regions. For similar reasons, all

turbulence models employed here will solve all the way through a boundary layer to the no-slip wall. While wall functions are often used to reduce grid size, they cannot be trusted to capture separation phenomena [7,10].

Two separate RANS codes are employed for this work: the Applied Fluid Technologies Incompressible Navier-Stokes (AFTINS) code, and NASA's OVERFLOW. The former is a purpose-built time-accurate incompressible flow code utilizing the discretization techniques presented in [5,6,8]. It has been extensively validated on a wide range of steady and unsteady problems [1,14,15,18]. OVERFLOW is a compressible code containing options for running at very low Mach numbers for steady state problems [3,13].

Both codes utilize the overset method to ease grid generation and improve solution accuracy. The approach allows grids to be developed around individual geometry components without regard to what happens in other parts of a grid. It maintains the accuracy advantages of structured grids, but allows many of the timesaving features of unstructured grids.

Figure 1 gives an example of surface grids generated using the overset technique. It depicts grid blocks around the intersection between a winglet, bulb, and bulb whale tail, and demonstrates how complex shapes can be easily resolved. Each color represents a separate block, but only faces corresponding to the yacht solid surface are shown.

The AFTINS code also includes capabilities for solving free surface problems. The method works by computing solutions to the nonlinear kinematic boundary condition during every RANS time step. Any grid face that contains free surface boundary conditions is projected onto this new free surface and re-intersected with the original hull definition. The hull is then




Figure 1: Typical Overset Surface Grid

truncated above the waterline, and a new grid generated. Boundary conditions for flow variables within the RANS code are provided by the zero stress condition. Pressure is set to atmospheric, and velocity gradients are combined in such a way as to yield zero shear.

The approach is more complicated than the Volume Of Fluid (VOF) method currently gaining popularity, although it does have advantages. The VOF approach, for example, doesn't really resolve a free surface. It interpolates it from solutions of a continuum variable. But like any discrete computational method, the variables representing a VOF free surface diffuse, and the free surface becomes "thick." Its exact location has to be interpolated out of the data, and the actual boundary conditions end up applied at some approximate location. The method can prove quite accurate, but requires substantial grid resolution near the free surface.

The main advantage a VOF method has over resolving a free surface is that it doesn't require regriding at every time step. Regriding using conventional tools can prove a very painful (and slow) process. With overset methods, however, this is no longer the case, and the VOF method loses many of its advantages.

Provisions for modeling sinkage and trim are incorporated (in AFTINS) by integrating forces on the hull at each time step. How far the yacht is out of balance is computed by comparing to input displacement and moments. Sensitivity gradients are computed for each degree of freedom (e.g. how trim moment changes with unit sinkage), and the resulting 3 by 3 system solved for the motions needed to balance.

### **Integrated Flow Prediction System**

While RANS codes are quite complex pieces of software themselves, they require equally sophisticated support packages to create a complete flow simulation system. Because they work on discrete definitions of a flow domain, RANS codes require a three-dimensional, curvilinear computational grid to model the flow space around a body. Generating this grid can be as difficult as obtaining the flow solution, but has to be repeated for every new geometry or run condition. (RANS codes, at least, only have to be written once.)

Gridding for complex geometries can be extremely demanding. It is not at all uncommon for expert users to require several days for meshing a domain. While this may prove acceptable for one-off flow analyses, such

throughputs will never satisfy the demands of time-critical design applications.

Grids that are developed from scratch have bigger problems than just the required man-hours. Every discrete solution to a continuum-mechanics problem contains some level of truncation error. While small, this error will inevitably start to look significant when compared to some of the smaller design trends. If grid generation is not carefully controlled, truncation error will take on a randomness that overrides any real differences in the flow.

Automated grid generation utilities have the advantage that truncation error-inducing grid attributes like cell size, skew, aspect ratio, expansion ratio, and number of points can all be controlled to insure every grid is similar. Truncation error may still be larger than the differences being investigated, but at least they won't be random.

Auto-gridders have an additional advantage for high throughput applications. Topologically similar grids will have similar stability requirements during RANS code convergence. Once one flow solution has reached convergence, it can be reasonably assumed they all will. Since instabilities in even one run of a design series delays the entire series, robust convergence of the RANS code is essential.

Unfortunately, achieving topological similarity is not always trivial. Unstructured grids utilize tessellations or advancing fronts that introduce randomness in their point spacings. This obviously works against any attempt to control grid quality. Unstructured grids also have difficulty placing sufficient points in the boundary layer for high Reynolds number applications.

Unstructured grids with prismatic boundary layer sections fare significantly better, but still have some randomness introduced when two-dimensional distributions of unstructured surface points are laid out to anchor the prisms.

Structured grids are fine for simple geometries, but often provide very high levels of truncation error as geometric complexity increases. Even a small amount of randomness in a large truncation error could overshadow some design trends.

The best alternatives for maintaining true topological similarity are overset grids [13,14, 16]. Overset grids utilize ordered hexahedra within any given block, but impose no

connectivity requirements between blocks. Communication between blocks is handled by conservative interpolation, and the only boundary restrictions are that the union of all blocks spans the computational domain. Since each block is generated using conventional structured gridding approaches, maintaining topological similarity is straightforward.

Overset grids have all the accuracy advantages of structured grids (e.g. boundary layer resolution), but all the simplicity advantages of unstructured. The domain about each geometry element is gridded without regard to surrounding geometry elements or resolution requirements. Gridding complexity does not increase with geometric complexity. The resulting grids tend to be more point efficient because regions of high resolution are not wasted on satisfying requirements of connectivity.

Finally, overset grids are easy to generate using automated techniques. Because each geometry element is independent, hyperbolic techniques can be used. These require only that the geometry surface discretization and its normal point spacings be specified.

The overset method requires some grid pre-processing before it can be properly utilized by a RANS code. Since the blocks were developed without any knowledge of adjoining grids (or geometry elements), some portion of the domain will be covered by more than one mesh. It is also possible that grid cells around one geometry element lie inside another geometry element, or even outside the flow domain.

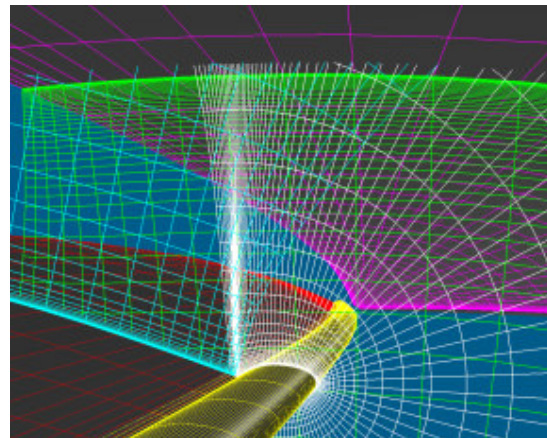
The process of removing redundant points from a grid is referred to as "blanking." Points that lie inside a closed body or outside the flow domain have to be blanked. Coarse cells that occupy the same volume in space as fine cells from another block should also be blanked.

Once blanking is complete, interpolation source cells have to be assigned so that one block can communicate with every other. Each non-blanked point overlapping into another block is identified, and has an interpolation partner chosen from the best candidate of all other blocks. Such points are treated as boundary conditions during the RANS solution, and updated from adjacent blocks after each time step.


The process is demonstrated in Figures 2 through 4, which show grid blocks being developed around the mast, mainsail, and Genoa. The view in each figure looks down from above, with the mast at bottom center and

the mainsail to the left. The Genoa windward surface shows near the top of each figure. Red lines depict the mainsail leeward side surface grid, purple the Genoa windward side surface grid, and yellow the mast surface grid. Blue lines depict one plane of the volume grid growing off the mainsail leeward side, green lines depict one plane growing off the Genoa windward side, and white depicts one plane growing off the mast.

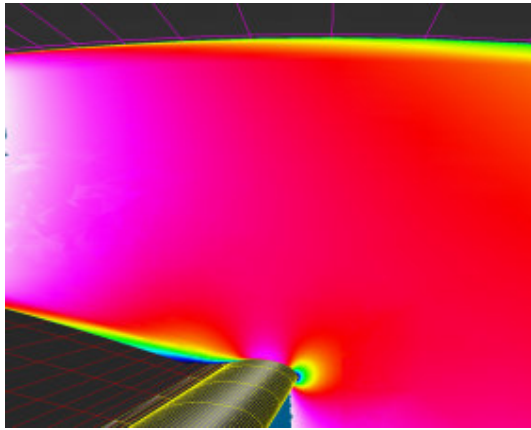
The situation in Figure 2 is confusing because the grid is shown before blanking. Green Genoa grid lines pass right through the mast and mainsail, and blue and white grids pass right through the Genoa. At some points between the main and Genoa, three grid blocks occupy the same space. Figure 3 shows the grid after blanking, and demonstrates that the best resolved grid cell is chosen for each point in space. It also shows that adequate overlap is allocated to allow interpolation without requiring data from boundary points. Figure 4 shows flow data on one plane of the grid colored by stream wise velocity. Note that interpolation has not disrupted solution smoothness.



**Figure 2: Overset Grid Before Hole Cutting**



**Figure 3: Overset Grid After Hole-Cutting**



**Figure 4: Flow Velocities on Overset Grid**

Once high-quality, topologically controlled grids are combined with a robust, accurate RANS package, only a few additional utilities are required to build practical systems for designer use. First, a capability must be provided to either generate geometry or convert it from CAD-provided form into usable gridder inputs. Second, input files must be generated to control grid resolution, blanking inputs, and RANS boundary conditions. Third, the RANS code has to be monitored during run time to check for convergence and runtime errors. Lastly, a post-processing utility must be provided to retrieve converged forces, convergence residuals, run time errors and messages, and any detailed flow data required for analysis.

Even with all these automated utilities, however, keeping hundreds of RANS analyses running is too much for most humans. One final utility is needed (in the form of UNIX scripts for the current example) to coordinate the required steps of matching input geometries to run conditions, starting grid generation and pre-processing, running the RANS code, and post-processing results. The utility must also locate available computational resources, move the required codes and geometries to that computer, perform the run, and notify the user when complete.

Two UNIX shell scripts are used to perform these final integration tasks. The first is referred to as the "run management script," and is responsible for managing the individual software elements required to complete a single RANS run. The second is called the "executive script," and is responsible for managing the complete set of analyses queued, monitoring available computer resources and launching individual run management scripts.

Running the executive script requires only a single input file containing one line of data for each required RANS analysis (i.e. unique design candidate at one test condition). The line contains a unique ASCII identifier for tagging output files and data. It also contains directory paths to each required geometry element (sails, mast, topsides for aero runs; or hull, keel, bulb, winglet, rudder for hydro runs). The line ends with yacht attitude information such as roll, yaw, heave, traveler or tab setting, and rudder angle.

The executive script monitors available CPU's on the host computer, and launches a job whenever resources become available. Once a particular run management script is queued, all individual components of that job (grid generation, pre-processing, RANS, post-processing) execute serially until complete. When a CPU frees itself, the next job (line) of the executive script control file is launched.

When all runs of the executive script input file have completed, the script executes a pre-arranged series of post-processing utilities. As a minimum, the three forces and three moments of each geometry element are collected for every run, and output to a single file. Utilities also exist to collect surface pressures and shears, chord wise integrations of force for lifting elements, and near-surface velocities for creating streamline plots. Any of these data can be consolidated into a single file for easy downloading and viewing.

### **Technical Approach -- Aerodynamics**

Geometries for aerodynamic studies; sail design exercises, or mast shaping efforts are all taken directly from their respective CAD packages. Sail shapes, for example, derive from either North Sails' DESMAN (when design shapes are tested) or MEMBRANE (when flying shapes are required). Mast, deck and topsides geometries are taken from IGES definitions of design or as-built shapes.

An automated grid generation system was developed to convert each geometry element into surface discretizations of prescribed resolution. The elements are then rolled, yawed, pitched, and heaved to values prescribed by the executive script input file.

The surface grids are output to file, and used as input for NASA's HYPGEN hyperbolic grid generator (4). The resulting volume grids are generated independently, so the process is both robust and simple. Because individual elements are gridded separately, the resulting meshes

exhibit very high quality, and therefore support accurate computation. Since surface resolution parameters and HYPGEN input spacings are all hard-wired, meshes of similar geometries are guaranteed to be geosyms. Grid-induced truncation error is thereby controlled, and successive results can be reliably compared.

Four blocks are used to resolve each sail; one on the windward side, one to leeward, one around the foot, and one wrapping up the leech and over the head (to resolve wakes). Three blocks are used to resolve a head foil or mast. Appropriate overlaps are established from the head foil to Genoa blocks, and from the mast to mainsail blocks. The deck and topsides require five more blocks, and a moderate resolution mesh is added encompassing the whole assemblage so that wakes are resolved far downstream. A final Cartesian far-field block is used to carry boundary conditions out to about five hull lengths. The complete grid typically requires close to three million grid points.

Normal spacing for any block face section that contains no-slip boundary conditions are rigidly controlled within HYPGEN. Since wall functions are avoided, near-wall resolutions of about  $y^+$  equal to 1 are specified. For America's Cup Class aerodynamics, this fixes the near-wall normal spacing at about 0.01 mm. All attempts are made to maintain normal wall stretching below a ratio of 1.2, which results in about 10 points being placed in the viscous sub-layer. Twenty-five points are typically placed across an entire boundary layer. The  $k\epsilon$  turbulence model option is used for all aerodynamics work.

Far field and water surface boundary conditions are prescribed to properly resolve the Earth's boundary layer. Since a boat-fixed grid is used, this necessitates a constant (but non-zero) velocity on the water. Three components of velocity are prescribed upstream, and introduce the required apparent wind vertical gradient and twist. Left, right, and top boundaries use conditions where the velocity direction is fixed by the Earth's boundary layer direction, but velocity magnitude is free to float. Downstream and pressure boundary conditions are all Neumann.

Once the grid and boundary condition setup files are available, the run management script starts the grid pre-processor to establish hole cutting and interpolation information. It then checks for grid errors and automatically launches the RANS run. On a 4-processor 833 MHz HP/Compaq ES40 (one-sixteenth of the 64-processor alpha cluster), the whole process takes between 12 and 24 hours.

When all the runs for a particular design candidate have converged, the executive script queues an automated post-processing utility to collect force and flow detail information from all nine data sets. The information is collated into a small number of files to be downloaded and studied by local designers. When all geometries of a particular design series are available, they get sent to the VPP for analysis.

This process can be repeated for any number of design variables (or combinations of variables). For the Oracle/BMW 2003 campaign, for example, the following design series were investigated:


- Genoa camber;
- Mainsail camber;
- Mainsail draft location;
- Genoa entry angle and curvature;
- Mainsail roach
- Mainsail twist and traveler settings;
- Genoa sheeting angle;
- Genoa leech hook angle;
- Standing rigging drag effect.

Each study typically contains a series of five to ten candidates. A few studies investigated multiple design parameters simultaneously, and required substantially more analyses.

Each candidate in a run series is analyzed over a range of heels and true wind angles selected to span the anticipated operating conditions for a particular design. Forces from each run in the range are fed to a VPP so that boat performance can be ascertained at the converged sailing angles and speed. A matrix of three heels and three wind angles is usually found sufficient.

Figure 5 shows flow field data from a typical RANS solution. Sails are colored by pressure jump (i.e. windward side pressure minus leeward side pressure), and wake cross flow cuts are colored by the stream wise component of vorticity. The wake data shows that the mainsail foot vortex system is strongest, and that the Genoa head vortex system is moderately strong. The mainsail head vortex is quite weak.


Figure 6 shows vertical distributions of drive force, side force, and roll moment resulting from a typical RANS run. Force areas and moment volumes (as opposed to coefficients) are used so that designs of different surface area can be compared. Genoa and mainsail loads are typically shown separately so that their interaction can be better studied. Loads from a corresponding potential flow simulation are included for comparison.



**Figure 5: Overview of Sail and Rig Aerodynamics (Pressure on Sails, Vorticity in Wake)**

mainsail head and under predict them near the Genoa foot. This is easy to understand near the mainsail head because angles-of-attack can be expected to increase when the Genoa ends. This will cause a viscous flow to stall and lose power, but not a potential flow.

The decrease in potential flow-predicted load at the Genoa foot is understandable if deck models are not included to mimic an end plate effect. It's interesting though; that the effect is seen (although to a lesser extent) even if the deck is paneled to provide an end plate (14). Figure 7 shows an explanation for this. Free stream flow is angled from the boat centerline (by the apparent wind angle), and separates at the deck windward edge. Flow to leeward of the rail rolls up into a vortex, and induces additional cross flow near the Genoa foot. If the vortex is resolved (as in a RANS simulation), it has the effect of increasing Genoa foot angle of attack, and therefore load.




**Figure 6: Vertical Load Distribution -- RANS versus Potential Flow**

This discrepancy poses an important question. Potential flow force predictions have long been used to drive VPP's, and are found to predict boat speed reasonably well. How can RANS be right if it doesn't agree with these well-validated panel codes?


The answer is demonstrated in Figure 8, which shows RANS versus potential flow total force predictions. It can be seen that potential flow under predictions of drive and side on the Genoa foot are almost exactly cancelled by over prediction at the mainsail head. The only significant difference is in roll moment, and is expected since potential flow predictions have the effect of moving load up the rig. The important finding here is that just because panel codes do a good job predicting boat speed doesn't necessarily make them the best tools for guiding detailed sail shape decisions.

Two observations are immediately obvious from the plots. The first is that a mainsail does very little near its foot. This is understandable since the sail is fairly flat over its bottom half, and because the Genoa has turned local flows to align (approximately) with the mainsail foot.

It is also obvious that RANS and potential flow load predictions differ substantially. Potential flow tends to over predict loads near the



**Figure 7: Effect of Rail on Genoa (Sails Colored by Pressure, Cross Plane Colored by Velocity)**



**Figure 8: Total Aero Force -- RANS versus Potential Flow**

### Technical Approach -- Hydrodynamics

Like the aerodynamic studies described above, geometries for all under-the-water components are taken directly from designers' CAD systems in the form of IGES entities. A range of attitude parameters (roll, yaw, heave, pitch, tab setting, and rudder angle) is provided for each candidate design, and grids are generated automatically using a system similar to the aero version. For hydrodynamic calculations though, the auto-gridder performs two additional steps. First, the tab and rudder are deflected to their prescribed angles. Second, the hull surface entity is trimmed off at the water surface. This last can be accomplished in one of two ways.

If potential flow calculations are available and wave resistance is not a RANS priority, the free surface shape can be taken from panel code results. Otherwise, the yacht's roll, heave, and trim are adjusted to provide the correct displacement and buoyancy moments, and the hull entity truncated at the still water level. Dynamic free surface boundary conditions are then applied during the RANS calculation, and kinematic conditions are used between iterations to update the evolving free surface.

Outputs from the auto-gridder come in the form of surface discretizations with prescribed resolution. These are used as inputs for the HYPGEN (4) volume gridder to generate independent meshes around each separate geometry element.

The resulting near-field grids consist of one block for the hull, one each for the keel and its wake, two for the bulb (the tail being gridded separately with greater resolution), and three each for the two winglets and rudder. Intermediate resolution and far-field boundary condition blocks similar to those used for aerodynamics are added last. Total grid size is around 3.5 million points.

Boundary conditions are straightforward since the far-field flow is now a simple free stream. Inflow is set based on boat speed; and left, right, and bottom walls are treated as no-penetration. Neumann conditions are again used for outflow and pressure boundaries. If wave shape has been fixed based on panel code results, the free surface is treated as a no-penetration face with prescribed pressure. Otherwise, the zero-stress free surface boundary conditions are solved within the RANS code, and the kinematic solver run after each time step to update free surface locations and grids.

Near-wall spacing follows a similar set of guidelines as described in the aerodynamics section. A similar run management script is also provided to perform corresponding functions for grid development, hole cutting, interpolation setups, and RANS run initiation. Both Spalart-Allmaras one-equation and  $k\epsilon$  two-equation turbulence models have been employed – although only one at a time is used during any complete design study.

Runs typically require between 18 and 20 hours using four processors of the HP/Compaq alpha cluster. Once all runs of a particular geometry are complete, the executive script launches an automated post-processor to collect and download forces and detailed flow outputs.

Oracle/BMW utilized this system for a wide range of hull and appendage studies. Design series were investigated in the following areas:

- Bulb shaping;
- Quantification of hull viscous drag;
- Hull form factor model development;
- Winglet placement;
- Winglet angle of attack and twist;
- Dillet, fillet, and strake design;
- Rudder performance.

Each candidate design was tested across two speed and attitude ranges: one representative of upwind sailing and one of downwind. The ranges were chosen to cover a realistic set of operating conditions, and vary heel, yaw, heave, tab setting, and rudder angle across with speed. Ten to fifteen operating attitudes were usually tested for each design. A design study typically included between five and ten candidates.

Since hydro analysis requires such a large number of operating points, performing the hydro equivalent of a 3 by 3 heel/TWA matrix at each point would prove impractical. Rather than using the VPP to predict equilibrium sail angles, designs are ranked by interpolating performance results to a constant side force or speed. In instances when seconds per mile comparisons


are necessary, the method of “linear deltas” is employed. The technique assumes that benefits (or penalties) due to unit changes in each of the forces or moments can be computed independently, and then added together.

Figure 9 shows flow field data from a typical hydrodynamic RANS solution. Contours of constant stream wise vorticity on one cross flow plane are shown in the yacht's wake. Streamlines are included to show where individual vorticies originate. The results show just how complex the keel/bulb/winglet flow is, and demonstrate that not all vorticity shed from a keel ends up absorbed as winglet lift.


Other examples of using RANS to highlight design detail can be found in Figures 10, 11, and 12. Figure 10 shows streamlines near the underside root of the leeward winglet. The winglet and bulb are colored by pressure, and indicate a relatively high winglet load (i.e. steep suction gradients along the leading edge). Despite the load, however, the streamlines show little sign of separation, and would seem to indicate a sound design.

Loads for three different winglet concepts are shown in Figure 11. Design 3 shows the highest leeward side loading and would therefore do the best job at increasing effective keel span. Since windward side lift can actually decrease performance, these plots could be used to guide a search for wings that work only where needed.

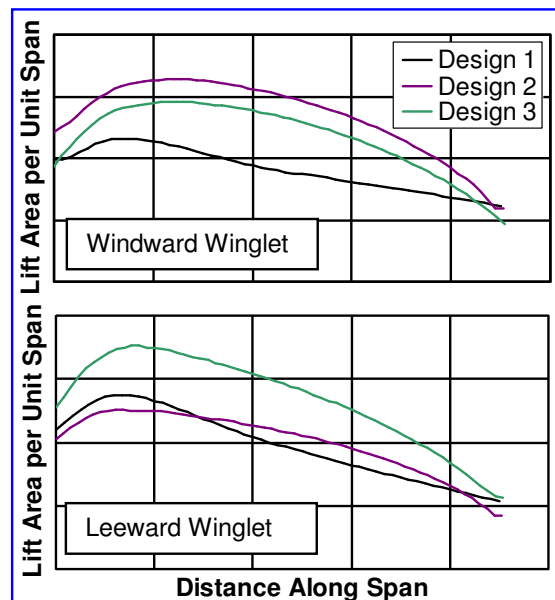
Figure 12 shows rudder lift/drag polars, and demonstrates the extreme asymmetry caused by the keel vortex. The symbols on each curve represent calculations made at the same angle-of-attack for each rudder. A designer can use this information to maximize rudder efficiency while minimizing any chances of stall.



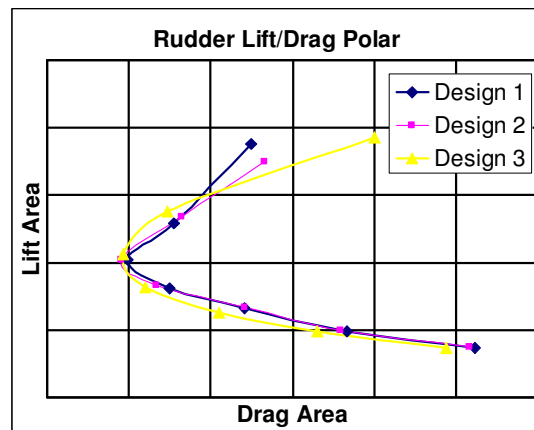
**Figure 9: Overview of Hydrodynamic Flow (Contours of Streamwise Vorticity)**



**Figure 10: Detail of Winglet Juncture Flow (Surface Pressure and Steamlines)**



**Figure 11: Winglet Loading**



**Figure 12: Asymmetric Rudder Polars**

## Design Example 1 – Genoa Camber

The optimization of Genoa camber makes for an excellent first example for the use of RANS in design. Boat performance is fairly sensitive to camber, but potential flow methods are known to provide potentially misleading results. Since changing Genoa camber will cause load tradeoff between fore sail and main, any employed analysis technique must resolve significant force distribution detail. It must also be trusted to accurately trend form drag as camber is altered.

The RANS study is formulated starting with a series of baseline designs covering a range of True Wind Speeds (TWS). Three sail code ranges are selected for the present example: 7 to 11 kts TWS; 10 to 14 kts TWS, and 13 to 17 kts TWS. For short, the sails will be referred to as the 9-kt sail, 12-kt sail, and 16-kt sail.

Optimization begins by designing sails in each wind range that are one step flatter and one step fuller than the baseline shapes. A range of heel angles and True Wind Angles (TWA) is then selected to bracket the expected sailing angles. Typically, a 3 by 3 matrix of heels and TWA's is found sufficient. RANS analyses are then performed for each design at every point in the matrix. For the case of 3 wind ranges, 3 cambers, 3 heels, and 3 TWA's a total of 81 RANS calculations are required. On a 64-processor 833 MHz HP/Compaq Alpha Cluster, this requires about 5 days of CPU time.

The resulting forces are fed into a VPP, and converged sailing angles and speeds obtained. The results are shown in terms of seconds per mile in Figure 13. Note that since "faster" means "less seconds," the better sails sit lower on the plot. The curves indicate that camber of the 9-knot baseline design needs to increase, but that camber of the 12-knot sail needs to decrease. The 16-knot sail needs to be fuller in its lower TWS range, but flatter at the upper end. This may indicate the sail is being expected to operate over too broad a wind range.

Figure 13 also does a good job showing how performance changes as a sail gets away from its design point. Curves with steeper slopes (like the 9-knot sails) indicate that performance degrades more rapidly than sails with flat curves. An obvious design goal should be to achieve the flattest possible curve. Finally, Figure 13 could be used to help guide selection of sail crossover points. If a crossover falls at a particularly bad wind speed, for example, these results could be used to guide the selection of new wind speed ranges.




Figure 13: Genoa Camber Study -- Optimum Camber and Crossovers

If boat speed is considered the "bottom line," calculation can stop and sail construction start. But if more detail is required (to guide the next design step for example), a final analysis set is required. The VPP places each sail in a slightly different attitude, so it is difficult to compare detail at any one point in the heel/TWA matrix. The sails must be reanalyzed at the respective sailing points, and new force data extracted. The nine additional runs take a day of CPU time.

Vertical load distributions from the 12-kt sails at their VPP point are shown in Figure 14. Genoa and mainsail loads are shown separately, and the fastest sail is indicated by attached symbols. The drive force plot (top) shows that the flat Genoa actually has the least drive force. Interestingly, however, mainsail drive increases with decreasing Genoa camber, and the two just about cancel out. Total drive force is relatively insensitive to Genoa camber.

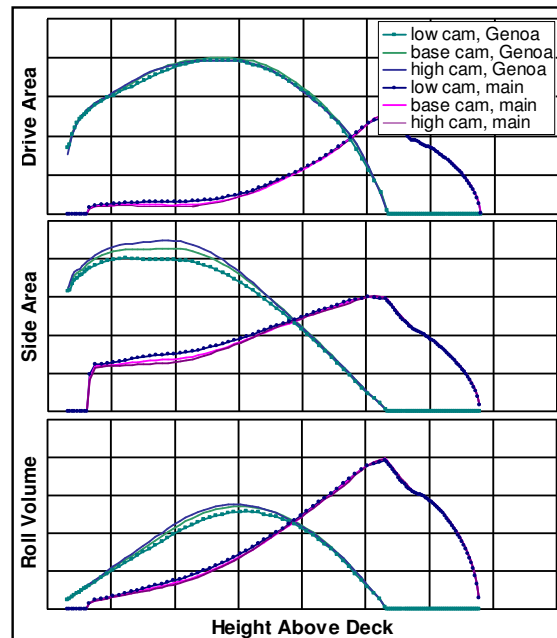



Figure 14: Genoa Camber Study -- Vertical Load Distribution



**Figure 15: Genoa Camber Study -- Main Effect**

Side force is a different story, and the flatter Genoa shows substantially less of it. Mainsail side force does increase (since mainsail load increases with decreasing Genoa camber), but to a much lesser extent. Total side force ends up as a reduction. Roll moment also reduces (although only slightly) – and is mostly attributable to the reduction in Genoa load.

For the 12-knot case then, a flat sail is best because it produces less side force for the same drive. Many sail design trends are trickier, and can only be deciphered with the aid of a VPP. If both side and drive decrease, for example, conclusions cannot be drawn based solely on changes to drive and side forces.

This interaction between Genoa and mainsail load distribution indicates that one sail cannot be optimized without taking into account the other. The point is borne out by the results of Figure 15, which show the 12-knot Genoas combined with various depth mains. Values on the vertical axes represent changes in seconds per mile from the “baseline main, baseline Genoa” case (negative being faster). Note that the “baseline main, flat Genoa” is the fastest sail from Figure 13. The curves show that while the flatter Genoa is indeed fast, it is almost universally better when combined with a flatter main. Only at the lowest end of its wind speed range does a flat/flat combination become underpowered.

**Design Example 2 – Mast Tube Shaping**

Design of the mast/mainsail combination is one of the most aerodynamically challenging issues facing an America’s Cup team. Placing a bluff mast at the leading edge of a thin airfoil is obviously disruptive to the flow. But minimizing the disruption has to be traded off against structural integrity and sail control. As a mast’s section is reduced, its ability to support a straight luff above the forestay is compromised. The mainsail will twist off, and possibly lose more performance than gained by the smaller mast. A design approach has to account for this.


The problem is further complicated by substantial flow differences above and below the Genoa head (sometimes referred to as the “hounds”). Mainsail angle-of-attack below the hounds is fairly low, and flow reattachment is assisted by sail interaction (e.g. the slot effect). Unfortunately, the angle-of-attack increases dramatically above the hounds, and the slot effects disappear just when most needed.

Modeling mast flow physics is therefore complex. None of the essential flow attributes are resolved by potential or inviscid flow methods. Wind tunnel tests are difficult due to wall effects, and can be misleading due to Reynolds scaling. Even two-dimensional techniques aren’t applicable since the flow has huge span wise components. Only three-dimensional RANS analyses can produce optimal designs.


An approach similar to that employed for Genoa camber is utilized. A series of masts is created that vary shape parameters like leading edge radius, trailing edge thickness, and length to width ratios. Each section is designed so as to provide similar structural and weight properties.

A series of RANS simulations is performed to spans the possible range of true wind angle; roll angle, and traveler setting. The resulting forces are fed into a VPP, and the boat’s actual sailing angles and speed computed. The process is repeated for each design candidate, and the best boat speed option is selected. Some trade offs are required because the best mast at low apparent wind speeds may not be the best at higher wind speeds. Differences in boat speed were found to be substantial for some designs.

Surface flow patterns between two of the candidate designs are shown in Figures 16 (mainsail head leeward side streamlines), and 17 (windward side surface streamlines). The left side of each figure shows the flow on Design 1, whereas the right side depicts Design 2.



**Figure 16: Mast Section Study – Leeward Flow**



**Figure 17: Mast Section Study -- Windward Flow**

The large concentration of leeward side streamlines visible approximately half way between the hounds and head marks the extent of attached flow. Flow above that line is filled in by the tip vortex rolling over from the windward side. Free stream flow (at least this close to the sail) never reattaches after separating from the mast. Flow below the concentration line is provided by attached flow entering the head region from lower on the sail. Although both designs exhibit significant vertical flow, Design 2 shows substantially more leeward side attachment than Design 1.

The effect to windward is similar. Now though, the vertical concentration of streamlines just aft of the mast depicts the reattachment point. Flow aft of that point reattaches and flows back (and up) on the sail. Flow forward of that point spirals up the mast to fill in the separation bubble from below. Note that the separation bubble on Design 2 is smaller than Design 1. Due to the improvement on both sides of the sail, Design 2 proves substantially faster.

It is interesting to note that the same study was attempted using two-dimensional RANS simulations. Unfortunately, the severe span wise flow makes these 2D analyses suspect and misleading. The resulting performance predictions differ wildly from those in three-dimensions, and section shape is actually guided in the wrong direction.

### Design Example 3 – Bulb Shaping

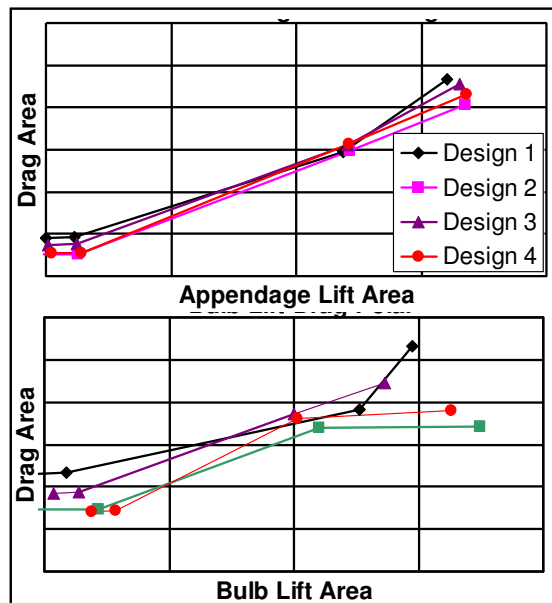
Of all components on an America's Cup yacht, perhaps none allows more design latitude than bulb shape. For a given volume of lead, a bulb can be either long and narrow or short and fat. Its cross section can be symmetric, oval, or even pear shaped (to concentrate lead at the bottom). When the multitude of nose and tail shapes are considered (to say nothing of winglet style, angle, and location), one quickly sees there are more options than can be tested by towing tank.

All these design options involve trade-offs between center-of-gravity, viscous drag, and form drag. Bulb shape also affects winglet efficiency, and is therefore closely tied to overall appendage performance. RANS is the perfect tool for investigating these trade-offs.

Finding an optimal bulb requires a huge number of RANS analyses. Each shape style must be considered, and a range of options investigated across that style. Every candidate must be tested for upwind performance (keel moderately loaded), downwind performance (keel lightly loaded), and acceleration (keel heavily loaded). RANS-predicted lifts and drags have to be input into a VPP for ranking options on boat speed.

A truncated example of one such study is shown in Figures 18 and 19. Four designs are considered that represent different section shapes. Each is tested at representative upwind, downwind, and acceleration points. Force results are shown in the form of lift/drag polars in Figure 18, and presented in terms of force area so that designs of different surface area can be compared.

The upper portion of Figure 18 depicts performance of the complete appendage set (keel, bulb, winglets), whereas the lower isolates just bulb performance. Lightly loaded cases fall to the left (low lift area) and highly loaded cases to the right. For a given load condition, all designs are tested at the same combination of heel, yaw, and tab setting. The polars indicate that lift is as much affected by section as drag.



**Figure 18: Bulb Design Study -- Forces**

If just bulb performance is considered, Design 1 is seen to be consistently poor. Designs 2 and 4 are consistently best, with one being optimal at light load and the other at moderate and heavy load. The importance of keel/bulb/winglet interaction is demonstrated by looking at the upper portion of Figure 18. Design 2 is now consistently better than 4, but Design 1 suddenly appears best at the upwind sailing point!

Figure 19 shows VPP-predicted boat speeds in terms of seconds per mile faster or slower than Design 1 (negative being faster). If a design were to be heavily weighted for steady-state upwind sailing, Design 1 might be the optimal choice. Design 2 appears to be a good all-around choice, but design decisions would depend on how much upwind performance is actually sacrificed.

One aspect of bulb performance not included in this study is laminar flow. The relatively large bulb surface area, combined with its remoteness from the free surface, make laminar flow more probable on the bulb nose than anywhere else. Up until recently, though, no techniques were available for predicting the extent of laminar flow on complex bodies. Tolmen-Schlichting instability techniques cannot account for cross flow effects, and experimental techniques cannot reach yacht Reynolds numbers.

All studies presented so far utilize RANS calculations that fix transition to a presumed location. Fortunately, however, a new technique for predicting high-Reynolds number transition on complex shapes is being tested by Imas [12]. The approach uses flow data from a RANS solution to construct Dynamic Sub-Grid Scale (DSGS) turbulence models similar to what would be used in a Large Eddy Simulation (LES). The developed models are used to locate regions in the flow field where small turbulent scales feed larger ones, and therefore predicate unstable boundary layer situations. Initial tests with the method have shown considerable promise.

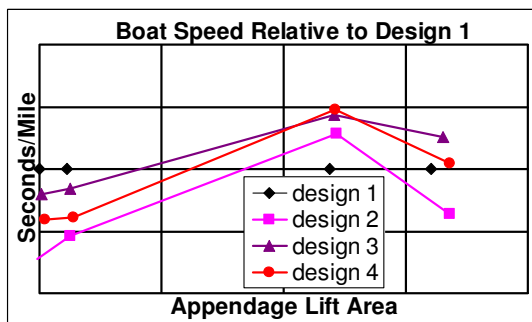


Figure 19: Bulb Shape Study -- Boat Speed

## Design Example 4 – Hull Form Factors

Traditional model test procedures use Froude scaling to apply tow tank-measured wave drag to full scale. Viscous effects, which do not scale in this way, are extrapolated using flat plate skin friction data. Since sailing yachts are not flat plates, they develop “form drag” as stagnation energy decreases in the boundary layer.

Form drag is usually predicted using a multiplier on flat plate drag. The multiplier is derived once at the beginning of a test, and applied for every model and every test condition. It seems obvious, however, that beamy or deep hulls have higher form drag than narrow ones. A yawed yacht will also have more form drag than one sailing straight. In fact, there’s no reason to think model form factors even apply at full scale.

A more precise approach can be developed using RANS codes. A given hull is analyzed at both model and full scale Reynolds numbers. A third run is made on the same grid, but with all no-slip boundary conditions removed to mimic inviscid flow. These last results give wave-making resistance, and can be subtracted from RANS-predicted values to yield total viscous drag. Total viscous drag contains both skin friction and form drag, and when divided by a flat plate values yield the factor. Both model scale and full-scale form factors are derived, and lead to more accurate extrapolation procedures.

This process was repeated over a range of hulls (varying beam, depth, prismatic coefficient, etc.), and over a range of test conditions (speed, yaw, heel). The resulting form factors are shown as a scatter plot in Figure 20. When regressed against design or test parameters, the data shows trends that allow models to be derived. These models are used to predict a form factor for any combination of beam, prismatic, heel, yaw, and scale. Model scale values are used to extract wave-making resistance from tow tank tests, and their full-scale counterparts are used to extrapolate total yacht drag.

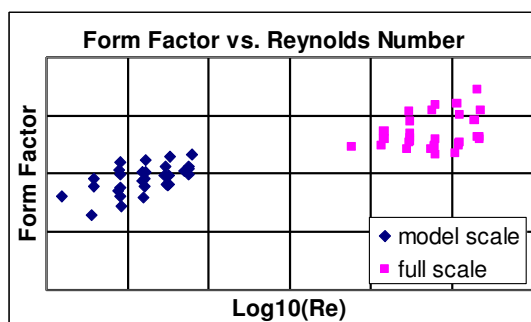


Figure 20: Hull Drag Study -- Tow Tank Data Extrapolation

## Future of RANS in AC Design

The level of productivity demonstrated herein indicates that RANS can be expected to play an expanding role in America's Cup design. Even as recently as the 2003, however, most publicly available literature on the subject described applications limited to relatively small numbers of mostly qualitative analyses [2,9,14,21].

Oracle/BMW's applications were unique in their ability to analyze large numbers of design options over a range of operating conditions. Yet the demonstrated applications are still limited to "design analysis" rather than design development. Even the best RANS simulations are still restricted to ranking a given set of designs rather than developing new ones. So the question remains, how can RANS transition from a design analysis role into an actual design tool?

One promising concept is to take advantage of the more robust and automated RANS codes to build "black box" analysis tools. The ability to go from geometry, to grid, to RANS forces, to VPP analysis (without human intervention) is 90% of what's needed for true design optimization. Once coupled with a method to automatically generate input geometries, a black-box tool could be set to work under the control of non-linear optimization executive programs.

Given the throughput already demonstrated, a point design could potentially be optimized in just a few weeks. This admittedly is a lot of compute cycles, but not necessarily outside the realm of current design time frames. The only foreseeable complication is that point designs are not likely to win an America's Cup. Optimization objectives must be crafted to cover a range of operating conditions. The resulting runs could require longer than a few weeks.

A second promising future application is to create RANS-based VPP aero-force databases. Most existing capabilities utilize combinations of test and potential flow data, and therefore exhibit the limitations of those technologies. Since virtually every design decision is governed by VPP, it makes sense to use the best data.

Two additional applications where RANS could play a future role are wave resistance and laminar flow predictions. While the first of these may at first seem like overkill (the viscous/free surface interaction of an AC yacht is quite weak) RANS might actually prove to be a significant break through. Most potential flow analysis codes are limited in their ability to resolve small details. But small changes in bow shape have

recently been shown to have large influences on performance. New techniques becoming available in the laminar transition area could also have an impact in AC design. Prediction of transition on complex shapes at high Reynolds numbers is no longer an unimaginable goal (12).

## Conclusions

The procedures and examples presented herein demonstrate that Reynolds-Averaged Navier-Stokes (RANS) simulations have become an indispensable part of the designer's toolkit. To truly support the design process though, any incorporated RANS capability must demonstrate sufficient throughput to handle a large number of design alternatives within reasonable turn-around times. It must also resolve very small differences in flow and performance. Finally, it must have sufficient functionality and robustness to enable its seamless integration with traditional design decision support codes such as Velocity Prediction Programs (VPP's).

The first requirement – that RANS be able to handle very high throughput rates – is important if RANS is to play a central role in ranking the large numbers of options facing an AC designer. It requires far more than a big computer. More importantly, it must:

- Incorporate simple (e.g. overset) gridding techniques to handle complex geometries without a need for human mesh developers;
- Consist of robust grid and CFD software able to complete hundreds of simulations without stability or accuracy problems;
- Contain extensive automation of problem setup and run management to minimize requirement for human intervention;
- Be applied by users experienced with both the software and the AC design process.

The second requirement – that RANS resolve very small differences – is crucial since yacht design has reached a very advanced state of technology. No one design attribute will be responsible for large differences in boat speed, and success will derive from a large number of small performance gains. Resolution of small differences requires that:

- Geosym grids be provided (e.g. by overset griders) to minimize mesh dependencies;
- Structured (or at least prismatic) near-wall grids be used to provide sufficient boundary layer resolution;
- Advanced turbulence models (e.g. one-or two-equation approaches) are used to insure accurate boundary layers and wakes;
- Wall functions are avoided so that separated flows are accurately resolved.

This paper demonstrates that the presented RANS capability is able to meet these demands. It has shown that RANS can be used in practical design environments and drive traditional design support codes. It can identify real performance improvements (both above and below the water), and resolve speed differences as small as 0.1 seconds/mile (0.03%). Used wisely, it can identify substantial cumulative performance gains. Finally, the paper has presented evidence that RANS will eventually evolve from today's design analysis tool into an even more practical methodology for generating optimal designs.

## References

- (1) Amromin, E., Miaine, I., Cook, L., Day, W., and Korpus, R., "High Speed Trimaran Drag: Numerical Analysis and Model Tests," *J. of Ship Res.*, to be published., 2004.
- (2) Azcueta, R., "RANSE Simulations for Sailing Yachts Including Dynamic Sinkage and Trim and Unsteady Motion in Waves," High Performance Yacht Design Conf., Auckland, 2002.
- (3) Buning, P.G., Jespersen, D.C., Pullium, T.H., Klopfer, G.H., Chan, W.M., Slotsick, J.P., Krist, S.E., and Renze, K.J., "OVERFLOW User's Manual," NASA Langley Research Center,
- (4) Chan, W.M., Chin, L.T., and Buning, P.G., "User's Manual for the HYPGEN Hyperbolic Grid Generator and the HGUI Graphical User Interface," NASA TM 108791, 1993.
- (5) Chen, H.-C., and Korpus, R., "A Multi-Block Finite-Analytic Reynolds-Averaged Navier-Stokes Method for 3D Incompressible Flows," ASME Summer Fl. Conf., 1993.
- (6) Chen, H.-C., Patel, V., and Ju, B., "Solutions of Reynolds-Averaged Navier-Stokes Equations for Three-Dimensional Incompressible Flows," *J. Comp. Physics*, 1989.
- (7) Chen, H.C., and Patel, V.C., "Near-Wall Turbulence Models for Complex Flows Including Separation," *AIAA Journal*, Vol. 26, No. 4, pp. 641-648, 1988.
- (8) Chen, H.C., and Patel, V.C., "The Flow Around Wing-Body Junctions," *Proc. 4<sup>th</sup> Sym. on Num. And Phys. Aspects of Aero. Flows*, 1989.
- (9) Cowles, G., Parolini, N., and Sawley, M.L., "Numerical Simulation Using RANS-Based Tools for America's Cup Design," 16<sup>th</sup> Chesapeake Sailing Yacht Sym., 2003.
- (10) Fan, S. Lakshminarayana, B., and Barnett, M., "Low-Reynolds Number k-e Model for Unsteady Turbulent Boundary Layer Flows," *AIAA J.* v. 31, n. 10, p. 1777-1784, 1983.
- (11) Hanjalic, K., and Launder, B.E., "Sensitizing the Dissipation Equation to Irrotational Strains," *ASME J. of Fluids Engineering*, Vol. 102, 1980.
- (12) Imas, Len, "A Model for Prediction of Boundary Layer Transition Around Sailing Yacht Appendages," Madrid Diseno Yates, 2004.
- (13) Jespersen, D.C., Pullium, T.H., and Buning, P.G., "Recent Enhancements to OVERFLOW," AIAA 97-0644, AIAA 35<sup>th</sup> Aero. Sci. Mtg., 1997.
- (14) Jones, P., and Korpus, R. "International America's Cup Class Yacht Design Using Viscous Flow CFD," Trans. 15<sup>th</sup> Chesapeake Sailing Yacht Symposium, 2001.
- (15) Korpus, R., Jones, P., Oakley, O., and Imas, L., "Prediction of Viscous Forces on Oscillating Cylinders by Reynolds-Averaged Navier-Stokes Solver," Tr. Intl. Soc. Offshore & Polar Eng., 2000.
- (16) Korpus, R., Hubbard, B., Jones, P., Stromgren, C., Bennett, J., "Hydrodynamic Design of Inte-grated Propulsor/Stern Concepts by Reynolds-Averaged Navier-Stokes Technique," Tr. 7<sup>th</sup> Intl. Symp. Practical Design of Ship & Mobile Units (PRADS), 1998.
- (17) Korpus, R., and Falzarano, J. "Prediction of Viscous Ship Roll Damping by Unsteady Navier-Stokes Techniques," OMAE, v. 1, Offshore Tech. ASME, pg. 241-247, 1996.
- (18) Kumarasamy, S., Korpus, R., and Balow, J., "Computation of Noise Due to the Flow Over a Circular Cylinder," Proc. 2<sup>nd</sup> Comp. Aero-acoustics Work-shop on Benchmark Problems, NASA Conf. Pub. 3352, pp. 297-303, 1997.
- (19) Schneider, A., Amone, A., Savelli, M., Ballico, A., and Scutellaro, P., "On the Use of CFD to Assist with Sail Design," 16<sup>th</sup> Chesapeake Sailing Yacht Sym., 2003.
- (20) Spalart, P.R. and Allmaras, S.R., "A One-Equation Turbulence Model for Aerodynamics Flows," AIAA 92-0439, AIAA 29<sup>th</sup> Aerospace Sciences Meeting, 1992.
- (21) Richter, H.J., Horrigan, K.C., and Bruan, J.B., "Computational Fluid Dynamics for Downwind Sails," 16<sup>th</sup> Chesapeake Sailing Yacht Sym., 2003.
- (22) Weems, K., Korpus, R., Lin, W.-M., and Fritts, M., "Near-Field Flow Predictions for Ship Design," 20<sup>th</sup> Sym. on Naval Hydro., 1994.