




THE 18th CHESAPEAKE SAILING YACHT SYMPOSIUM

ANNAPOLIS, MARYLAND, MARCH 2007

Performance Prediction without Empiricism: A RANS-Based VPP and Design Optimization Capability

Richard Korpus, Applied Fluid Technologies, United States



ABSTRACT

The use of Velocity Prediction Programs (VPP's) in sailing yacht design has been standard practice for years. VPP fidelity, however, continues to be limited by the accuracy of aero and hydro force data used to represent a particular yacht. Even the most advanced America's Cup VPP's usually derive sail forces from panel or vortex-lattice models, and hull forces from potential flow codes or experiment. Real world effects attributed to viscosity are added using simplified theoretical or empirical models that cannot resolve all the complexity of sailboat physics.

This paper describes a new approach for performance prediction and design optimization that relies solely on high-resolution Reynolds-Averaged Navier-Stokes (RANS) computational fluid dynamics. All aero and hydro forces and moments are generated by RANS, and therefore include the real-world flow features of boundary layers, separation, shed vorticity, and turbulence. RANS software and grid model requirements suitable for VPP applications are discussed, and sample aero and hydro solutions included. Examples from America's Cup design are used to demonstrate the technique's practicality and accuracy. Finally, since VPP's require forces from a large number of sailing conditions, the extensive development effort (undertaken through three America's Cup cycles) to transition state-of-the-art RANS into the practical realm is summarized.

NOTATION

γ	Leeway angle
ϕ	Heel angle
θ	Rudder angle
ρ	Density
Ω	Objective function for design optimization (usually VMG)
AWA	Apparent Wind Angle
AWS	Apparent Wind Speed
$d_1 d_2 d_3 d_4$	Arbitrary design variables specified by user to define an optimization process
F_{Xaero}	Component of aero force along boat track
F_{Yaero}	Component of aero force to starboard of track
F_{Xhydro}	Component of hydro force along boat track
F_{Yhydro}	Component of hydro force to starboard of track
M_{Xaero}	Component of aero moment about boat track
M_{Zaero}	Component of aero moment about vertical
M_{Xhydro}	Component of hydro moment about boat track
M_{Zhydro}	Component of hydro moment about vertical
Q_{aero}	Dynamic head for aero ($\rho_{air} \times AWS^2$)
Q_{aero}	Dynamic head for hydro ($\rho_{water} \times V^2$)
SFF	Side Force Factor (the combination of leeway, tab, and rudder used for a RANS simulation)
TWA	True Wind Angle
TWS	True Wind Speed
V	Boat speed
VMG	Velocity Made Good

INTRODUCTION

Performance-driven design requires that speed (not drag or some other design parameter) be used as a measure-of-merit. One of the most important tools in a designer's arsenal is therefore the Velocity-Prediction Program (VPP) – a numerical tool for calculating how design changes impact boat speed (Oliver, et. al., 1987, von Oossanen, 1993). In its simplest manifestation the VPP is a straightforward balance between a yacht's aerodynamic and hydrodynamic forces. In reality, however, the problem is more complicated because neither of these force sets is easy to estimate. The reason some researchers build entire careers in VPP development is not because a force balance is particularly difficult, but because obtaining accurate forces *is*. Even to this day VPP's are limited by our ability to provide adequate aero and hydro data (Jackson, 2001).

Commonly used VPP's (e.g. Claughton, 1999, Oliver and Claughton, 1995) derive their forces from a wide array of sources: towing tanks; wind tunnels; curve fits to standard series data (Keuning, 1999); and potential flow Computational Fluid Dynamics (CFD) (Rosen, et. al., 2000, and Euerle and Greeley, 1993). Each source has its strengths, but also its weaknesses. Experiments, for example, are perhaps best at incorporating real world details like wave breaking and unsteadiness, but are limited by the inaccuracies of scale effect and measurement error. Potential flow is fast, but cannot resolve turbulence, separation, or vortex wake detail.

No currently available VPP can provide the needed forces based on first-principles physics. Because aero and hydro forces are affected by waves, viscosity, turbulence, separation, and vortex wakes; the most appropriate form of CFD is Reynolds-Averaged Navier-Stokes (RANS). But while RANS techniques are unique in their ability to resolve complex flow attributes, they have (until recently) been too inefficient to provide the quantity of data needed to power a VPP.

Demonstrations of RANS' ability to assist yacht aero- and hydrodynamic design are becoming almost common (Cowles, 2003, Jones and Korpus, 2001). But documentation proving RANS is sufficiently practical for a time-critical design environment is still scarce. This is because VPP usage places great demands on the accuracy and quantity of force data, and therefore taxes both robustness and efficiency. Success depends more on how a RANS code is applied than which particular code is used. A high level of automation is required if the method is to prove efficient, and the ability to resolve very small design differences is essential to its very justification. Only very recently has RANS demonstrated the ability to meet all these demands simultaneously (Korpus, 2004).

This paper presents a new VPP built specifically to utilize high-resolution force data, and demonstrates its

application by employing the RANS methodology described in Korpus (2004). Aero and hydro forces are provided using a large number of RANS simulations to span the range of yacht sailing attitudes. Four degrees of freedom are included to solve for boat speed, side force generation (leeway), roll angle, and rudder angle.

It will be shown that the number of runs typically required to "VPP" a given hull or sail design can be obtained in approximately one week. While this is less lead-time than most model tests, it is substantially more than potential flow methods. So while RANS-based VPP's might be considered "practical" for some applications, they will certainly not prove so for all. The current method is therefore offered not as a replacement for existing VPP's, but as a supplement for applications with longer lead-time and more stringent resolution requirements.

JUSTIFICATION

Many designers are by now familiar with the valuable insights provided by RANS. But most of these are also aware that RANS can be a time consuming tool to apply. So, other than for the elegant physics resolved by RANS, why would a designer bother at all? The obvious answer is that for many applications, they should not. RANS is a "big hammer" approach not suited for all projects. If used for designing non-performance oriented cruising yachts, for example, RANS could be likened to "swatting a fly with a Howitzer." Even if the project had sufficient time, the cost could not be justified.

For design programs with more demanding levels of refinement and longer lead times (e.g. America's Cup, Volvo, TP52, performance maxi-yachts), RANS may be the only tool able to provide sufficiently detailed force predictions. America's Cup teams, for example, search for design advantages as small as 0.1% of boat speed. Such small differences cannot be resolved by experiment or (where affected by viscosity) potential flow methods.

A second shortcoming to existing VPP force predictions is their reliance on ad hoc modification factors. Because they are developed to work (by necessity) with often-insufficient input data, VPP's require empirical corrections to ensure convergence (e.g. "reef" and "flat" sail force corrections). These factors often mask the very design changes being studied, and in some instances should be avoided (Jackson, 2001). A RANS-based VPP uses forces obtained only from first-principles relations between a design and its surrounding flow field.

The following four examples are included to demonstrate when a RANS-based VPP is justified. All derive from previous America's Cup work, and are further described in Korpus (2004) and Jones and Korpus (2001).

One of the most obvious viscous flow problems is separation around the mast (and the associated effects on

mainsail design). Figure 1 depicts a RANS solution over the upper mainsail leeward side, and shows a massive amount of separation. The thick red, white, and blue ribbons show off-surface flow whereas thin lines represent surface limit streamlines. The red ribbon (originating closest to the luff) spirals up behind the mast and highlights the severity of mast separation. Surface limit streamlines on the upper sail progress down and forward (rather than up and aft), and demonstrate the extent of stall. Potential flow methods cannot accurately predict forces in this case. Wind tunnels cannot reach full scale Reynolds numbers (to guarantee full-scale similitude), and require purpose-built facilities to mimic the twist in a sail's distribution of AWA.

Figure 2 shows a typical vertical distribution of sail drive force, and demonstrates the differences between RANS and potential flow predictions. Four curves are included: one for each sail (main and Genoa); and one for each computational method. It is seen that RANS predicts more drive than potential flow near the Genoa foot, but less near the mainsail head. Note that the area under any curve is representative of the total force produced by that sail, and that this allows an interesting observation. The sum of drive over the two potential-flow sails is approximately equal to that of the RANS sails. So while a VPP would predict roughly equal speeds using the two methods, potential flow would not correctly resolve individual sail design differences near the Genoa foot or mainsail head.

Justification for RANS-predicted forces can also be found under the water. Figure 3 shows the distribution of shear stress over a typical America's Cup hull and appendages set. Note that the distribution does not mimic a flat plate boundary layer – the method used to derive shear drag by potential flow and model test techniques. Note also that asymmetries in shear are attributable to asymmetries in pressure, and therefore indicate that the hull is generating side force. The asymmetries are greatest near the keel and rudder, and demonstrate a concept referred to

as “lift carry-over.” Potential flow methods represent both effects using ad hoc methods that cannot be expected to resolve small design differences. Since drag is due in large part to shear, and since hull side force can be as much as 10% of the total, ignoring these effects can be devastating.

The final example is highlighted by the middle figure of this paper's front piece. It depicts the wake of the keel, bulb and wings, and demonstrates that rudder inflow is quite complex. Predictions of rudder lift and drag that do not account for this complexity will be inaccurate.

Even these examples, however, would prove insufficient to justify the time and cost of RANS in its usual manifestation. The problem arises because not one or two solutions are needed to run a VPP, but one or two hundred. Most RANS codes are meant for application in a research environment, and cannot reach this level of throughput. The RANS system applied herein was developed, validated, and improved over three America's Cup cycles. It has been integrated with geometry, gridding, and post-processing support codes to achieve a high level of automation and efficiency, and is therefore uniquely suited for VPP applications.

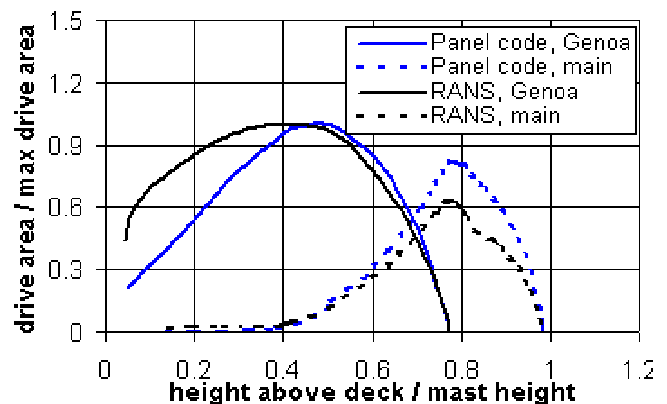


Figure 2: Comparison of RANS and Potential Sail Force

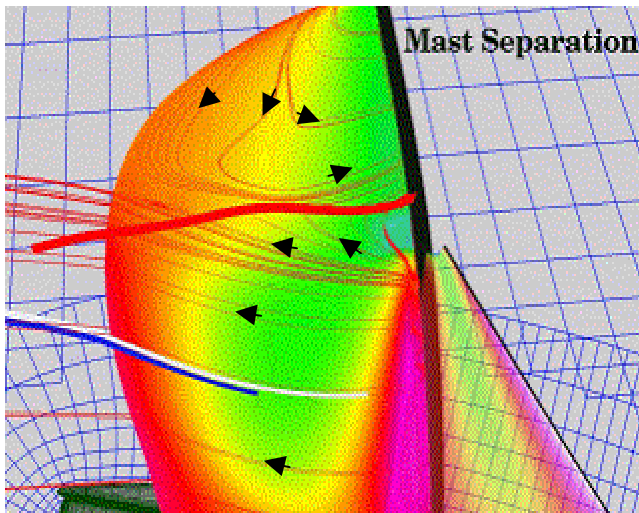


Figure 1: Separation over the Mainsail Leeward

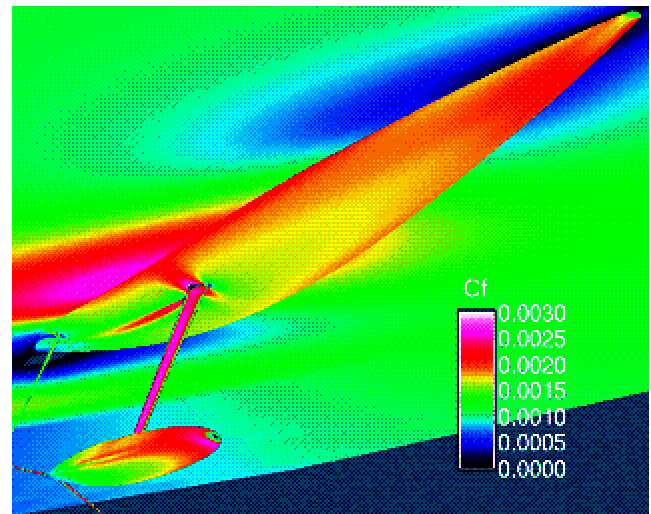


Figure 3: Distribution of Shear Stress

APPROACH

General Description of the Method

The approach followed here is unique in that it relies solely on RANS for quantifying hydro and aero force, and avoids the use of empirical corrections like “reef” and “flat”. Other than these important differences, however, the presented VPP is not substantially different from other performance prediction strategies. The basic structure consists of three sets of nested loops: an innermost set to converge the force degrees of freedom; a middle set to span the desired range of TWS and TWA; and an outer set to provide for optimizing over a range of design options. Figure 4 describes the software flow.

Four degrees of freedom are solved by the innermost loops: 1) the balance between drive and drag force; 2) between aero and hydro side force; 3) between roll moment and righting moment; and 4) between aero and hydro yaw moment. The remaining two degrees of freedom (sink and trim) are treated as they are in most VPP’s – i.e. assumed fixed throughout the solution process. Values of sink and trim are computed by providing RANS-predicted sail forces to the RANS hydro simulation, and are therefore “built in” to the grid and solution set. If later converged VPP results show substantially different aero forces, RANS hydro simulations may have to be repeated.

Four force degrees of freedom allows us to use four independent variables for defining the “state” of a yacht at each point in time. For a given TWA and TWS, the state consists of boat speed (V), a factor representing side force magnitude (SFF), roll angle (ϕ), and rudder angle (θ). The concept of SFF is best explained as follows.

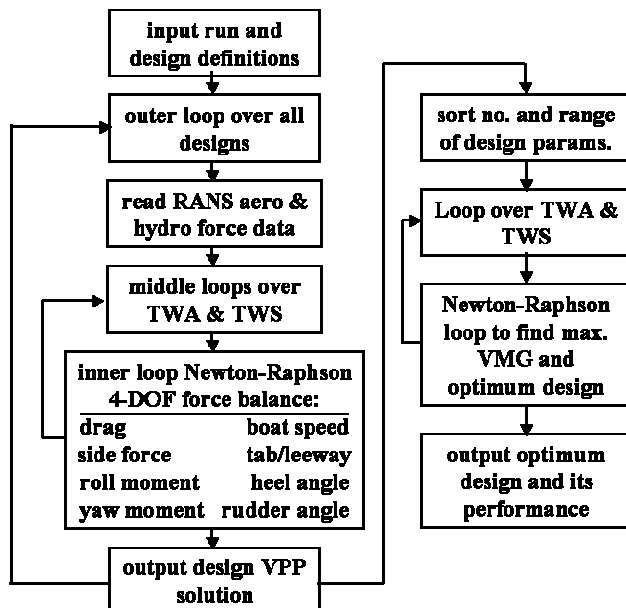


Figure 4: VPP Software Flow Chart

For each design, RANS forces have to be provided that span the range of expected sailing conditions. But if leeway, tab, and rudder are considered independent state variables, and just four points are used to span each variable, a total of $4^6 = 4096$ RANS solutions would be needed. Since RANS tends to be time-consuming, the number and range of state variables is best kept limited.

The concept of SFF is borrowed from the towing tank (where it was developed for the same reason), and helps reduce the degrees-of-freedom by realizing a yacht is required to carry a known range of side force at each wind speed. A designer chooses to counter this force through some combination of tab, leeway, and rudder. The concept of SFF takes advantage of this by assuming tab angles are prescribed at each TWS, and the force balance is obtained by varying leeway. Appendage angles can therefore be built into the RANS solution to limit that degree-of-freedom. The reason SFF is used (instead of γ) is that the prescribed angles are functions of TWS and heel.

Note that by its very nature, RANS considers a “design” to be one unique geometry with all external surface shapes prescribed. The geometry is free to move as a rigid body within the VPP, but not deform. This implies that tab and rudder (at least for prescribed values of V , ϕ , and θ) are fixed as described above. The same thing can be said of sail trim parameters like traveler angle and twist. The word “design” is therefore used a little differently in this paper than what is usually implied. This is a subtle but important distinction to understand when using this VPP.

Because the use of SFF fixes rudder angle to something “close” to that needed for steady state sailing, we can also eliminate θ as one of the state variables treated by RANS. The approach works by redefining the VPP state variable θ as the difference between that needed to balance the boat, and the “hard-wired” value built into the RANS configuration for a given SFF. If the difference is sufficiently small, its effect can be incorporated by linear airfoil theory. The combined use of SFF and rudder angle reduces the number of runs needed to define a design (again assuming four cases per state) from 4096 to $4^3 = 64$.

Relationships between the remaining state variables and forces must all be resolved by RANS simulation. For aero simulations the operating range is typically defined as those heels and AWA’s that bracket the expected steady state sailing points. The effect of AWS (Reynolds number) on non-dimensional aero force is typically quite weak, and can be ignored for any given sail shape. It is important to realize, however, that sail shape changes substantially with wind speed, so complete performance predictions require a suit of sails to cover the expected AWS range. Hydro simulations are performed over a range of boat speeds, heel angles, and SFF’s.

Each RANS simulation is a fairly large computation (4 to 6 million grid points with 6 unknowns per point), but can be automatically completed using the approach of Korpus (2004). Once all simulations for a given design are finished, a post-processor is run to compute shear stress and pressure-induced forces. Since RANS data is non-dimensional, the post-processor converts the forces to real-world values and writes a file of force components suitable as input to the VPP. VPP solutions are performed in the coordinate system shown in Figure 5. Although this paper is not intended to focus on the RANS process itself, it is important for a user to understand any limitations of that process. A summary of technical requirements will therefore be presented in a later section.

The VPP's outermost loops are included so that a range of designs can be evaluated with a single run. Once all designs are analyzed, the VPP begins an optimization process to identify the best combination of design variables. While this might be considered an "option" for some VPP's it is essential for this one. As mentioned earlier, a complete performance analysis requires that a number of sail designs be used to cover the range of AWS. Any one sail shape converges only over a very narrow range of AWA and AWS. The optimization loop must be used with sail design variations in twist and sheeting angle to identify appropriate shapes for each wind speed.

This optimization process is entirely general. The variables handled by the outer loops can be true design parameters (such as prismatic and beam), or trimming parameters (such as mainsheet tension and tab angle). They could even be some combination of both, and/or include a mixture of aero and hydro variables. VPP solutions for each design are provided, but the outer-most loops also report performance predictions and design values for the "optimal" boat.




Figure 5: VPP Coordinate System

Force Specification

Given that two degrees of freedom (sink and trim) are pre-specified, the four remaining hydrodynamic forces and moments can be uniquely defined by a four-term "state vector" $[\mathbf{V}, \text{SFF}, \phi, \theta]^T$. More specifically:

$$\begin{aligned} F_{X\text{hydro}} &= f_1(\mathbf{V}, \text{SFF}, \phi, \theta) \\ F_{Y\text{hydro}} &= f_2(\mathbf{V}, \text{SFF}, \phi, \theta) \\ M_{X\text{hydro}} &= f_3(\mathbf{V}, \text{SFF}, \phi, \theta) \\ M_{Z\text{hydro}} &= f_4(\mathbf{V}, \text{SFF}, \phi, \theta) \end{aligned} \quad (1)$$

where f_1 through f_4 are nonlinear functions defined by RANS data. Aerodynamic force and moment counterparts are similar, but more generally written in terms of state variables AWA, AWS and ϕ . The two sets of state variables are closely related, however, since AWA and AWS are functions of \mathbf{V} , γ , TWA and TWS (through the "wind triangle"), and since γ has a prescribed relation to SFF. Aero forces can therefore be restated using the same state vector as hydro forces – a transformation made automatically inside the VPP. Aero and hydro force data is input exactly as derived from the RANS simulations.

Each of the functions f_1 through f_4 (as well as their aero counterparts) must be uniquely defined. This is necessary because the VPP requires every state variable, force, and moment to be a continuous function. RANS simulations, on the other hand, are available at only discrete combinations of state variables.

Continuous functions are built from discrete RANS data using cubic splines. Hydro data is usually provided over a range of \mathbf{V} , SFF, and ϕ , so three-dimensional splines are used. Aero data also uses three-dimensional splines, but over the state variables AWA, AWS, and ϕ . AWA, and AWS are computed from TWA, TWS, \mathbf{V} , SFF, and γ before entering the spline routines.

The author is aware that some VPP's utilize more complex interpolation basis functions, and that this might someday be required here as well. To date, however, only America's Cup work has been performed, and such boats operate over fairly small ranges of speed. The forces generated by these boats are quite smooth throughout the operating range. Furthermore, while more complex basis functions might improve interpolation accuracy for specific cases, any one basis set will never prove optimal to all types of data. Every data stream exhibits slightly different shape, and may be best suited to some basis function other than the one programmed. The only universal solution is to use more points.

The number of discrete RANS simulations required to accurately define an interpolation depends on the range of state variable being investigated and the complexity of data represented. Hydro forces, for

example, change quickly with speed and heel, but are almost linear with leeway. Some knowledge of sailing physics is also important. While resolving heel angle to five-degree increments may be sufficient for some designs, it would prove unacceptable for resolving the effects of AWA on upwind sail force. Finally, it should be noted that including at least one point above and below the expected range of state is essential. Two points is recommended.

Forces and moments associated with rudder angle do not require interpolation because most of the rudder influence is already included in the RANS data. Some angle was prescribed when SFF was defined and the RANS simulations performed. The intention is that these angles lie close to the final rudder operating angles, and that only small perturbations are required to balance the boat. The rudder state variable, therefore, can be thought of as a small correction to rudder angle, and can be accurately accounted using linear theory.

If we define θ_{SFF} as the rudder angle already included in the RANS simulation, and θ as the angle required to balance the boat, than the difference $\theta - \theta_{SFF} = \delta\theta$ can be used to estimate the incremental rudder lift and drag. If $\delta\theta$ is sufficiently small (say < 3 degrees) and the rudder is not near stall, then the incremental lift and drag are estimated by:

$$\begin{aligned}\delta L &= Q_{hydro} A_{rudder} \delta\theta 2\pi / (1 + 2 / \Lambda) \\ \delta D &= 4 L_{rudder} \delta\theta / (\Lambda + 2)\end{aligned}$$

where A_{rudder} is the rudder area, Λ its aspect ratio (including hull end plate effects), and L_{rudder} the rudder lift at RANS simulation point. Incremental roll and yaw moments are easily computed knowing the horizontal and vertical distances between rudder center and moment center (values entered at input). Note that while these corrections are, strictly speaking, empirical (with apologies to the title), the intention is that $\delta\theta$, δL , and δD will be quite small if θ is properly chosen for the RANS simulations.

Force Balance (Inner Loop)

Because the functions f_1 through f_4 (and their aero counterparts) are nonlinear, their solution is not trivial. Fortunately, the Newton-Raphson method is well suited for such problems, and although it requires iterations, tends to converge very quickly. The system to be solved is:

$$\vec{F} = \vec{F}_{aero} + \vec{F}_{hydro} = \begin{bmatrix} F_{Xaero} + F_{Xhydro} \\ F_{Yaero} + F_{Yhydro} \\ M_{Xaero} + M_{Xhydro} \\ M_{Zaero} + M_{Zhydro} \end{bmatrix} = \begin{bmatrix} F_X \\ F_Y \\ M_X \\ M_Z \end{bmatrix} = \vec{0}$$

where the F and M entries come from one of the nonlinear relations in Equation (1) (or the aero counterparts).

Newton-Raphson starts with a “guess” to the state vector $[V, SFF, \phi, \theta]^T$ that “almost” solves the above system (in practice almost any guess with a state lying inside the range of provided RANS data proves sufficient). If the guess is referred to as “state n” and the actual solution as “state n+1”, then the last equation can be perturbed about the actual solution using a four-dimensional Taylor series. Truncating the Taylor series to linear terms yields:

$$\vec{F}^{n+1} = \vec{0} = \begin{bmatrix} F_X \\ F_Y \\ M_X \\ M_Z \end{bmatrix}^n + \quad (2)$$

$$\begin{bmatrix} \partial F_X / \partial V & \partial F_X / \partial SFF & \partial F_X / \partial \phi & \partial F_X / \partial \theta \\ \partial F_Y / \partial V & \partial F_Y / \partial SFF & \partial F_Y / \partial \phi & \partial F_Y / \partial \theta \\ \partial M_X / \partial V & \partial M_X / \partial SFF & \partial M_X / \partial \phi & \partial M_X / \partial \theta \\ \partial M_Z / \partial V & \partial M_Z / \partial SFF & \partial M_Z / \partial \phi & \partial M_Z / \partial \theta \end{bmatrix}^n \begin{bmatrix} \delta V \\ \delta SFF \\ \delta \phi \\ \delta \theta \end{bmatrix}$$

Since we know state n, F^n can be computed from Equation (1) (plus the aero counterparts) and all 16 partial derivatives of the matrix found by finite difference methods. The 4x4 system is then inverted to find the state perturbation vector $[\delta V, \delta SFF, \delta \phi, \delta \theta]^T$. Finally, the “guess” can be updated to state n+1 using $V^{n+1} = V^n + \delta V$, $SFF^{n+1} = SFF^n + \delta SFF$, $\phi^{n+1} = \phi^n + \delta \phi$, and $\theta^{n+1} = \theta^n + \delta \theta$.

Convergence is achieved by iterating the above process until all state variable updates are sufficiently small. The method has proven very robust in practice, and converges to speeds less than 0.01 knots and angles less than 0.01 degrees in just a few iterations. The method is so efficient that an entire array of TWS’s and TWA’s will converge in less than one second on an Itanium-2 Altix 350 server.

The procedure can be summarized as comprised of the following 12 steps:

1. Make a reasonable guess for the initial state
2. Compute AWA and AWS from the wind triangle
3. Interpolate hydro forces from the input RANS solutions to the given state vector
4. Interpolate aero forces from the input RANS solutions to the given AWA, AWS, and γ
5. Compute derivatives by finite difference (e.g. $\delta F_Y / \delta \phi = F_Y(V, SFF, \phi + \delta \phi, \theta) - F_Y(V, SFF, \phi - \delta \phi, \theta) / 2\delta \phi$)
6. Solve Equation (2) for the state perturbation vector
7. Update the guess to state n+1
8. Convergence is obtained when $\delta V < 0.01$ knots, $\delta SFF < 0.01$, $\delta \phi < 0.01$ degree, and $\delta \theta < 0.01$ degree
9. If convergence criteria not yet met, return to step 2
10. Repeat steps 1 through 9 for the next TWA
11. Repeat steps 1 through 10 for the next TWS
12. Repeat steps 1 through 11 for the next design.

Output solutions are written after analysis of each design is complete, and includes converged values of AWA, AWS,

V, SFF, ϕ , θ , VMG, seconds per mile, and the difference between θ and corresponding SFF rudder angle. Runtime condition flags are also output to show those cases where force extrapolation occurred outside the range of input RANS data, or where the force balance exhibits questionable convergence.

Design Optimization (Outer Loop)

Design optimization capabilities were specifically developed to maintain complete generality in the types of studies allowed, and can handle up to four independent variables. The variables can be classical design parameters (such as prismatic coefficient or mainsail camber) or parameters more traditionally thought of in terms of operating and trimming (such as tab angle or traveler location). The method can even handle combinations of design and operating parameters, and/or combinations of hydro and sail parameters.

Each of the (up to) four parameters can span an arbitrary number of discrete designs. The method is therefore more limited by the number of RANS solutions achievable within a reasonable time frame than anything else. The method was originally developed because a “design” consists of only one sail shape, and the range of state variables over which this will converge is fairly limited. Traditional VPP’s use “flat” and “reef” parameters to circumvent this problem, but such an approach would have negated one objective of this VPP.

To achieve convergence over a more practical range of TWA and TWS, a wider range of sail shapes must be considered. But to run a range of RANS simulations spanning traveler angle, sheeting angle, mainsail twist, and Genoa twist would require too much time to repeat for every single sail design one would like to check. Instead, a “sail trim” study is performed only once, and the optimization approach used to define the best trimming parameters for a range of TWA and TWS. Later design studies are then performed over a smaller range of trim parameters centered on this optimum.

The optimization process is similar to that used to converge forces in the inner loops. It starts once all designs of a given run set have been analyzed through the VPP, and utilizes those results to build an objective functional. Since only upwind and downwind cases have been studied to date, VMG is usually used for the objective.

Optimization occurs independently at each TWA and TWS using the variational calculus concept of “first variations” (i.e. that an extrema exists when the partial derivatives of an objective with respect to each design variable are zero simultaneously). If we call the optimization objective Ω , the first variation can be written:

$$\delta\Omega = 0 = \begin{bmatrix} \partial\Omega/\partial d_1 \\ \partial\Omega/\partial d_2 \\ \partial\Omega/\partial d_3 \\ \partial\Omega/\partial d_4 \end{bmatrix} \quad (3)$$

where d_1 , d_2 , d_3 , and d_4 are the user-specified design variables. A penalty function is added to Ω to enforce convexity and inequality constraints (so design variables stay within the range of RANS input) (Hildebrand, 1965).

We are again faced with solving four simultaneous nonlinear equations, and turn to the Newton-Raphson method. Expanding Equation (3) in a 4D Taylor series and truncating to first-order terms, we have:

$$\begin{bmatrix} \partial\Omega/\partial d_1 \\ \partial\Omega/\partial d_2 \\ \partial\Omega/\partial d_3 \\ \partial\Omega/\partial d_4 \end{bmatrix}^{n+1} = \bar{0} = \begin{bmatrix} \partial\Omega/\partial d_1 \\ \partial\Omega/\partial d_2 \\ \partial\Omega/\partial d_3 \\ \partial\Omega/\partial d_4 \end{bmatrix}^n + \begin{bmatrix} \partial^2\Omega/\partial d_1^2 & \partial^2\Omega/\partial d_1\partial d_2 & \partial^2\Omega/\partial d_1\partial d_3 & \partial^2\Omega/\partial d_1\partial d_4 \\ \partial^2\Omega/\partial d_2\partial d_1 & \partial^2\Omega/\partial d_2^2 & \partial^2\Omega/\partial d_2\partial d_3 & \partial^2\Omega/\partial d_2\partial d_4 \\ \partial^2\Omega/\partial d_3\partial d_1 & \partial^2\Omega/\partial d_3\partial d_2 & \partial^2\Omega/\partial d_3^2 & \partial^2\Omega/\partial d_3\partial d_4 \\ \partial^2\Omega/\partial d_4\partial d_1 & \partial^2\Omega/\partial d_4\partial d_2 & \partial^2\Omega/\partial d_4\partial d_3 & \partial^2\Omega/\partial d_4^2 \end{bmatrix} \begin{bmatrix} \delta d_1 \\ \delta d_2 \\ \delta d_3 \\ \delta d_4 \end{bmatrix} \quad (4)$$

The procedure works exactly as before except that now we use Ω (e.g. VMG from all the previous VPP solutions) instead of force from RANS solutions. Since the objective functional is known only at a discrete number of design points, it must first be rendered continuous through interpolation. All first and second derivatives are found by second-order accurate finite difference formulae. Finally, Equation (4) is solved for the design variable perturbation vector, and the process repeated until converged.

Like most optimization processes, robustness depends on the problem. Simple and well-behaved systems (like the sail trim study mentioned previously) converge quite nicely. But in general, such success should not be anticipated for every combination of design variables.

Runtime Input Requirements

The vast majority of effort required to use this VPP is expended on generating RANS solutions to provide force input. Not all RANS codes are efficient and robust enough to meet these needs, so a separate section in this paper is dedicated to describing the essential requirements. This current section, therefore, focuses on information needed to run the VPP once RANS forces are available. This turns out to be a remarkably simple task.

Figure 6 shows how simple a typical input deck actually is. The only inputs are the range of TWA’s and TWS’s, and geometric parameters necessary to define the rudder force perturbation model. Inputs also include an option for correcting RANS force data that may have used

separate moment centers in the aero and hydro simulations. The only complication is to ensure that the range of TWA's and TWS's selected fall within the expected sailing conditions represented within the input force data. If, for example, 8-knot TWS sails are used to generate RANS aero data, but the VPP is run at 16 knots, excessive heel angles or lack of convergence could result.

Figure 7 shows a sample design optimization control input file. The first few lines provide the number and names of design variables being optimized. Since the VPP works with any combination of variables, providing names is essential to document results. Line four specifies the number of designs in the optimization set, and is followed by one line describing each design. Description lines include the design name, the values of each design variable, and file names for input aero and hydro force data. Designs can be input in any order -- the only restriction being that the matrix of design variables is not sparse. That is, if two variables are optimized -- one spanned by four designs and the other by two -- the inputs must include eight designs covering every combination.

Figure 8 shows a simplified sample of the input RANS hydro force data. Each geometry component is given separately, and forces in all six degrees of freedom are included. Separate lines within each geometry component give results for one set of state variables. In general there would be many more lines required to span the likely range of operating states, but this file has been shortened to limit required page space (the number of significant digits has also been reduced). Files for the RANS aero forces look similar, and both files are standard outputs from AFT's automated RANS post processors.

The number of RANS simulations needed to populate the VPP (and the time required to obtain them) is an essential element in planning its use. If tab and rudder angles can be approximated from previous experience, then the SFF state variable approach can be used, and the hydro problem is not severe. Hydro force data is usually well represented by 7 speeds, 7 heel angles, and 3 leeway angles; so a total of 147 runs can cover both upwind and downwind sailing. The 16-processor Itanium-2 Altix-350 server at Applied Fluid Technologies completes these runs in 8 days. While this is certainly a substantial amount of computer time, it is considerably faster than building and testing a tow tank model.

Later analysis of similar designs can be completed more efficiently by solving RANS only for those combinations of state variable that lie close to converged VPP solutions of the lead design. Force data associated with state variables far from a converged operating point are needed only during the inner loop force convergence process. Such points will not affect the converged result, so data can be "recycled" from the first design. These "follow-on" designs are usually completed in 2 to 3 days.

```

16 number of TWA's required
25.0 minimum TWA
40.0 maximum TWA
1 number of TWS's required
12.0 minimum TWS
12.0 maximum TWS
0.0 long. distance from aero to hydro moment center
0.0 vert. distance from aero to hydro moment center
8.0 long. dist. from rudder to hydro moment center
1.0 vert. dist. from rudder to hydro moment center
1.0 rudder planform area in square meters
20.0 rudder aspect ratio (including hull end plate)

```

Figure 6: Sample VPP Run Control Input Deck

```

2 number of design variables
design_var_1 name of first variable
design_var_2 name of second variable
8 number of designs to optimize
var1_1_var2_1 1. 1. v1_1_v2_1.aero v1_1_v2_1.hydro
var1_1_var2_2 1. 2. v1_1_v2_2.aero v1_1_v2_2.hydro
var1_1_var2_3 1. 3. v1_1_v2_3.aero v1_1_v2_3.hydro
var1_1_var2_4 1. 4. v1_1_v2_4.aero v1_1_v2_4.hydro
var1_2_var2_1 2. 1. v1_2_v2_1.aero v1_2_v2_1.hydro
var1_2_var2_2 2. 2. v1_2_v2_2.aero v1_2_v2_2.hydro
var1_2_var2_3 2. 3. v1_2_v2_3.aero v1_2_v2_3.hydro
var1_2_var2_4 2. 4. v1_2_v2_4.aero v1_2_v2_4.hydro

```

Figure 7: Sample VPP Optimization Input Deck

```

hull
caseID area drag side vertical roll pitch yaw
state_1 58. 0.1 -0.2 -17. -6.5 9.9 -0.1
state_2 58. 0.1 -0.3 -17. -6.7 9.4 0.3
keel & tab
caseID area drag side vertical roll pitch yaw
state_1 4. 0.0 -0.9 -0.5 2.3 0.3 -0.4
state_2 4. 0.0 -1.1 -0.6 2.8 0.4 -0.5
bulb
caseID area drag side vertical roll pitch yaw
state_1 10. 0.0 0.0 -1.7 3.0 1.0 0.0
state_2 10. 0.0 -0.0 -1.7 3.1 1.1 0.1
starboard (windward) winglet
caseID area drag side vertical roll pitch yaw
state_1 0. -0.0 0.0 -0.0 -0.0 0.0 0.0
state_2 0. -0.0 0.0 -0.0 -0.0 0.0 0.0
port (leeward) winglet
caseID area drag side vertical roll pitch yaw
state_1 0. -0.0 -0.0 0.0 0.1 -0.1 -0.1
state_2 0. -0.0 -0.0 0.0 0.1 -0.1 -0.2
rudder
caseID area drag side vertical roll pitch yaw
state_1 2. 0.0 -0.3 -0.2 0.5 1.7 -2.6
state_2 2. 0.0 -0.3 -0.2 0.6 1.9 -3.1
complete yacht
caseID area drag side vertical roll pitch yaw
state_1 77. 0.3 -1.0 -19. -0.2 13. -3.3
state_2 77. 0.3 -2.0 -19. 0.0 12. -3.5

```

Figure 8: Sample Hydro Input Force File (RANS coordinate system, Significant Digits Truncated)

The problem becomes more time-consuming if tab and rudder angles are considered independent state variables. Tab and rudder angle optimization is therefore usually performed just once, and the best values chosen at each TWS for all follow-on analyses. The optimization might be repeated if the design evolves extensively, but many weeks of RANS runs are required. Such an approach is not practical for every design.

Similar arguments hold for aero force analysis. If sheeting angle and sail twist are approximated a priori, only the AWA, AWS, and ϕ state variables need be considered. For windward/leeward type racing, experience has shown such a state variable space can be spanned with 3 AWA's, 5 AWS's, and 3 heel angles – a total of 45 RANS runs for upwind sails and 45 for downwind. AFT's computers require about 6 days to complete the 90 runs, but follow-on designs (if similar) can be completed faster. A complete set of TWA polars would require substantially more time.

If optimal Genoa and mainsail sheeting angles and twists are not known a priori, the process takes substantially longer. Each sail shape or trim parameter must be treated as an independent design variable, and the optimization process completed. Huge numbers of RANS runs are required, and the approach is not practical to repeat for every design. As with tab and rudder angles, the optimization is performed just once, and the best trim values used over the wind range for all similar designs.

RANS CODE REQUIREMENTS

Four essential elements are mandatory for RANS to prove practical at powering VPP's. First, it must achieve a level of throughput able to analyze a design within some "reasonable" timeframe. The definition of "reasonable" of course depends on the project – but days to weeks are more likely than months or years. Second, it must prove sufficiently refined to resolve very small design changes (otherwise we could rely on simpler computational or experimental methods). RANS' niche is its ability to differentiate very subtle differences, so this goal is paramount. Third, it must be robust enough to complete hundreds of simulations without failure. And fourth, it must be integrated with a high level of automation.

Meeting these demands 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 a high degree of robustness so that no time is wasted debugging unstable runs. A system is required to automate geometry handling, grid generation, and RANS setup because only by "removing man from the loop" can high throughput be achieved and human error avoided. Finally, resolution of small differences requires an unprecedented degree of grid repeatability, discretization accuracy, and turbulence model sophistication.

A RANS package meeting these requirements has been developed and validated by AFT over the last three America's Cup cycles. 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 16-processor Itanium-2 SGI Altix-350 parallel supercomputer. Recent increases in processing speed (combined with substantial decreases in cost) have made the ownership of such powerful computers practical for even small companies like AFT. Production levels up to 500 three-dimensional (4 to 6 million point grid) RANS analyses per month are regularly obtained on AFT-owned hardware.

At its innermost level, the system utilizes an in-house RANS code known as AFTINS (Applied Fluid Technologies Incompressible Navier-Stokes). AFTINS is a purpose-built, time-accurate, incompressible flow code utilizing the discretization techniques presented in Chen (1988, 1989) and Chen and Korpus (1993). It has been extensively validated on a wide range of steady and unsteady problems (Amromin, et. al., 2004, Korpus, 2005, 2004, 2000, Korpus, et. al. 1998, 1996, and Kumarasamy et. al. 1997). It utilizes the overset method to ease grid generation and improve solution accuracy, and $k\epsilon$ turbulence models for reliable separation modeling.

These last two features are critical for practical application, and worthy of additional description. While many highly capable RANS codes meet the demands of discretization accuracy (Chen, 1989, Cowles, 2003), not all incorporate sufficiently robust turbulence modeling. Because sailboat flows exhibit large degrees of separation (e.g. mainsail luff, Genoa leech, spinnakers), they require sophisticated turbulence models to accurately resolve boundary layer attachment. Two-equation models are typically regarded as the minimum level sufficient to capture complex separation physics over a wide range of applications. Zero-equation models (such as Baldwin-Lomax) are robust and easy to use, but not particularly accurate in separated flows. One-equation models show remarkable accuracy in some flows, but are less impressive in others (Spalart and Allmaras, 1992). For similar reasons, all turbulence models need to solve all the way through a boundary layer and viscous sub-layer to the no-slip wall. Wall functions are used by some practitioners to reduce grid size, but cannot be trusted to capture separation phenomena (Fan et. al., 1983, Chen and Patel, 1988).

While RANS codes are quite complex pieces of software, 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, and 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 uncommon for even expert users to require many days for meshing a domain. While this may prove acceptable for one-off flow studies, it could never satisfy the demands of VPP applications to time-critical design problems.

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 more subtle 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 it 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 execution. Once one flow solution has reached convergence, it can be reasonably assumed they all will. The resulting robustness is an essential time saver.

The best meshing technique for maintaining true topological similarity is overset gridding. 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 restrictions are that the union of all blocks spans the computational domain. Since each block is generated using conventional structured gridding, maintaining topological similarity is straightforward.

Overset grids have the accuracy advantages of structured grids (e.g. boundary layer resolution), but the simplicity advantages of unstructured grids. The domain about each geometry element is meshed without regard to surrounding geometry elements or their 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 spacing be specified. A more detailed description can be found in Korpus (2004).

The approach is summarized in Figures 1 and 2. The first shows one grid plane cut through the mast, mainsail, and Genoa, and the second the corresponding solution. The view in each figure looks down from above with the mast at bottom center and mainsail to the left. The Genoa is visible at upper right. Independent mast, mainsail, and Genoa grid blocks are visible, as are the arbitrary way in which they overlap. The RANS solution is continuous and smooth across jagged block boundaries.

Once high-quality, topologically controlled grids are combined with robust, accurate RANS solvers, only a few additional utilities are required to build a practical RANS system. First, a capability must be created to convert geometry from designer-provided CAD to grid-friendly input. Second, input files must be generated to control grid resolution, overset pre-processing, 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-processor must be provided to retrieve forces, convergence residuals, run time errors, and any other detailed flow data required for analysis.




Figure 9: Typical Overset Grid Near Mast & Sails




Figure 10: Typical Solution Near Mast & Sails

Even with all these automated utilities, however, keeping hundreds of RANS analyses running is too much for most humans. One final utility is needed 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 automation tasks. The first is a “run management script” responsible for handling the individual software modules required to complete a single RANS run. The second is an “executive script” responsible for managing the complete set of analyses required for one design, monitoring available computer resources and launching the 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 set of state variables). 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 each set of state variables, 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 this data can be consolidated into a single file for easy downloading and viewing.

SAMPLE APPLICATIONS

The application of RANS-based VPP’s to design will be demonstrated using three practical design examples. The first will show how changes to an America’s Cup class hull, keel, and bulb impact VMG. The second will demonstrate the process of finding optimum sail trim -- an essential application since “flat” and “reef” are no longer

available. The third example demonstrates outer loop capability by combining a hydro design change with sail trim optimization. Reference to specific design variables or quantitative results is obscured to protect sensitive data.

Example 1 -- Hydrodynamic Design

Example 1 demonstrates the straightforward use of a RANS-based VPP in the traditional sense – to rank two candidate hull designs. Only upwind performance is considered. Each hull is analyzed by RANS at 72 sailing “states” covering a range of conditions, and then post-processed to build a table of VPP input hydro forces (similar to Figure 8). Sail forces are also provided by RANS, but cover a more limited range to mimic the fidelity of data available within other VPP tools. In particular, while different sail shapes are used across the TWS range, no attempt is made to find optimum trim at any one state. A range of heel and AWA’s are used to build the data set, but sail shape and trim are the same for each hull. The hull data sets required 6 days each of RANS analysis on AFT’s in-house computers. Aero force generation took two days.

Differences between the two designs look slight to the casual observer, but would be considered “substantial” by an America’s Cup designer. Beam ranges over 5%, for example, so that speed changes are easily discernable. Polar plot results are shown in Figure 11, and predictions from a potential flow VPP are included for comparison.

Two observations are immediately obvious. The first is that the difference between RANS and potential-flow predictions is larger than the design variation predicted by RANS. Although the general trends are similar, boat speed and VMG magnitudes are not. TWA at optimum VMG is also different.

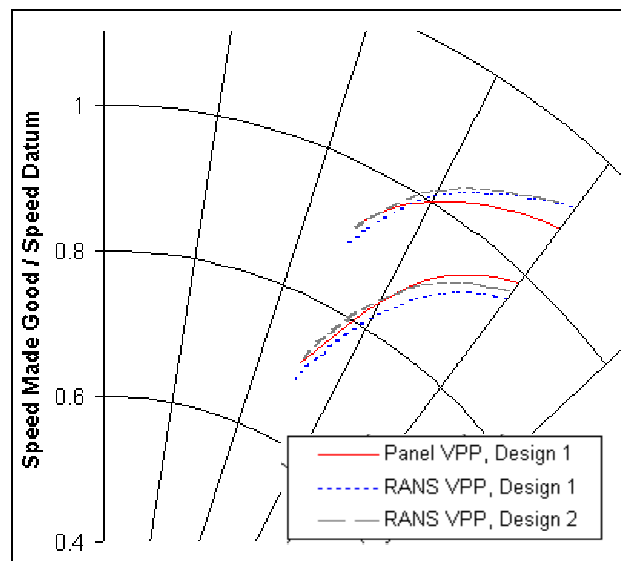


Figure 11: Speed Polar Comparing Designs and Analysis Technique

The second observation is that there is a “cross-over” between the two techniques. Compared to potential flow, RANS predicts higher boat speed, VMG and optimum TWA at high TWS, but lower values at low TWS. This difference is perhaps more important than the first because CFD is usually justified by its ability to predict design trends. While no CFD code has yet proven 100% accurate, RANS has been repeatedly demonstrated to adequately resolve relative differences. Yet here we see an example of low-resolution CFD (potential flow) being unable to capture the effect of decreasing sail performance with TWS. This trend is well known, and typically attributed to the panel code’s inability to resolve the effect of separation on a sail plan as Reynolds number decreases.

Figures 12, 13, 14, and 15 show additional outputs from the VPP, and depict the usual way information is portrayed to a designer. The figures give “target” values defined as those numbers a sailor would see when the boat is performing at optimum VMG. Boat speed is given in Figure 12, and clearly depicts crossovers when comparing either the analysis technique or the designs.

Figures 13 and 14 depict heel angle and optimum AWA respectively, and show the most startling differences between RANS and potential flow. Regardless of design, RANS predicts the boat should be sailed more “loaded” than potential flow (i.e., at higher heel angles and AWA’s). This is true across the tested TWS range, and is a trend confirmed by comparison against on-the-water sailing data. This indicates that potential flow either over predicts hull drag due to heel, over predicts the detrimental effect of heel on sail aerodynamics, or both.

Figure 15 shows (not surprisingly) that higher rudder angles are needed to sail the boat the RANS way. It also indicates that substantially different rudder angles are required for balance despite the fact that RANS predicts similar heel angles and uses identical sails.

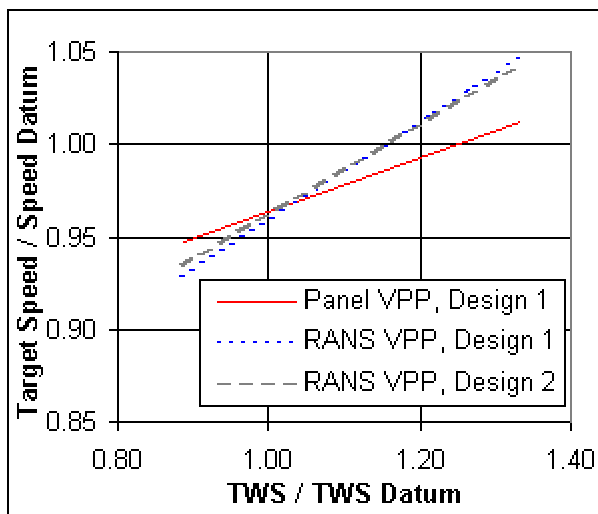


Figure 12: Comparison of Target Boat Speed

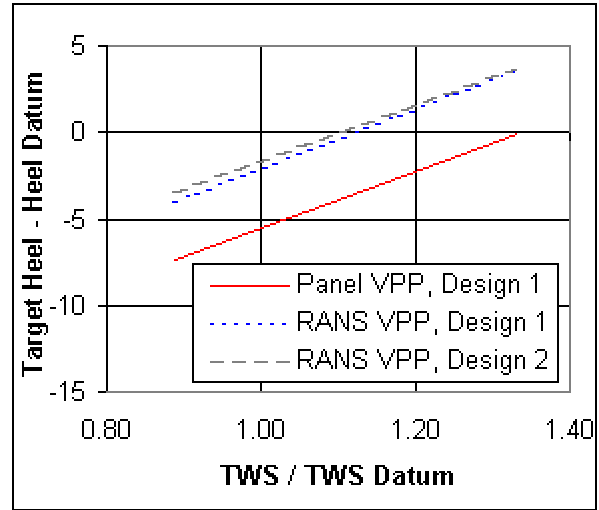


Figure 13: Comparison of Target Heel Angle (degrees)




Figure 14: Comparison of Target AWA (degrees)

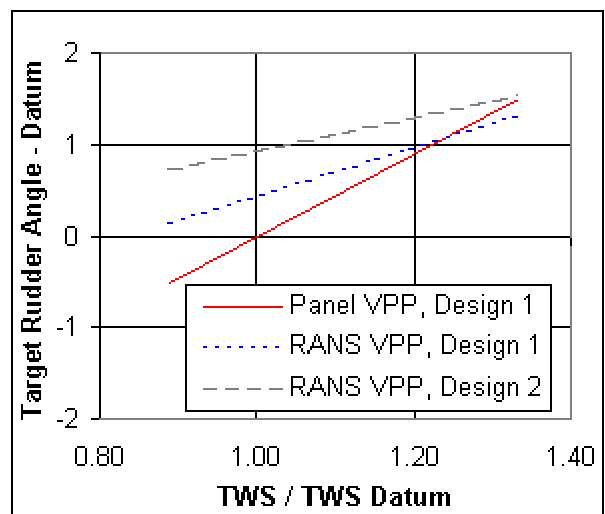


Figure 15: Target Rudder Angle (Degrees)

Example 2 -- Aerodynamic Design Optimization

Traditional VPP's use empirical correction factors to ensure that a limited set of sail forces can be used to converge over a wide range of sailing conditions. A factor known as "flat" is used to reduce overall force, and another known as "reef" is used to reduce roll moment. Such factors are entirely empirical, and have not been included here. This puts huge demands on the input force data since a sail designed for 8 knot TWS will likely overpower a boat at 16 knots TWS. Similarly, a sail trim used for "pinching" will not prove particularly good for "footing." It is essential, therefore, that a range of sail shapes and trims be included in a RANS-based performance analysis.

Such a capability is introduced into the current VPP by performing RANS simulations over a range of sail twist, Genoa sheeting angle, and mainsail boom angle (traveler position). The VPP outer loop optimization process is then used to identify the best trim combination for each design and sailing state. It is important to emphasize that this approach is NOT intended to define optimum sail trim. Computers will likely never prove better at trimming than a sailor constantly adjusting for subtle changes in wind direction, speed, and gradient. Instead, the approach was developed solely to ensure that design decisions are made using only those conditions centered near a boat's optimum sailing state.

The example presented here accounts for changes in trim using three variables: Genoa sheeting angle; mainsail sheeting angle; and sail twist. Genoa and mainsail twist are varied in unison in order to reduce the number of analyses required. Each sail trim is tested over a range of heel angles and AWA's, and repeated at three TWS's. The required 324 runs take 20 days on AFT's computers. Hydro data from Design 2 of the previous example is used.

Figure 16 shows the resulting polar, but for clarity includes only a small subset of the sails analyzed. Lines of one color and pattern correspond to one particular TWS, and VMG performance at each is seen to have dramatically different character. At low TWS, for example, VMG changes relatively slowly with sail trim and AWA. At higher wind speeds performance is more sensitive to sail trim and AWA. Off-design performance degradation at the highest wind speed is dramatic.

Performance at optimum sail trim is plotted using heavy lines that correspond to the individual wind speeds by color and pattern. Optimization effectiveness is demonstrated by the fact that heavy-line results are clear envelopes bracketing the individual sail shapes. At each TWA and TWS, optimum VMG is greater than or equal to that of all the sail shapes that make up the individual sail designs. The sail shape that corresponds to this optimum changes dramatically with TWS and TWS. Optimum TWA decreases with increasing TWS as expected.

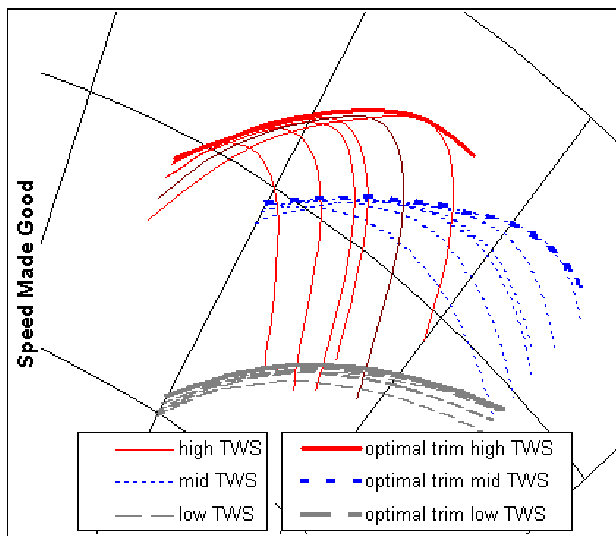


Figure 16: VMG Polar over a Range of Sail Trim

Target AWA and heel angle are shown in Figures 17 and 18 respectively. As before, only a small subset of the total sail shape inventory is included for clarity, and results corresponding to optimum sail trim are given by a heavy line. The overall trends make sense when compared to on-the-water observations. Heel angle, for example, is seen to increase rapidly until a fairly high value is reached near the middle TWS range. At this point the boat is "fully powered" and heel starts to level off. Additional heel could add more sail power, but not enough to overcome additional hull drag.

There is no such obvious pattern to describe the relationship between individual sail performance predictions and optimum sail performance predictions. AWA starts near the upper end of the range of input sails at low TWS, but drops into the middle of the range as TWS increases. Optimum heel angle starts near the upper end of the input range, but ends close to the lower end.

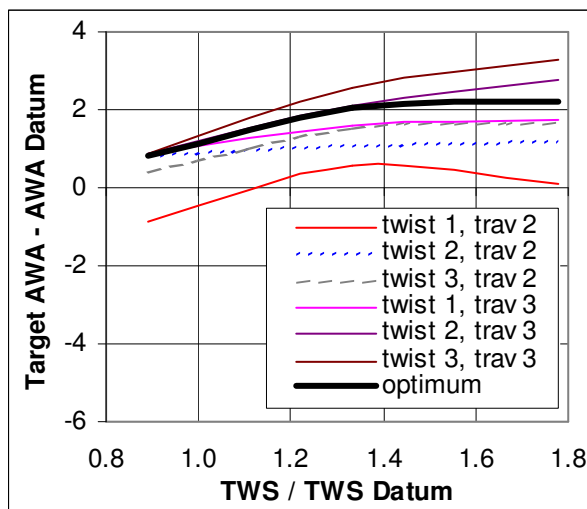


Figure 17: Target AWA versus Sail Trim (Degrees)

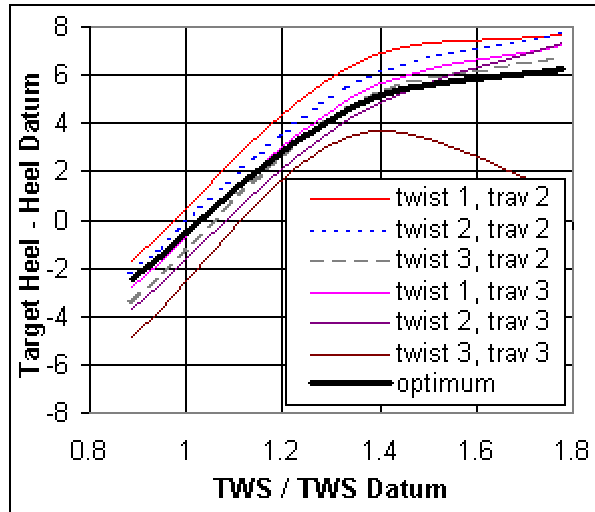


Figure 18: Target Heel versus Sail Trim (Degrees)

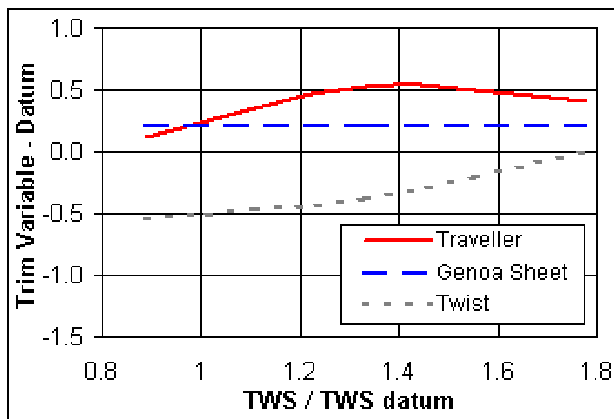


Figure 19: RANS-Predicted Optimum Sail Trim (Degrees)

Optimum sail trim as predicted by RANS is shown in Figure 19. This particular design appears insensitive to Genoa sheeting angle, but otherwise normal compared to other boats. As TWS increases the first required trim change is to drop the traveler. Mainsail twist is next increased to control heel. At the highest TWS's twist has increased to the point where lowering the traveler has to stop in order to keep the mainsail leech working.

As mentioned earlier, these results do not necessarily dictate "on the water" trim due to nuances in real world sailing conditions. They could however, be used to check on-the-water test data; to speed tuning of a new boat; or to provide a starting point for polar development.

Example 3 -- Simultaneous Optimal Design

The last example combines the previous two, and highlights the importance of simultaneously optimizing sails and hull. It uses the 324 sail trim variations of example two to analyze the hulls of example one. The process is repeated twice: once with a fixed sail shape and a second time using sail shapes optimized for each hull.

Figure 20 shows VMG versus TWA as predicted using the first approach. Three wind speeds are included, and the results show that like the earlier examples, performance is fairly insensitive at low TWS. Performance at the higher wind speeds is more sensitive, and also indicates a crossover between the two designs. Design 2 is still best at optimum TWA (maximum VMG), but Design 1 is a better boat for "footing."

Figure 21 shows results for the analysis repeated using optimization to select the best sails for each design. As expected, VMG's are consistently higher than those of Figure 20. At the higher TWA's though, optimization has improved Design 2 more than Design 1, and the crossovers are eliminated. This provides a graphic demonstration of the importance of simultaneous optimization. It also highlights the shortcomings of trying to predict hull design trends using simplified sail force data.

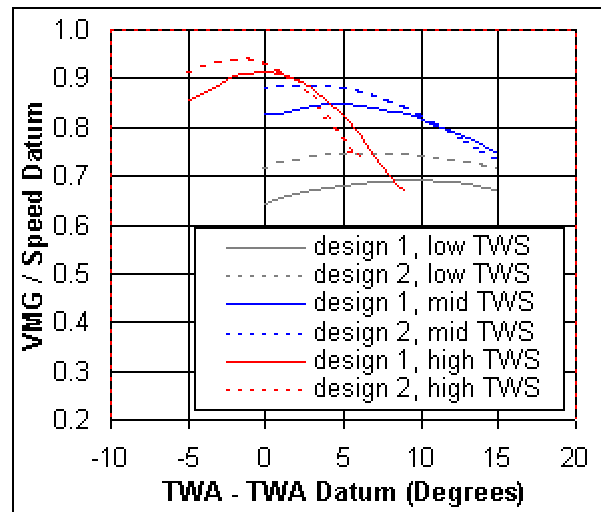


Figure 20: Performance Comparison with Identical Sails

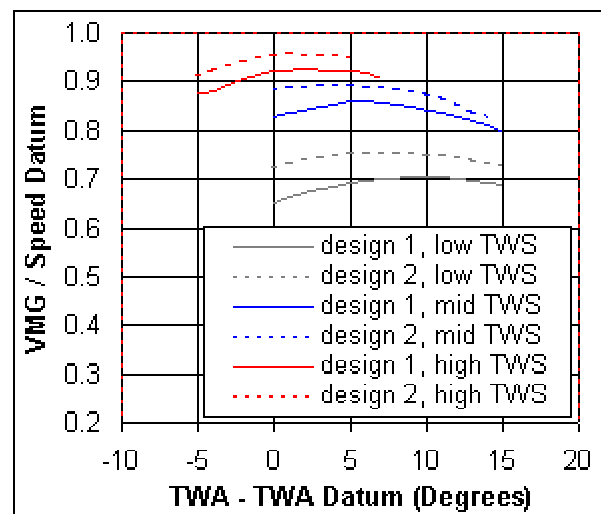


Figure 21: Performance Comparison with Sails Optimized for Each Design

This example also highlights the importance of using VMG to rank a design rather than relying on raw force data. Figure 22 shows targets obtained using optimal sail trim for each hull design, and includes both boat speed and VMG. Although VMG is better for Design 2, the boat speeds themselves are very close; indicating that their drag is likely comparable. The advantage of Design 2 is its pointing ability rather than its low drag.

Figures 23 and 24 show the corresponding target AWA's and heel angles respectively. Taken together with Figure 22 they indicate that while Design 2 is consistently better, it is for different reasons across the wind speed range. At the higher wind speeds Design 2 is able to sail both with more power (greater heel) and with better pointing (lower AWA). At the lower wind speeds Design 2 has to sail at higher angles, but manages to power up faster.

Figure 25 compares the optimal trims needed by each design. The format is similar to Figure 19, except that trend lines are now included for each of the two hulls. Like the results of Figure 19 we see that performance is fairly insensitive to Genoa sheeting angle. The other trim parameters, however, differ quite a bit. Both traveler and twist are consistently lower for Design 2, and indicate the tighter trimming required to sail closer to the wind (as demonstrated by Figures 23 and 24). Design 2 continues to keep its mainsheet tighter (twist lower) even as wind speed increases to its maximum. Because Design 2 twist is maintained to lower levels, the traveler does not have to be let down as far as Design 1. The fact that Design 1's traveler actually comes back towards centerline at high wind speeds indicates that twist must be excessively high at that point.

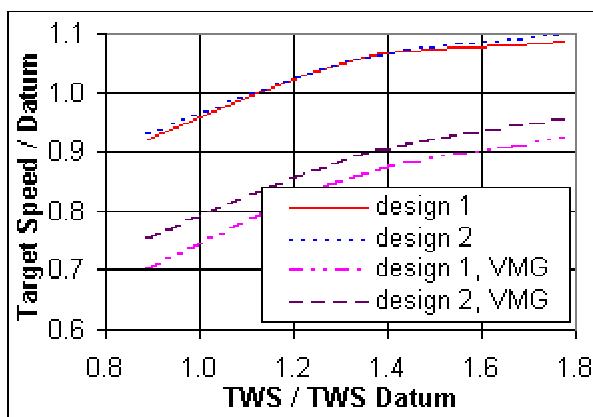


Figure 22: Target Boat Speed with Optimal Sails




Figure 23: Target AWA with Optimal Sails (Degrees)




Figure 24: Target Heel with Optimal Sails (Degrees)

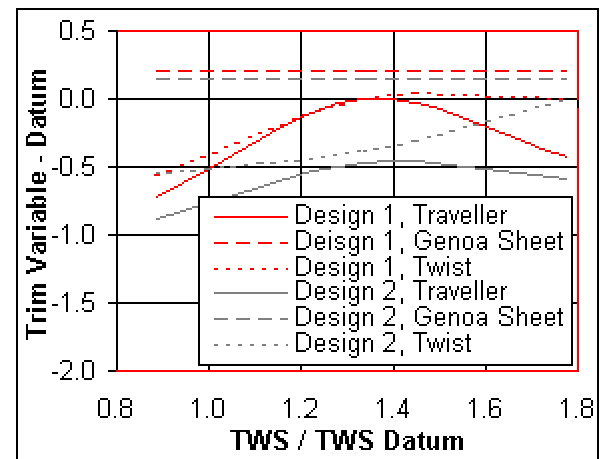


Figure 25: Comparison of Trim with Optimal Sails (Degrees)

CONCLUSIONS

The concepts presented in this paper demonstrate that it is now possible to drive sailboat performance prediction using aero and hydro forces based on first-principles, high-resolution, RANS-based CFD. The advantages of force predictions that resolve viscosity, turbulence, separation, and shed vorticity are shown to be substantial. Reynolds-Averaged Navier-Stokes methods are demonstrated to ably provide these forces in a robust, accurate, and efficient manner.

It has also been demonstrated that, for some applications, RANS-driven VPP's might even be considered practical. This is not to say that such VPP's are suitable for every application, or that they will replace existing VPP's. But when used wisely for demanding applications where no other force-generation procedure is adequate, they can lend invaluable design guidance. Examples of such applications include (but are not limited to): sail trim or shape optimization when separation is present; hull design when small changes to viscous drag are expected; and mast tube section design.

Sample calculations presented herein demonstrate that high-resolution force data from RANS also alleviates many of the numerical difficulties common with some VPP's. Convergence is rapid and robust, and even design optimization (for the few examples tested) proves simple.

The coupling of RANS to VPP's provides a significant step towards bringing viscous flow CFD into the forefront of yacht design. Performance-based design programs, especially those where small improvements are paramount, will benefit immediately from this new high-fidelity velocity prediction capability. As the method continues to mature, it is inevitable that an even larger portion of the yacht design community will eventually embrace RANS-based VPP's.

ACKNOWLEDGEMENTS

The author would like to acknowledge the long-term encouragement and support of Nils Salvesen. Dr. Salvesen was one of the pioneers in applying CFD to sailing yacht design (Stars & Stripes, 1987), and believed as long ago as 1988 that RANS would someday become a practical design aid. His confidence in advocating what (at the time) was as immature research technology was tireless. Due in large part to his perseverance, RANS development for maritime applications has thrived, and finally matured into a useful design tool. The accomplishments presented in this paper would not have been possible without the strength of his convictions.

REFERENCES

- Amromin, E., Miaine, I., Cook, L., Day, W., and Korpus, R., "High Speed Trimaran Drag: Numerical Analysis and Model Tests," *J. of Ship Res.*, v. 47, n. 2, 2004.
- Chen, H-C., and Korpus, R., "A Multi-Block Finite-Analytic Reynolds-Averaged Navier- Stokes Method for 3D Incompressible Flows," ASME Summer Fluids Conf., Wash. D.C., 1993.
- Chen, H.-C., Patel, V., and Ju, B., "Solutions of Reynolds-Averaged Navier-Stokes Equations for Three-Dimensional Incompressible Flows," *J. Comp. Physics*, 1989.
- 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.
- Claughton, A., "Developments in the IMS VPP Formulations," 14th CSYS, Annapolis, MD 1999.
- Cowles, G., Parolini, N., and Sawley, M.I., "Numerical Simulation Using RANS-Based Tools for America's Cup Design," 16th CSYS, Annapolis MD, 2003.
- Euerle, S.E., and Greeley, D.S., "Towards a Rational Upwind Sail Force Model for VPP's," 11th CSYS, Annapolis, MD, 1993.
- 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.
- Hildebrand, F.B., Methods of Applied Mathematics, Prentice-Hall, Englewood Cliffs, NJ, 1965.
- Jackson, P., "An Improved Upwind Sail Model for VPP's," 15th CSYS, Annapolis, MD, 2001.
- Jones, P., and Korpus, R. "International America's Cup Class Yacht Design Using Viscous Flow CFD," 15th CSYS, Annapolis, MD, 2001.
- Keuning, J.A. and Sonnenberg, U.B., "Approximation of the Calm Water Resistance on a Sailing Yacht Based on the Delft Systematic Yacht Hull Series," 14th CSYS, 1999.
- Kerwin, J.E., "A Velocity Prediction Program for Ocean Racing yachts," MIT Ocean Eng. Report n. 78-11, Cambridge, MA 1978.
- Korpus, R. and Liapis, S., "Active and Passive Control of SPAR Vortex-Induced Motions," Proc. OMAE, Halkidiki, Greece, 2005.
- Korpus, R., "Reynolds-Averaged Navier-Stokes in an Integrated Design Environment," Madrid Deseno de Yates, Madrid, ESP, 2004.
- Korpus, R., Jones, P., Oakley, O., and Imas, L., "Prediction of Viscous Forces on Oscillating Cylinders by Reynolds-Averaged Navier-Stokes," ISOPE, Seattle, WA, 2000.

Korpus, R., Hubbard, B., Jones, P., Stromgren, C., Bennett, J., "Hydrodynamic Design of Inte-grated Propulsor/Stern Concepts by Reynolds-Averaged Navier-Stokes," 7th Practical Design of Ship & Mobile Units (PRADS), Le Hague, NL, 1998.

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.

Kumarasamy, S., Korpus, R., and Balow, J., "Computation of Noise Due to the Flow Over a Circular Cylinder," Proc. 2nd Comp. Aero-acoustics Work-shop on Benchmark Problems, NASA Conf. Pub. 3352, pp. 297-303, 1997.

Oliver, J.C. and Claughton, A.R., Development of a Multi-Functional Velocity Prediction Program (VPP) for Sailing Yachts," Int. Conf. CADAP95, RINA, 1995.

Oliver, J.C., "Performance Prediction for Multi-Hull Yachts," 9th CSYS, Annapolis, MD, 1989.

Oliver, J.C., Letcher, J.K., and Salvesen, N., "Performance Prediction for Stars and Stripes," Tr. SNAME, NY, 1987.

Rosen, B.S., Laiosa, J.P., and Davis, W.H., "CFD Design Studies for America's Cup 2000," AIAA-2000-4339, 2000.

Spalart, P.R. and Allmaras, S.R., "A One-Equation Turbulence Model for Aerodynamics Flows," AIAA 92-0439, AIAA 29th Aerospace Sciences Meeting, 1992.

Von Oossanen, P., "Predicting the Speed of Sailing Yachts," Tr. SNAME, v. 101, 1993.