OPERATIONS RESEARCH
Vol. 52, No. 1, January–February 2004, pp. 1–16
ISSN 0030-364X | EISSN 1526-5463 | 04 | 5201 | 0001

DOI 10.1287/opre.1030.0078
© 2004 INFORMS

# A Simulation Model for Predicting Yacht Match Race Outcomes

**A. B. Philpott**
Department of Engineering, University of Auckland, Private Bag 92019, Auckland 1020, New Zealand, a.philpott@auckland.ac.nz

**S. G. Henderson**
School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York 14853, shane@orie.cornell.edu

**D. Teirney**
Orion Systems NZ Limited, Auckland, New Zealand, david.teirney@orion.co.nz

We outline the development of a model for predicting the outcome of a yacht match race between two competing designs. The model is a fixed-time-increment simulation that accounts for the dynamic performance of each yacht. The wind speed and direction are modelled using hidden Markov chain models. Each yacht is assumed to follow a fixed sailing strategy determined by a set of simple decision rules. The simulation models both yachts simultaneously and accounts for interactions between them—for example, when they cross. The model is illustrated by applying it to International America's Cup Class designs.

*Subject classifications*: Simulation applications. Recreation and sports: yacht performance analysis.
*Area of review*: Simulation.
*History*: Received May 2002; revision received February 2003; accepted March 2003.

## 1. Introduction

Computer modelling has become an important part of high-performance yacht design. In this paper, we describe the development and use of a computer program for modelling yacht match races. As distinct from fleet races between many boats, a match race is held between a pair of nearly identical boats. The classic example is the America's Cup, which is held every three years between a challenging yacht and a defending yacht. Since 1992, the yachts competing for the America's Cup must satisfy the International America's Cup Class (IACC) design rules. Although these are intended to yield designs with similar performance, there is considerable scope for design variations within the constraints laid down by the IACC rules. In the America's Cup, computer models are used extensively to evaluate different design options.

The key computer model currently used by yacht designers is called a *velocity prediction program* (VPP). A VPP combines models for the hydrodynamic forces on the hull and the aerodynamic performance of the sails to predict a yacht's performance in a given weather condition. Velocity prediction programs were devised over 20 years ago by yacht designers and handicappers. The seminal paper in this area is by Kerwin (1978), who formulated the equilibrium behaviour of a yacht as a set of simultaneous nonlinear equations equating the forces and moments on the vessel. Because one of the variables in these equations is the velocity of the yacht, a solution to these equations will give a prediction of the velocity.

A number of authors have proposed improvements (see, e.g., Letcher 1974, Schlageter and Teeters 1993, van Oossanen 1993, Philpott et al. 1993) to Kerwin's basic model, and much effort has been devoted to improving the force models that provide the appropriate equations in the VPP. In particular, a VPP requires models for both the hydrodynamic forces on the hull and the aerodynamic forces on the sails. The hydrodynamic model uses information from computational fluid dynamics (CFD) programs and towing-tank data to fit equations that define the forces and moments on a hull moving through the water at a given velocity and trim. Similarly, the aerodynamic model uses CFD, scale-model measurements in wind tunnels, and full-scale measurements to derive models for the lift and drag on the sails and rig, as a function of apparent wind speed, wind angle, and sail trim. All of these models are dependent on some geometrical description of the sails, as well as the rig, the hull, and its appendages—namely the rudder, keel, and (possibly) a bulb and winglets.

A design evaluation tool that is closely related to the VPP is the *race modelling program* (RMP). The output from a VPP, along with historical weather data, is used in an RMP to assess different yacht designs by racing candidate designs against each other over a range of weather scenarios. A win/loss probability is then estimated for each pair of yachts.

Although high-performance yacht designers now routinely simulate the performace of candidate designs for ocean races, the use of RMPs in short course races has

been confined primarily to the America's Cup. The use of RMPs was first reported by the 1987 *Stars & Stripes* syndicate in a famous *Scientific American* article (Letcher et al. 1987). The *Stars & Stripes* design team realized early in the campaign that

> yacht racing has an essentially random component in that the relative performance of two yachts depends on the wind speed and the sea conditions, which vary randomly from day to day. VPP results by themselves are therefore inconclusive and possibly misleading for determining the order of merit of two candidate yachts over a series of races (Letcher et al. 1987, p. 39).

For the *Stars & Stripes* campaign there were two RMPs developed. The first was a simple probabilistic model using the predicted time difference between the two yachts as a function of wind speed, and a distribution for the wind speed to determine the winning probability. The second RMP improved on the first, allowing for wind speed distributions for each leg, and also taking into account interaction effects when the yachts are very close. The output from each RMP was used to analyse the win/loss probabilities of two yachts over a specified course.

Since 1987, RMPs have featured in several different America's Cup challenges. For example, one of the major programs for the Partnership for America's Cup Technology (PACT), which was founded in 1990, was to gather site-specific environmental data in San Diego—the location for the next America's Cup in 1992. These data were used in the creation of "a statistical weather model . . . for use in conjunction with the RMP which was developed for evaluating the probability of success for various designs in conditions likely to be experienced off San Diego during the trials and the America's Cup" (Gretzky and Marshall 1993, p. 203). PACT researchers also deployed a wave-measuring buoy to gather sea state spectra, which could then be correlated with local meteorological conditions so that the RMP could be run with rough water effects (Gretzky and Marshall 1993).

The approach to race modelling pioneered by the *Stars & Stripes* design team estimates times around a course in different wind and wave conditions, sampled for each leg of an America's Cup course. It is well known that the leading boat in a match race has an enormous advantage owing to its ability to "cover" its opponent by sailing a course that keeps between the trailing boat and the next mark. When sailing upwind, this advantage is enhanced by the ability to spill turbulent "dirty" air onto the trailing boat. Race-modelling programs of the type described above can account for this advantage to some extent by conditioning the probability of a yacht being in front at the end of each leg on whether it is in front at the beginning, but these effects are difficult to quantify even with a large amount of experience.

This paper describes an alternative approach that uses a fixed-time-increment simulation model of each leg that accounts for wind fluctuations and interactions between the two boats. Previous RMPs have used equilibrium performance predictions for the speed of the yacht, and probability distributions for the wind speed estimated from historical data. In a fixed-time-increment simulation, the speed of the yacht changes dynamically as the weather conditions vary. This means that if the equilibrium performance characteristics of two yachts are similar, but one accelerates more quickly, then this is reflected in the result.

Previous RMPs do not consider the tactical advantages that a faster yacht has over a slower yacht. The speed advantage is simply translated into a faster time around the course. However, "when the tactical advantages of the faster boat are considered as well, the effects of the speed differences are even greater" (Milgram 1993, p. 13-7), and therefore current RMPs may underestimate the probability that a faster yacht wins. Indeed, "during match racing between two yachts, the leader has more opportunities for subsequent gains than does the follower" (Milgram 1993, p. 13-3). In a simulation of a match race the tactical advantages that a leader has can be modelled, thereby improving the estimate of the win/loss probability when comparing yacht designs.

Finally, most RMPs do not take into consideration the interaction between the yachts. The design program for *Stars & Stripes* in 1987 realised that "being faster than the opponent is not always enough to ensure victory. Mistakes, luck and other unforeseen events affect the outcome of a race. To be certain of winning a race requires a margin of victory large enough to overcome these unpredictable occurrences" (Schlageter and Teeters 1993, p. 300). This was reflected to some extent in the RMP developed for the *Stars & Stripes* design program, where a yacht had to beat the opposition by some specified time margin to be sure of a victory. In a simulation of a match race the interaction between the yachts can be modelled, and thus some of the unforeseen events are removed. Of course, there will always be some unforeseen events that are too complicated or unlikely to be included in the model. A central question is to determine what factors can be ignored in the model without biasing the results.

One of the difficulties in creating a simulation is ensuring that there is no bias within the simulation, so that there is no bias in the results. Another difficulty is to ensure that the true result is not swamped by noise, which is usually reflected in the variance of the results. A simulation is still a probabilistic tool that requires many runs (race simulations) to get an accurate estimate for the true probability of winning. It is therefore imperative that any models created for the simulation are as fast as possible.

Our main contribution in this paper is to develop a methodology for predicting match-race outcomes that is an improvement over existing approaches. We wish to emphasise that our intention in this paper is not to specify the optimal design parameters of an America's Cup yacht. Indeed, because many of the inputs for the simulation model will depend on the race location, the person helming the boat,

and the attitudes and preferences of the design team, an optimal yacht design will be contingent upon these.

The outline of the paper is as follows. In the next section, a model is derived to compute the dynamic performance of a yacht. This allows the speed and location of the yacht to be updated at each time step, dependent on the observed weather conditions. Section 3 describes the methods used to generate the weather conditions observed on a single yacht at each time step. Section 4 describes the approach used to model tactical decisions that would be made by the helmsman on a yacht in response to the position of the yacht on the course, the position of the other yacht, and the changing weather conditions. In particular, a method is described that uses penalties relating to different tactical situations that occur during match racing to determine the best tack for a yacht in its current situation. Methods for determining the optimum heading to sail using one of two common sailing strategies are also discussed. We conclude by giving some results of running the simulation model on some candidate designs.

An online companion to this paper in the Online Collection at the *Operations Research* home page http://or.pubs.informs.org/pubs/collect/html provides further details on some of the aspects of the simulation that are merely motivated and sketched here. Included in this companion are three avi files that show animations of the simulation results. These are particularly useful to illustrate some of the effects we seek to model in this paper.

Our approach is a heavy user of notation, so for the sake of clarity we reuse notation in several places. We hope that the meaning of the notation will be clear from the context, so that no confusion will arise.

## 2. The Yacht Model

### 2.1. The Yacht Dynamics

The dynamic model of the yacht uses three state variables. These are the position $s$, speed $V_s$, and heel angle $\phi$ of the yacht. These evolve according to a system of difference equations that are integrated forward in time. The speed of the yacht is updated every time step during the simulation, where each time step is separated by a discrete constant time interval $\Delta t$. The time interval $\Delta t$ is chosen small enough so that the acceleration of the yacht over $\Delta t$ can be assumed to be constant.

The acceleration is a function of the net thrust $F$ and the total mass being accelerated, which is made up of the combined mass $m$ of the yacht, gear, and crew, and an added mass $M$ due to a body of water surrounding the yacht being accelerated as well. The added mass $M$ is a constant that is estimated for each yacht using measured performance data, and differs for acceleration and deceleration. For time step $k$ this gives

$$V_s^{k+1} = V_s^k + \frac{F^k}{m+M}\Delta t. \tag{1}$$

To update the location of a yacht each time step, it is moved in a straight line in the direction it is travelling (assumed to be its heading). The distance $s^k$ that is travelled by a yacht during the time step $k$ is

$$s^k = \frac{V_s^k + V_s^{k+1}}{2}\Delta t.$$

The thrust $F$ is a function of $V_s$ and heel angle $\phi$. Therefore, to evaluate the right-hand side of (1) we must compute $\phi$. The dynamics governing changes in $\phi$ are complicated owing to the effects of added mass terms and moments of inertia (see Davies 1990). For this reason we have chosen to approximate the heel dynamics by moving the heel angle $\phi^k$ for the yacht at the current time step $k$ closer to an equilibrium heel angle $\bar{\phi}$ for the observed weather conditions so that

$$\phi^{k+1} = (1-p)\phi^k + p\bar{\phi}, \tag{2}$$

where the parameter $p \in (0, 1)$ determines how fast the heel angle approaches its equilibrium.

The three equations above define the difference equations to be satisfied by the yacht. It remains to specify models to determine $F^k$ and $\bar{\phi}$. These will both depend on the observed weather conditions, and the heading of the yacht. One approach to modelling these is to develop a full set of difference equations for the yacht dynamics, as described in Davies (1990). This requires a large number of coefficients to be estimated for each candidate design, a costly exercise. A compromise is to use the data obtained from VPPs to produce an approximate dynamic model. The advantage in this approach is that designers will use VPPs routinely to investigate designs, and their software can easily be configured to deliver the information required for the approximate dynamic model. Henceforth, we shall refer to an underlying VPP model that can be used to obtain information on the forces acting on any candidate design.

The values of $F^k$ and $\bar{\phi}$ are determined in the simulation for a given true wind speed $V_t$, and a true wind angle $\beta_t$, being, respectively, the wind speed and direction observed by a stationary observer facing in the direction of the yacht's tack. Given $V_t$ and $\beta_t$, a VPP will determine the equilibrium speed of the yacht and its equilibrium heel angle $\bar{\phi}$ in those wind conditions. However, in the simulation the current yacht velocity $V_s^k$ will not equal the equilibrium yacht velocity, and this will produce a force imbalance that can be used to derive $F^k$.

Suppose, then, that the yacht is in state $(s^k, V_s^k, \phi^k)$ and we observe $V_t$ and $\beta_t$. This gives an apparent wind speed $V_a$ and an apparent wind angle $\beta_a$ that can be calculated using the apparent wind triangle shown in Figure 1. As shown, the wind on the sails gives a lift $L$ (defined to be the aerodynamic force perpendicular to the wind direction) and a drag $D$ (defined to be the aerodynamic force in the direction of the wind). This results in a thrust $T$ and a side force $S_A$ defined by

$$T = L\sin\beta_a - D\cos\beta_a,$$

**Figure 1.** Resolving lift and drag to thrust and side force.

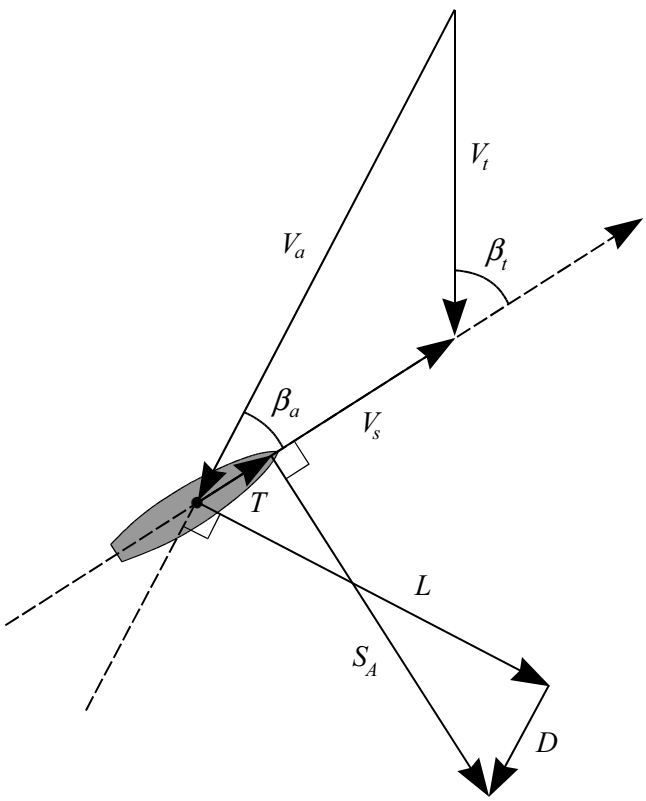


and

$$S_A = [L\cos\beta_a + D\sin\beta_a]\cos\phi^k.$$

We now assume that the aerodynamic side force $S_A$ generated by the wind pressure on the sails is matched instantaneously by $S_H$, the hydrodynamic side force currently generated by the hull and appendages. On the other hand, the thrust $T$ will in general differ from the hydrodynamic resistance $R$ experienced at $V_s^k$ and $\phi^k$; this difference $(T-R)$ will give $F^k$.

We shall not explore the details of the hull and sail models here, but give a sketch only. The reader is referred to van Oossanen (1993) for a comprehensive description of the modelling involved. The lift and drag forces on a foil in fluid of density $\rho$ and velocity $V$ are calculated using the dynamic pressure $q = \frac{1}{2}\rho V^2$, the plan area $A$ of the foil, and coefficients $C_L$ for the lift and $C_D$ for the drag, both of which are functions of the angle of attack $\alpha$. In particular,

$$L = qAC_L(\alpha),$$

and

$$D = qAC_D(\alpha).$$

For the sails the angle of attack $\alpha$ is effectively the same as $\beta_a$, but must be corrected for heel and leeway. Moreover, the air velocity must be adjusted to account for variations of wind speed with altitude, and to account for the component of airflow towards the tip of the mast.

The lift and drag are modified by two normalized sail-trim variables that are standard in most VPPs. Flat($\bar{f}$) models a reduction in the camber of the sail by reducing its lift by the fraction $\bar{f}$, and reef ($\bar{r}$) models shortening the sail by reducing the sail area $A_S$ by the fraction $\bar{r}^2$. The aerodynamic lift comes from a tabulated maximum lift coefficient $C_{L\max}$, combined with flat and reef to give

$$L = qA_S C_{L\max}(\alpha)\bar{f}\bar{r}^2.$$

The aerodynamic drag is a function of three separate sail-force coefficients: the parasitic drag coefficient $C_{Dp}$, the induced drag coefficient $C_{Di}$, and the windage drag coefficient $C_{Dw}$. The parasitic and induced drag coefficients are both functions of the apparent wind angle, whereas the windage drag coefficient is a constant. The parasitic drag is modified by reef and the induced drag coefficient is modified by both flat and reef. The aerodynamic drag is calculated by summing these components, giving

$$D = qA_S\bar{r}^2 C_{Dp}(\alpha) + qA_S\bar{r}^2 C_{Di}(\alpha)\bar{f}^2 + qC_{Dw}.$$

Here, the values of $\bar{f}$ and $\bar{r}$ are those that maximize $V_s$ for $V_t$ and $\beta_t$, as determined by the VPP. Observe that because reef lowers the centre of effort of the sailplan, the optimal choice for $\bar{r}$ in any given true wind speed $V_t$ and true wind angle $\beta_t$ will affect the optimum equilibrium heel angle $\bar{\phi}$ in this wind condition. This is accounted for in the heel dynamics (2).

The hydrodynamic resistance of a yacht is usually separated into three components:

1. Upright resistance $R_u$, which is the resistance of the hull and appendages with zero heel angle;
2. Heeled resistance $R_h$, which is the additional resistance caused by heeling the yacht;
3. Induced resistance $R_i$, which is the additional resistance caused by the leeway angle and is related to the side force generated by the hull and appendages.

In our model, we set $R_u$ and $R_h$ to be the values calculated by the VPP for state $(V_s^k, \phi^k)$. The induced resistance component is calculated using

$$R_i = \frac{(S_A)^2}{q_w 2\pi[\bar{B}_e(V_s^k, \phi^k)]^2}.$$

Here, we use the fact that the hydrodynamic sideforce $S_H$ exactly matches $S_A$. In this formula, $q_w$ is the dynamic pressure of water, and the equilibrium effective span $\bar{B}_e$ of the hull and appendages is generated by the VPP in terms of the speed and heel angle of the yacht. The total hydrodynamic resistance $R$ is then obtained by summing the separate resistance components, giving

$$R = \bar{R}_u(V_s, \phi) + \bar{R}_h(V_s, \phi) + R_i.$$

## 2.2. Tacking

Tacking occurs when a yacht changes direction from its current course to a new course having the wind on the opposite side. In general, yachts perform tacks when sailing both upwind and downwind.

For simplicity, we shall restrict our discussion here to tacking upwind. The downwind case is very similar. When a yacht executes a tack upwind, the yacht moves by some distance $s_{\text{tack}}$ upwind, loses some amount of speed $\Delta V_s$, and takes some amount of time $\Delta t_{\text{tack}}$ to do so.

One attempt to model these functions is given by Bilger (1995), who derives formulae for the loss in yacht speed during the turning part of a tack, which is approximated to be a circular arc. An alternative approach fits

$$s_{\text{tack}} = v(\Delta t_{\text{tack}}, V_t, \beta_t, V_s)$$

to data recorded from tacking manouevres. The function we use approximates the distance $s_{\text{tack}}$ that a yacht moves upwind during a tack using its *velocity made good* ($V_{\text{mg}}$) which is defined to be $V_s \cos\beta_t$. The value of $V_{\text{mg}}$ changes during a tack (see Marchaj 1996) as shown in Figure 2.

> The hump at the beginning of the manoeuvre, i.e., temporary increase in $V_{\text{mg}}$ is caused by the boat shooting into the wind. The valley afterward corresponds to the loss in $V_{\text{mg}}$ as the boat heads off and loses speed. After a time, the apparent wind angle $\beta_a$, the boat speed $V_s$, and subsequently $V_{\text{mg}}$ return to their pre-tack values (Marchaj 1996, p. 250).

We assume that $V_{\text{mg}}$ at the start of the tack is approximately the average $V_{\text{mg}}$ for the duration of the tack because the integral of the curve shown in Figure 2 over the interval $\Delta t_{\text{tack}}$ is approximately equal to the constant $V_{\text{mg}}\Delta t_{\text{tack}}$. (In fact, as shown in Figure 2, there will be some small loss in distance made good, but we treat this as a second-order effect.) Thus, the empirical function $v$ we use to calculate the distance that a yacht moves upwind during a tack is

$$s_{\text{tack}} = V_s \cos\beta_t \Delta t_{\text{tack}}.$$

This approach is supported by Bilger's finding "that there is little loss in distance made good during the turning part of the tack due, in main, to the shooting up effect of turning up into the wind" (Bilger 1995, p. 269).

The speed $\Delta V_s$ that a yacht loses during a tack is chosen to be

$$\Delta V_s = U,$$

where $U$ is a specified constant that differs for upwind and downwind sailing.

A drawback of using empirical models to define the behaviour of a yacht during a tack is that the state variables for a yacht are not easily updated at each time step, because the length of time $\Delta t_{\text{tack}}$ taken to tack is unlikely to be the same as (or even a multiple of) the time interval $\Delta t$ between time steps in the simulation. To deal with this, the yacht needs to stop "time stepping" with the rest of the simulation when it begins to tack, and then needs to restart once the tack is finished and all the state variables have been updated. This stopping and starting of time stepping has to be done carefully so that the yacht remains synchronised with the rest of the simulation. A yacht should restart time stepping at the next time step after a tack is finished. Therefore, a yacht needs to move forward for some length of time after a tack is completed so that it restarts time stepping at the correct location. Figure 3 illustrates this situation. Here, if the decision to tack is made at time $t$ (point **A**), then the tack will be completed at time $t + \Delta t_{\text{tack}}$ (point **B**). Time stepping will restart at the next time step, which is at time $t + 2\Delta t$ (point **C**), and therefore the time to move the yacht forward after completing the tack is $2\Delta t - \Delta t_{\text{tack}}$.

**Figure 2.** Graph of $V_{\text{mg}}$ vs. time during a tack.

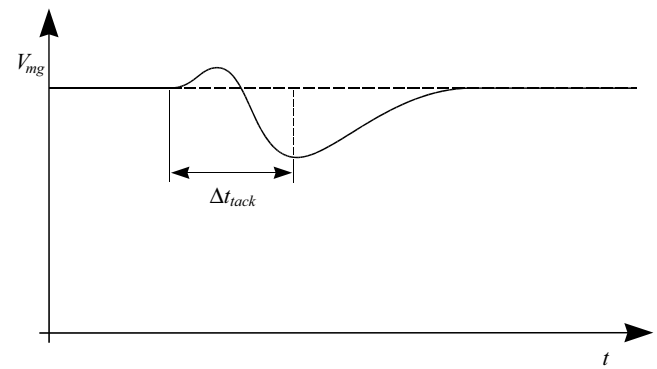


**Figure 3.** Synchronising tacking with the rest of the simulation.

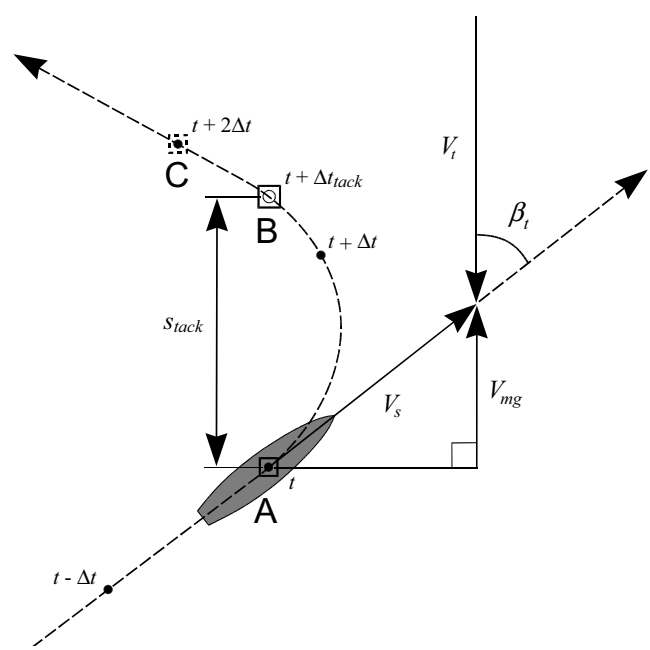

### 2.3. Mark Rounding

In contrast to tacking, mark rounding has a number of complicating factors that require careful attention when modelling. We sketch the essential ideas of our model for mark rounding here, and refer the reader to Appendix A in the online companion for full details. In yacht races, marks have to be rounded in a certain direction, and so the actual path that a yacht may take in rounding a mark can be rather complex. In particular, it depends on the tack of the yacht as it approaches the mark, the desired tack after rounding the mark, and a number of other factors including the wind direction at the time of mark rounding. These complexities have led us to adopt a similar approach to that used in tacking, whereby the yacht stops "time stepping" with the rest of the simulation when it begins to round a mark, and is resynchronized with the rest of the simulation after it has rounded the mark.

## 3. The Weather Model

The dynamics described in the previous section are driven by a model of weather, which is described in this section. The model described here is complex. It captures the dynamic (in time) behaviour of wind direction and speed at the two yacht locations.

A far simpler model of wind would require one to merely sample a wind speed and direction at the start of a race, and then leave it as a constant throughout the race. Unfortunately, this approach would lead to biased design decisions. To see why, note that, for example, it would rarely be optimal to tack with such a weather model. Yacht designs that trade straight-line speed for loss of speed during a tack would be favoured. However, in match racing, a yacht that loses a great deal of speed during a tack fares poorly in tacking duels and in general racing.

Another important issue in constructing a model of wind speed and direction is the issue of *dependence*. In particular, when two yachts are very close to one another, then they observe basically identical wind direction and speed. However, when they are separated, they can observe very different wind conditions. Indeed, leading yachts often employ a "loose cover" to try to ensure that the trailing yacht cannot exploit a local wind effect and, as a result, pull ahead. It is important to capture this correlation for design decisions because, as we show in §5, the probability that one yacht design beats another depends quite strongly on the degree of this correlation.

We begin by describing a model for wind speed and direction as seen by a stationary observer, and then show how this can be transformed into a model for a moving yacht. We then derive a model for variations in wind speed and direction on each of two yachts; these must be constructed so as to reflect the correlation between these variations when the yachts are close together. Finally, we discuss how the position and orientation of a yacht affects the wind speed and direction as observed on its competitor.

**Figure 4.** Sketch showing eddy representation of turbulence passing measuring stations.

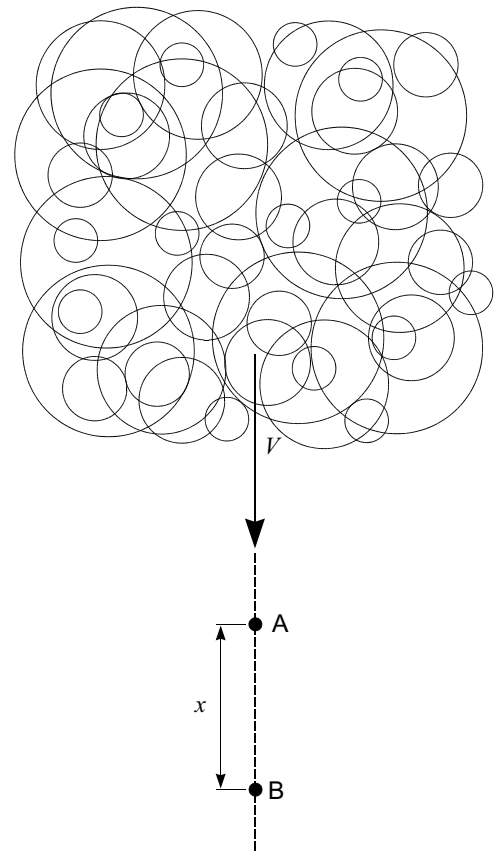


*Note.* From T. V. Lawson, *Wind Effects on Buildings: Statistics and Meteorology*, p. 144. © 1980 Applied Science Publishers Ltd., London, U.K.

Throughout this section we base our weather models on the Taylor hypothesis.

> [The Taylor hypothesis] is a concept of turbulence as a random mixture of eddies of all shapes and sizes as if etched on a sheet of perspex, the perspex moving with the mean wind speed $[V]$ (see Figure 4 (reproduced from Lawson 1980, p. 144)). An anemometer at point **A** would record exactly the same turbulence as an anemometer at **B**, only at an instant $x/V$ seconds earlier (Lawson 1980, p. 143).

Hence, we can think of the weather system as an unchanging giant wind field (which contains eddies that make the weather conditions deviate from their global means) moving at a global mean wind speed in a global mean wind direction. The Taylor hypothesis also assumes that the wind field is isotropic, which means that it has the same statistical properties in all directions.

### 3.1. Weather Models for Stationary and Moving Observers

Suppose that a giant wind field remains stationary, rather than moving, and that there are two stationary observers within it. The wind speed and wind direction seen by the stationary observers remain constant. Now suppose that the

two observers begin to move in the same direction, one moving twice as fast as the other, and that they both travel a set distance. The faster observer will cover the set distance in half the time that the slower observer will cover the same distance.

If one were to create a time series from both observers, where observations were recorded at the same discrete time intervals, the faster observer would only have half as many observations as the slower observer. Even though each observer has seen the same physical observations, the apparent rate at which observations were recorded is different and, therefore, one is able to construct the faster observer's time series from the slower observer's time series by taking every second observation. Thus, the apparent rate at which observations are recorded is important when moving through a wind field.

We call the changing weather conditions as seen by a stationary observer (at the normal rate) a *stationary observer weather series*. A stationary observer weather series is used to create the changing weather conditions that a moving yacht sees. We call the changing weather conditions as seen by a moving yacht a *weather realisation*.

The weather realisation observed by a moving yacht is created from a stationary observer weather series in accordance with the Taylor hypothesis. To do this, suppose that we have a weather time series $W = \{W(T) : T \geqslant 0\}$ as seen by a stationary observer, and we wish to convert this to the time series as seen by a moving observer. As discussed earlier, a moving observer will see the same weather time series $W$ as a stationary observer, although at a different rate. This amounts to a simple time change. To make this time change, we define two different clocks. First, there is a *race clock* that determines the "normal" time for a moving observer that would be seen, for instance, on a stopwatch timing the race. Second, there is a *weather clock* that determines the time at which the moving observer is seeing weather conditions from a stationary observer weather series.

If an observer is moving at an apparent speed $V_a$ relative to the wind field, and the global mean wind speed (of the wind field) is $\overline{V}$, then the observer sees the weather time series $W$ at rate $V_a/\overline{V}$. For instance, in a time interval of $\Delta t$ seconds in the race clock (for a moving observer), the weather clock (for a stationary observer) will advance by

$$\frac{V_a}{\overline{V}}\Delta t \tag{3}$$

seconds. The apparent speed of the yacht $V_a$ relative to the wind field is calculated using the vector subtraction of the yacht's speed $V_s$ from the global mean wind speed $\overline{V}_t$ as shown in Figure 5.

Suppose $t$ is the current race clock time, $T$ is the time in the weather clock that corresponds to $t$, and we seek the weather at the position of the yacht at $t + \Delta t$. As illustrated in Figure 5, we compute

$$\Delta\widehat{T} = \frac{V_a}{\overline{V}_t}\Delta t, \tag{4}$$

**Figure 5.** Generating new weather conditions observed on a yacht.

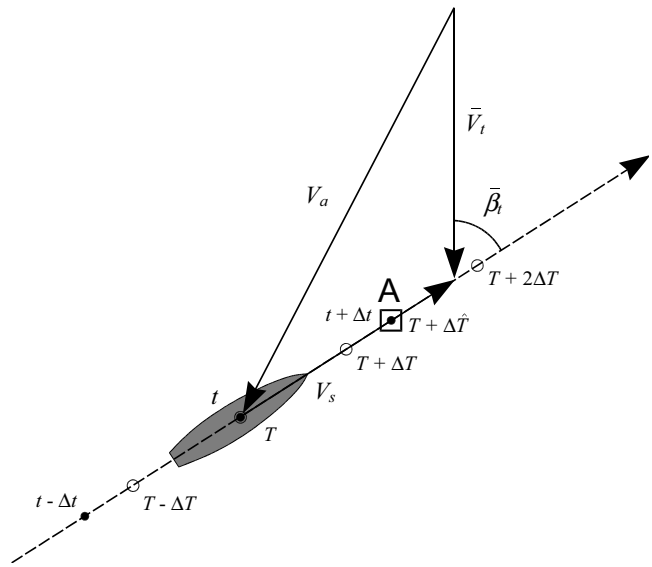


and use linear interpolation to calculate the weather conditions at the time $T + \Delta\widehat{T}$ in the weather clock. In Figure 5, the required time advance $\Delta\widehat{T}$ (at point **A**) is somewhere between $\Delta T$ and $2\Delta T$. Therefore, linear interpolation is used between the weather conditions observed at times $T + \Delta T$ and $T + 2\Delta T$ in the (stationary observer) weather series.

### 3.2. Weather Model for a Stationary Observer

We now turn our attention to constructing a weather model for a stationary observer. The major assumption we use when generating this weather series is that there is no correlation between the wind speed and the wind direction. On a large scale (in both time and space), this assumption clearly does not hold. Geographical and climatological conditions often determine the relative strength of the wind given a predominant wind direction. However, on a small meteorological scale, the assumption has some validity. This assumption provides a much simpler model in which wind speed and direction can be represented independently. Notwithstanding, the models obtained under this assumption must be quite sophisticated to represent real weather data.

A sample of a wind realisation as seen by a stationary observer is shown in Figure 6. We applied standard time-series fitting software to the wind speed for a stationary observer, and found that an excellent fit was provided by an autoregressive moving average (ARMA) process

$$V_t^k = \alpha + \varphi_1 V_t^{k-1} + \varphi_2 V_t^{k-2} + \vartheta_1 \varepsilon^{k-1} + \varepsilon^k, \tag{5}$$

where the sequence $\{\varepsilon^k : k \geqslant 0\}$ consists of independent, identically distributed normal random variables with mean 0 and variance $\sigma^2$ ($\varepsilon^k \sim N(0, \sigma^2)$).

**Figure 6.** Stationary observer wind speed and direction series.

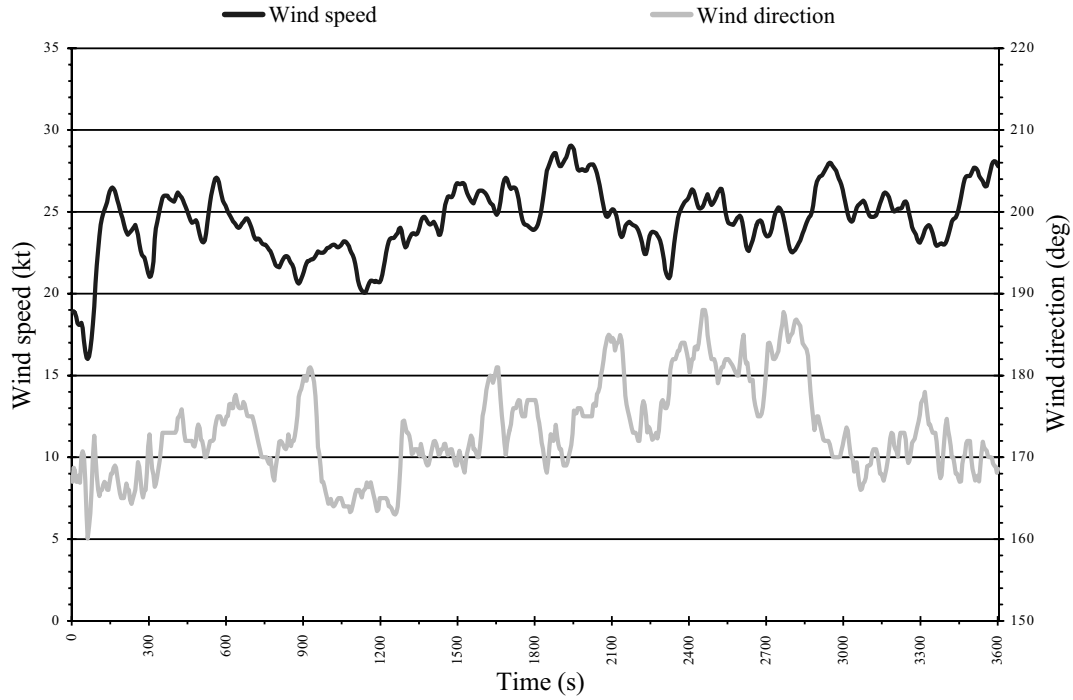


Modelling the wind direction requires more care, because at irregular time intervals there are large movements in wind direction known as wind shifts. We model wind shifts using a Markov chain with specified wind directions as states. To create the Markov chain for the Markov process, the deviations from the global mean wind direction are used. These deviations are calculated using the overall mean wind direction for the stationary observer wind direction series as an approximation to the true global mean wind direction. The wind direction deviations are then allocated to *wind direction bins* of a certain width. Each wind direction bin is considered as a state in a Markov chain and, therefore, the one-step transition matrix is of the form $P(i, j)$, where $i$ and $j$ vary over the possible wind direction states. The mean wind direction for each Markov state is approximated to be the centre of the associated wind direction bin.

Once the wind direction for each observation has been allocated to a bin, and therefore a state, the transition matrix $P$ for the Markov chain can be estimated from $N(i, j)$, the number of transitions observed from bin $i$ to bin $j$ using a maximum likelihood estimator (MLE). Defining $n(i) = \sum_j N(i, j)$, this gives the MLE

$$P(i, j) = \begin{cases} \dfrac{N(i, j)}{n(i)} & \text{if } n(i) \neq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

It is important that the wind direction is not biased towards one side of the course for two main reasons. First, if the overall mean wind direction is perceived to have moved during a race, then the course is also moved to ensure that the mean wind direction flows directly down the course. Second, if the yachts start on different sides of the course, then with biased wind, one yacht receives a tremendous advantage over the other, and this advantage can completely dominate differences in design that we are attempting to quantify.

To avoid biased wind direction distributions, the steady-state distribution for the Markov chain should be symmetric about zero. Unfortunately, this is not guaranteed by the estimator (6). A simple sufficient condition for a symmetric stationary distribution is that $P(i, j) = P(K - i, K - j)$, assuming $K + 1$ bins numbered $0, \ldots, K$. Under this constraint, the MLE is

$$P(i, j) = \begin{cases} \dfrac{N(i, j) + N(K - i, K - j)}{n(i) + n(K - i)} & \\ & \text{if } n(i) + n(K - i) \neq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

A Markov chain that evolves according to the transition matrix $P$ defined by (7) models the discrete jumps in the mean wind direction, but it does not model smaller random fluctuations in the wind direction. These random fluctuations are important to model, perhaps primarily for reasons of *accreditation*, i.e., to ensure that end users have faith in the model and its predictions. To model the random fluctuations, a mean-reverting process is used, similar to the wind speed model. If the Markov chain is in state $i$ (with mean wind direction $M_i$), then the new wind direction $M_t^k$ at time step $k$ in the weather clock is generated using

$$M_t^k = \alpha_i + \varphi_1 M_t^{k-1} + \varphi_2 M_t^{k-2} + \vartheta_1 \varepsilon^{k-1} + \varepsilon^k, \tag{8}$$

where $\varepsilon^k \sim N(0, \sigma^2)$, and $\alpha_i = M_i/(1-\varphi_1-\varphi_2)$. The noise terms $\varepsilon^k$ generate the small fluctuations we are seeking. To ensure that the process is symmetrical, the parameters $\varphi_1$, $\varphi_2$, $\vartheta_1$, and $\sigma$ are chosen to be the same for each state $i$.

### 3.3. Correlated Weather

Now consider the modelling of both wind speed and direction on two yachts at different locations. We give a full description of how this is done in Appendix B in the online companion, and give just the key ideas here. Again, we model the wind speed and wind direction as independent quantities.

The wind speed process is a vector process because there is a wind speed for each yacht. When the yachts are far apart, it is reasonable to model the component processes as being independent. However, when the yachts are close together, the component processes should be identical. When the yachts are moderately far apart, some combination of these extremes seems reasonable. So, we model the wind speed as being a vector ARMA process, where the correlation between the two components is a function of the separation of the boats.

Our model of the wind direction process is similar, but is more complicated, owing to the fact that we employ a hidden Markov model to generate wind directions. First, consider the wind direction bin process. Each component of this process is a Markov chain, as described in the previous section. When the yachts are far apart, we allow the two components to evolve independently. When the yachts are close together, the component processes are identical. And again, when the yachts are moderately far apart, we use a convex combination of these two extremes of behaviour. Recall that we also model smaller fluctuations in the wind direction using a mean-reverting ARMA process. We use a vector process approach similar to that used for the wind speed to model these smaller fluctuations.

### 3.4. Yacht Interactions

According to the Taylor hypothesis, the weather conditions that a yacht observes will continue unchanged in the direction of the global mean wind direction at the global mean wind speed. Therefore, the weather conditions observed on a downwind yacht should be related to the weather conditions that are observed on the upwind yacht at an earlier time. In this section, we briefly discuss our approach for dealing with this issue.

To keep our correlated weather model consistent with the Taylor hypothesis, we generate the upwind yacht's weather using the model for a single yacht, and then generate the downwind yacht's weather conditional on a history of the weather observed by the upwind yacht. (This is statistically equivalent to generating both yacht's weather simultaneously, using a two-yacht model.)

To do this, we require a record of the weather conditions observed on the upwind yacht. We call this a *wind trail*. Because either yacht could be the upwind yacht at some stage in a match race, wind trails are created for both yachts. A wind trail is a series of wind records, where each wind record corresponds to an observation in the stationary observer weather series for each yacht. Thus, a wind record is "dropped" every time a new stationary observer weather condition is generated, which corresponds to a time step in the weather clock.

Instead of generating the weather only at the locations of the yachts, as our model does, it is perhaps more intuitive to generate the weather over the entire course. Indeed, one way to do this is to generate the weather at a series of "source locations" located at the top of the course. One can then allow these weather conditions to propagate down the course, exactly as in the case of wind trails. One can then simply "look up" the weather for a yacht at a particular yacht location. In fact, we considered this approach first, but abandoned it for two main reasons. First, generating the weather conditions over the entire course is computationally more demanding than generating weather conditions at only two yacht locations. Second, generating the weather conditions over the entire course requires detailed modelling of both spatial and temporal correlations. In particular, one needs to account for the dependence of the weather conditions at each source location on other source locations, and on previously generated weather conditions that are propagating down the course. Our two-location model does not completely avoid the spatial dependence problem, but we believe that it is a computationally efficient reasonable approximation of wind behaviour at two moving locations under the Taylor hypothesis.

For full details on how the wind trails are maintained within the simulation, and the method used to generate wind conditions on the downwind yacht conditional on the values in the wind trail of the upwind yacht, see Appendix B of the online companion.

### 3.5. Bent and Turbulent Air

The weather conditions observed on any yacht are affected locally by the presence of a second yacht in its vicinity. Weather conditions observed downwind of another yacht are different from those that would prevail in the absence of this yacht. The disturbance created by the upwind yacht depends on its course and sail trim. We attempt to model this disturbed airflow in two situations; first, when the yacht disturbing the air flow is sailing close hauled (upwind), and second, when the yacht disturbing the air is running with the wind (downwind). In both these situations, we assume that it is only the downwind yacht that is affected.

As shown in Figure 7 (reproduced with permission from Twiname 1983, p. 18), the wind passing around the sails of a beating yacht is bent. If yachts are sailing in close quarters, the downwind yacht's wind direction appears to be bent when compared to the wind direction it would observe if the upwind yacht was not present. We call this the *bent air effect*. If the downwind yacht is sailing on the same tack

**Figure 7.** Bent air flow when beating—created using a water model.

*Note.* From E. Twiname, *Start to Win*, p. 18. © 1983 Adlard Coles Ltd., London, U.K. Reprinted with permission.

**Figure 8.** Bent air region model.

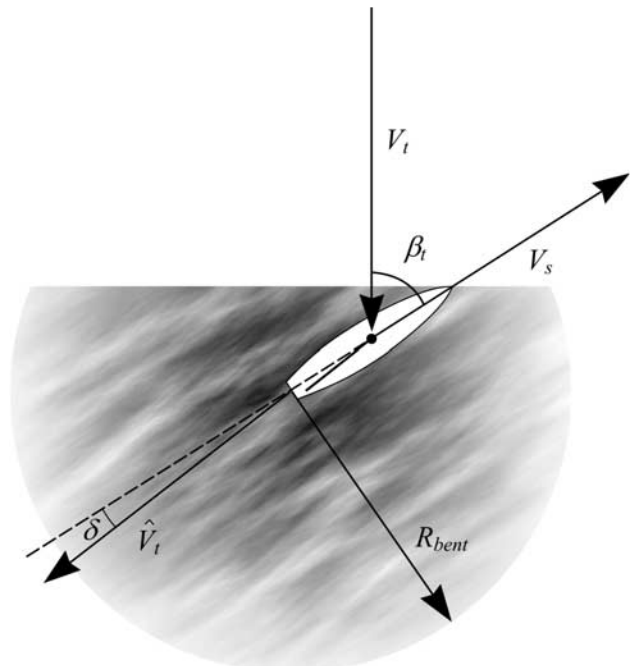


as the upwind yacht, then the bent wind direction observed on the downwind yacht is undesirable because it results in some loss of performance for the downwind yacht. It is therefore important that this effect is modelled.

In modelling the bent air effect, it is assumed that the wind speed is not affected in the bent air region. The bent air effect is modelled solely by altering the wind direction that a downwind yacht observes when sailing upwind and in close proximity to the upwind yacht. The bent air effect is most severe at the trailing edge of the mainsail on the upwind yacht and gradually decreases in severity until it is zero at some distance from the trailing edge of the mainsail. The wind direction at the most severe point is approximately in the direction of the trailing edge of the mainsail.

A simple model for the change in the wind direction is one that decreases linearly with distance from the trailing edge of the mainsail (assumed to be the back of the upwind yacht) until it reaches zero at some distance $R_{\text{bent}}$. At this distance the wind direction returns to normal. Therefore, our model for the bent air region is a truncated circular region that is centered on the stern of the upwind yacht, as seen in Figure 8. The bent air region is truncated because the bent air effect does not extend farther forward than the upwind yacht. Outside of this region the wind direction is what would be observed if there was no upwind yacht present.

The bent air direction at the centre of the bent air region is modelled as the approximate direction of the trailing edge of the mainsail. If one assumes that the mainsail is very flat, then the direction of the trailing edge can be approximated as the direction of the boom. The direction of the boom will be slightly to leeward at some small angle $\delta$ from the centreline of the yacht when sailing to windward. Hence, the direction of the bent air $\widehat{V}_t$ at the trailing edge of the sail is opposite to the yacht heading, with some small offset angle $\delta$ added to leeward to account for the boom angle (see Figure 8). If the downwind observer is at some radius $r$ from the back of the upwind yacht and within the bent air region radius $R_{\text{bent}}$, then the bent air $\widehat{V}_t$ has wind direction $\widehat{M}_t$ given by

$$\widehat{M}_t = M_t + (\beta_t - \delta)\left(1 - \frac{r}{R_{\text{bent}}}\right).$$

A similar model is used to represent turbulent air downwind of a yacht while running downwind. Here the wind is effectively stalled, causing the air flow to be turbulent for some distance downwind. See Appendix B of the online companion for full details.

## 4. Tactics

The previous sections have dealt with modelling the dynamic performance of a yacht and generating the changing weather conditions observed by the two yachts. This section discusses the last major part of the simulation, namely, the modelling of tactical decisions that are made by the helmsman on a yacht. The tactical decisions made by the helmsman need to take into account the performance capabilities of the yacht, the observed weather conditions, and also any possible tactical advantage that can be gained over the opposition.

Tactical decisions made while sailing are an important factor in determining the outcome of a match race. According to Twiname, "a race won is a series of [tactical] decisions correctly made" (1983, p. 37). Conversely, a race

lost is perhaps just one incorrect tactical decision made. To compare yacht designs with different design trade-offs, there must be no bias introduced with regard to different tactical abilities of the two helmsmen. Consequently, each yacht uses an identical tactical decision model.

Although the tactical decision models are identical, it is still important that a realistic model is created so that the yachts sail competitively against one another while attempting to sail each yacht design to its optimum capability. There are two parts to each tactical decision that is made in the simulation. The first part is to choose the best tack for the yacht. The second part is to then determine the optimum heading to sail. With these two parts combined, they effectively tell the yacht the best point of sail for the current situation.

There are three possibly conflicting goals that need to be taken into account when choosing the best tack for a yacht. First, the tack should be chosen to optimize the route and, therefore, minimize the time spent sailing on the course. Second, the tack should maximize the tactical advantage over the opposition (or minimize the tactical advantage of the opposition) by *covering*. (Covering is where a yacht changes tack to reduce the separation between the yachts.) Third, the yacht must obey the rules of yacht match racing so that collision is avoided. When deciding the best tack for the yacht, there needs to be a trade-off between these three goals.

To compare the benefits of staying on the same tack with the benefits that can be obtained by changing to the opposite tack, a series of *penalties* have been created for each of the tactical goals of the helmsman. Each penalty relates to a particular situation that can occur within the scope of each goal. The numerical value that a penalty takes for either the current or the opposite tack depends on the current situation for the yacht with regard to the part of the goal being modelled.

Once the penalties $P_i$ for all of the goals have been calculated, they are traded off against each other using weightings $w_i$ that give the relative importance of each of the situations modelled by the penalties. The weightings $w_i$ for each of the penalties $P_i$ therefore define the decisions that will be made by the helmsman of the yacht in response to different situations. By choosing sensible weightings for the penalties, the decisions that will be made can be reasonably realistic. The total penalty $P^c$ for the current tack is given by

$$P^c = \sum_i P_i^c w_i,$$

where each $P_i^c$ is calculated for the current tack. The total penalty $P^o$ for the opposite tack is given by

$$P^o = \sum_i P_i^o w_i,$$

where each $P_i^o$ is calculated as if on the opposite tack. The tack (current or opposite) with the smallest total penalty is chosen to be the best tack for a yacht in its current situation.

Many of the penalties have been created as Boolean (zero or one) variables. This means that each possible situation modelled by these penalties either exists or it does not exist. In these cases, the choice of weightings to define the decisions made by the helmsman can be easily made to reflect the relative importance of each situation. For example, a collision penalty is very important, so its weighting will be much higher than weightings that apply to less important situations.

We will now describe in the following subsections the different penalties that have been created for the situations that can occur within each of the conflicting goals. It is useful to point out that only the route optimization penalties are needed to choose the best tack for a yacht that is sailing on its own. The covering penalties—which relate to obtaining a tactical advantage over the opposition—and the rule adherence penalties are only needed when sailing against another yacht, such as during a match race.

### 4.1. Route Optimization Penalties

Because a race is sailed in varying weather conditions, it is especially important that a yacht sails the optimal route for the observed weather conditions. This means that a yacht should tack on wind shifts, as long as the expected benefit of being on the opposite tack is greater than the loss incurred during tacking. Because of the loss incurred during a tack, there must be an associated penalty for changing tack. The *tacking penalty* is constructed as a Boolean variable, where the penalty is zero for the current tack and one for the opposite tack.

There are several other penalties relating to route optimization. These penalties are briefly described below. See Appendix C in the online companion for full details.

**Wind Direction.** The wind direction is constantly varying. If the wind direction moves towards the direction that the yacht is moving, then the yacht is then pointing more into the wind and some loss of speed results. This is called a header, and should be avoided if possible, so it is common to change tack in such situations. However, if the change in wind direction does not persist, then changing tack is not recommended. The wind direction penalty penalizes a yacht that is receiving a header, and rewards a yacht that is receiving a lift (i.e., a yacht that is on the opposite tack). This makes the lifted tack more attractive.

**Layline.** When a yacht is close to the mark, the optimum tack changes. When a yacht reaches a *layline*, it is possible to reach the mark without further tacks as long as the wind direction remains constant, and the yacht's current tack is toward the mark. The layline penalty penalizes a yacht that has reached a layline but is not on a tack that carries it towards the mark.

**Comfort Zone.** If a yacht is close to a layline, but has not quite reached it, and is pointing towards the mark, then

one might think that it should tack once to reach the layline, and then tack again to point towards the mark. This consecutive tacking reduces the speed of the yacht dramatically, and so tacking towards the layline should be penalized. The comfort zone is a region that represents the "middle of the course," where tacking towards the layline is considered reasonable. The comfort zone penalty penalizes a yacht if it is on a tack that takes it towards the layline while the yacht is outside the comfort zone.

**Equilibrium Speed.** If a yacht is sailing below its equilibrium speed for the current wind conditions, then tacking is usually avoided if possible.

**Mark Rounding.** Marks need to be rounded in a certain direction. This penalty is designed to encourage a yacht that is close to a mark to be on the correct side of the mark for rounding.

### 4.2. Covering Penalties

To cover is "to maintain a position of advantage with respect to a competitor" (Walker 1976, p. 386). There are two types of covering that can exist during a match race—a loose cover and a tight cover. When the leading yacht becomes separated from the trailing yacht, it is difficult to use the lead to an advantage and therefore help to maintain it. Also, when yachts are separated by a long distance, the weather conditions observed on each yacht can be vastly different. The trailing yacht may observe much better weather conditions than the leading yacht and, as a result, take the lead.

A *loose cover* occurs when the leading yacht keeps in touch with the trailing yacht by keeping the separation between the yachts less than some distance. Thus, if the leading yacht becomes too separated from the trailing yacht due to being on diverging tacks, it will change tack to reduce the separation. If the lead is marginal, then it is common for both yachts to employ a loose cover. We have created a *separation penalty* to model loose covering. See Appendix C of the online companion for details on how this works.

We finish this subsection by formulating a penalty that is incurred if a yacht is in the other yacht's wind shadow. In the wind shadow of a yacht, the air flow is turbulent, so the thrust that can be produced by the sails is reduced compared to the thrust that could be produced when sailing in clean air. The turbulent air effect is most pronounced when sailing downwind, but also exists when sailing upwind. An upwind yacht can attempt to put the downwind yacht into its wind shadow to try and gain an advantage. This is called a *tight cover*. If an upwind yacht employs a tight cover, then the downwind yacht should attempt to get out of the wind shadow. We have created a *wind shadow penalty* to model a tight cover because being in the wind shadow of a yacht is undesirable.

The wind shadow regions created by a yacht are different when it is sailing on a broad reach (downwind) or when sailing to windward (upwind). The wind shadow region created when sailing on a broad reach is a cone that extends from the upwind yacht for some specified distance $R_{ws}$ downwind with the centreline of the cone at some specified angle (typically $\beta_a$) to the yacht track as shown in Figure 9. In our wind shadow region model, the angles to the centreline that specify the edges of the cone (marked by **A** and **B**) are defined to be the same for both upwind and downwind sailing, though this condition can be easily relaxed.

**Figure 9.** Downwind wind shadow region model.

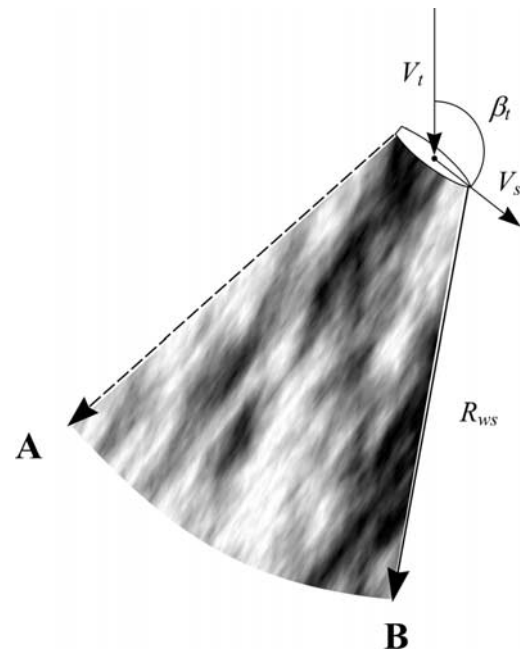


The wind shadow penalty penalizes the downwind yacht and rewards the upwind yacht. When calculating the total penalty, the weighting $w_i$ for the wind shadow penalty differs for both the upwind and downwind yachts. The bonus for the upwind yacht is usually not as important as the penalty that the downwind yacht incurs, and therefore the weighting for the upwind yacht is less. A bonus for the upwind yacht encourages tacking duels. A situation where this might happen is after two yachts cross when sailing to windward. Once the yachts have crossed, the upwind yacht can tack on top of the downwind yacht to put it in its wind shadow. This forces the downwind yacht to tack out of the wind shadow, and so the duel continues. For an animation of such a tacking duel, see TackDuel.avi in the online companion.

Observe that the wind shadow penalties are imposed to influence the tacking decisions of the helmsmen. The bent and turbulent air corrections outlined in §3.5 of this paper affect the aerodynamic forces on the downwind yacht, but the magnitude of these corrections have no effect in the simulation on tacking decisions. In some situations (e.g., when on a layline), a helmsman might have to remain in a wind shadow (because the wind shadow penalties are

outweighed by others) while still sailing in bent wind. This effect can be clearly seen at the end of the animation TackDuel.avi, when the blue boat trails the red boat into the weather mark and loses speed from the bent air behind the red boat. Further details on the wind shadow penalties, including the wind shadow region when sailing to windward (upwind), are given in Appendix C of the online companion.

### 4.3. Collision Avoidance Penalties

We ignore any collisions that could occur due to two yachts sailing on different legs, and concentrate on the two instances when collision could occur when sailing on the same leg. Collisions can happen when the yachts are either in a crossing situation or in close proximity, and in a passing situation. A possible crossing situation occurs when the two yachts are on opposite converging tacks. According to the rules of yacht match racing, in this situation "a port-tack boat shall keep clear of a starboard-tack boat" (International Sailing Federation 2001, p. 5). Consequently, if the yachts will not clear safely in a crossing situation, then the yacht on port tack is penalised. This forces the port-tack yacht to tack away. When this occurs, the starboard-tack yacht usually also tacks away to retain the tactical advantage of being on the right-hand side of the course. This behaviour is known as *bouncing*. However, because the starboard-tack yacht has the right of way, it will not usually bounce when it is close to a layline. As discussed earlier, is does not make sense to tack towards a close layline.

We have created a *crossing penalty* to model this behaviour. The crossing penalty is calculated as a Boolean variable, which is zero unless yachts are on opposite converging tacks and the yachts are within some specified clearance of each other so that a collision will occur. The crossing penalty is given to the port-tack yacht and also the starboard-tack yacht unless it is outside the comfort zone. In the latter circumstance, it is advantageous to remain on starboard tack unless other penalties make it unattractive.

It is conceivable that a yacht with no right of way will *dip* (i.e., pass behind) her competitor rather than tack, leading to the paths of the yachts crossing to opposite sides of the course. In practice, this happens rarely (typically when the skipper of a yacht with starboard advantage sacrifices this to exploit perceived better wind conditions on the other side of the course), so we do not model this in the simulation.

A passing situation occurs when the two yachts are in close proximity and the trailing yacht is travelling faster than the leading yacht. According to the rules of yacht match racing, "when boats are on the same tack and overlapped, a windward boat shall keep clear of a leeward boat. When boats are on the same tack and not overlapped, a boat clear astern shall keep clear of a boat clear ahead" (International Sailing Federation 2001, p. 5).

We have created a *passing penalty* to penalise the correct yacht when in a passing situation. The passing penalty is a Boolean variable, which is zero unless the yachts are on the same tack and get within a certain clearance so that a collision may occur. The penalty is given to the windward yacht if overlapped, and the trailing yacht if not. According to the rules of yacht match racing, a yacht is overlapped if one yacht is neither clear ahead or clear astern of the other. A yacht is clear astern "when her hull and equipment in normal position are behind a line abeam from the aftermost point of the other boats hull and equipment in normal position. The other boat is clear ahead. They overlap when neither is clear astern or when a boat between them overlaps both" (International Sailing Federation 2001, p. 149).

### 4.4. Determining the Optimum Heading

Once the optimal tack is determined by minimizing the total weighted penalty, each yacht's heading must be determined. There are many sailing strategies that can be employed to determine a suitable heading for a yacht. The strategies we consider choose a heading to maximize the component of the yacht's velocity in a specific direction.

One of the most common strategies for sailing on a windward-leeward course is to maximize the component of yacht velocity that lies in the wind direction, i.e., maximize $V_{\mathrm{mg}}$. We call this the $V_{\mathrm{mg}}$ *strategy*. A slightly different strategy that is often employed when the wind direction is oscillating from one side of the course to the other is to maximize the velocity of the yacht in the direction of the course, i.e., maximize $V_{\mathrm{mc}}$ (Kirkman 1987). We call this the $V_{\mathrm{mc}}$ *strategy*. Because the weather is expected to oscillate from one side of the course to the other, this is the strategy that is most commonly employed during the simulation.

The optimum headings to sail for each strategy are computed offline using the equilibrium information generated by the VPP. This is used to construct a lookup table that gives for each direction $\alpha$ in a range, the angle $\tilde{\beta}_t$ to sail with respect to the wind direction so as to maximize the component of velocity in the direction $\alpha$.

To follow the $V_{\mathrm{mc}}$ strategy in the simulation, a yacht has to be sailing on the lifted tack for the observed wind direction relative to the course direction. If, say for some tactical reason, a yacht is not sailing on the lifted tack, then the yacht follows the $V_{\mathrm{mg}}$ strategy. It is useful to point out that if the wind direction is directly down the course, then $\tilde{\beta}_t$ will be the same for both strategies. It is also useful to point out that the laylines are defined by the $V_{\mathrm{mg}}$ strategy.

If a yacht happens to overlay the next mark (sail past the layline) and the yacht is outside the comfort zone, then the optimum heading to sail is chosen to be the direction towards the next mark. However, if the yacht is outside the laylines and still inside the comfort zone due to a very large shift in the wind direction, better progress can be made by sailing as if still inside the laylines (i.e., following strategy $V_{\mathrm{mc}}$ rather than sailing directly to the mark).

Once $\tilde{\beta}_t$ (the desired heading with respect to the wind) has been determined, the desired heading $\widetilde{M}_s$ is then computed by adding or subtracting $\tilde{\beta}_t$ from the observed wind

direction in accordance with the chosen tack. Because a yacht cannot change its heading instantly to the desired heading $\widetilde{M}_s$, the best realisable heading is somewhere between the current heading and the desired heading. If the current heading for a yacht is $M_s$ and the desired heading is $\widetilde{M}_s$, then the best realisable heading $\widehat{M}_s$ for the yacht is modelled as

$$\widehat{M}_s = (1-p)M_s + p\widetilde{M}_s,$$

where $p$ is the heading transfer percentage.

## 5. Results

The final RMP is clearly somewhat complex, and so we spent a great deal of effort verifying that the code worked as intended, and validating the results of the simulation versus what can be expected in reality. We provide only a sketch of the steps here. For further details, see the online companion and Teirney (1999). The code was developed in C++, and was designed to be highly modular. Each module was tested in isolation before being linked to other modules. A particular example of this is the ability of the code to run in "time trial mode," where a single boat is timed around the course. The complexities of boat interactions are not part of this calculation. The boat dynamics were compared with predictions from the VPP on which it was based, and close agreement found. A sensitivity analysis (described in the online companion) was also conducted on the parameters that were most difficult to estimate.

The weather model output was compared both with input parameters to ensure agreement, and also with actual realisations. Plots of the weather model output and true realisations were indistinguishable. We ran special cases where the winning probabilities are known, such as when the boats are identical, and the expected results were obtained. But perhaps the key efforts in verification and validation involved following program traces using debugging tools, and examining the simulated races using an animation tool employed by Team New Zealand to replay races recorded from GPS systems on the yachts. A sample of these can be viewed as avi files in the online companion.

The correlation between the wind conditions experienced on each yacht had a strong influence on the results of simulated races. Recall that the wind speed and direction on the downwind yacht is correlated with the wind speed and direction taken from the closest point on the upwind yacht's wind trail, assumed to be distance $d$ from the downwind yacht. The estimation of the correlation parameters (for wind speed and direction, respectively) and their variation with $d$ is a delicate and challenging problem. We took the simple approach of setting $\lambda = e^{-\mu d}$, where a single $\mu$ (for both speed and direction) was chosen to give weather realisations on both boats that appeared to be close to observed weather observations. When this decay rate was large, the yachts always sailed in essentially independent weather. In this case, unequal yachts appeared more evenly matched than they did when the correlation decay was small. In the latter case, the yacht with a higher boat speed won a much higher proportion of the races. (Details of these experiments are provided in the online companion.)

The interactions between the boats that are represented in a match race simulation are also significant in affecting the probability of winning. This can be seen by comparing the results of a simulated match race and a simulated "drag" race (in which there is no interaction between boats).

Figure 10 shows both the drag and match racing curves with both full and zero correlation for a faster boat versus

**Figure 10.** Comparison between the drag and match race probability curves with both full and no weather correlation, and the predicted winning delta curve.

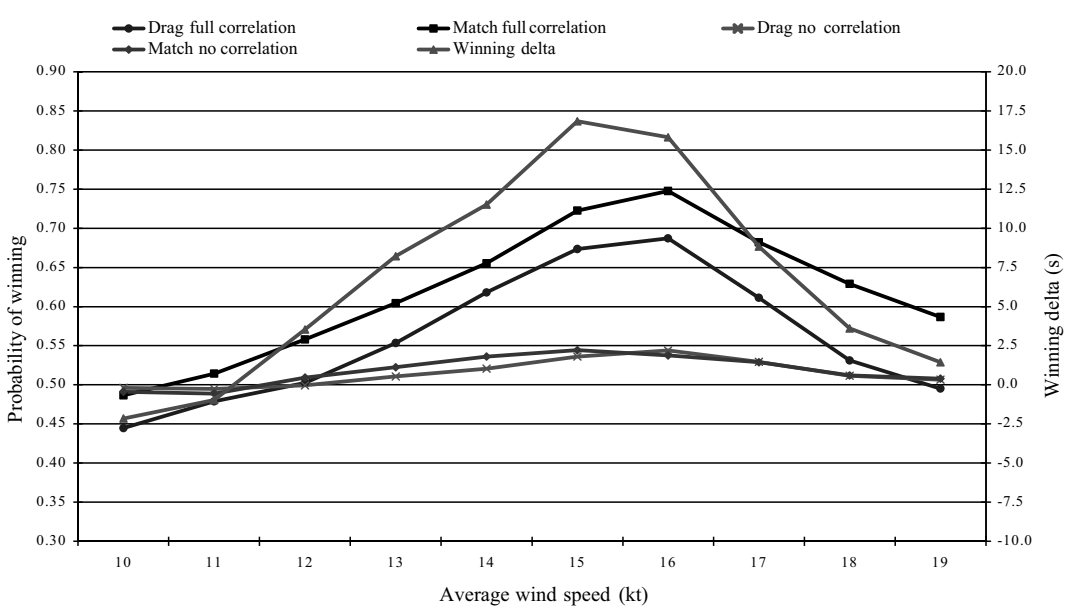

**Figure 11.** Probability of winning with length variation.

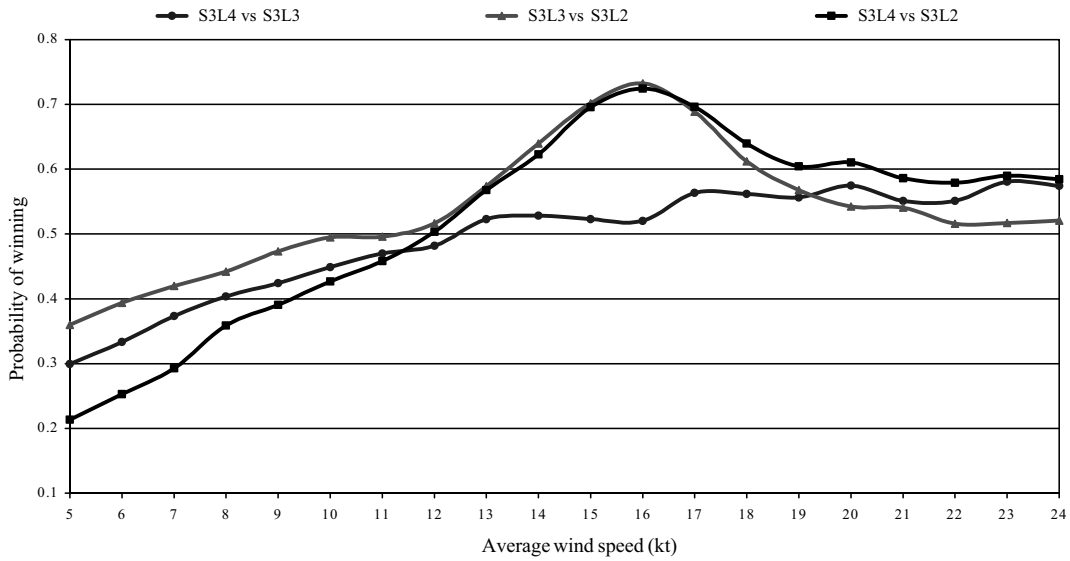


a slower boat. When there is perfect correlation, the yachts will appear to see identical weather, so the probability of winning agrees with the predicted winning delta. (The *delta* gives the time difference obtained by sailing the yachts around the course in a constant wind speed and direction with no interaction between the yachts.) The match racing results produce the same trend while uniformly increasing the probability of winning for the faster boat, reflecting the tactical advantages that the extra speed bestows.

In the case where the yachts see no correlation in the weather conditions it appears that the advantage of the faster yacht has been reduced due to the random nature of the weather observed on each yacht, which swamps the speed difference between the yachts.

The principal aim of the RMP was to provide an estimate of the probability that a candidate design would beat a given competitor in a given average wind speed. In the 1995 America's Cup in San Diego, CA, the wind speed had a low mean and low variance. The Hauraki Gulf, which was the venue for the 2000 and 2003 America's Cup regattas, experiences higher average wind speeds that are also more variable. Designers at Team New Zealand were interested in the trade-offs to be made in designing a heavy-weather boat and a light-air boat. These are illustrated in Figure 11,

**Figure 12.** Delta values of course times for pairs of yachts with length variation.

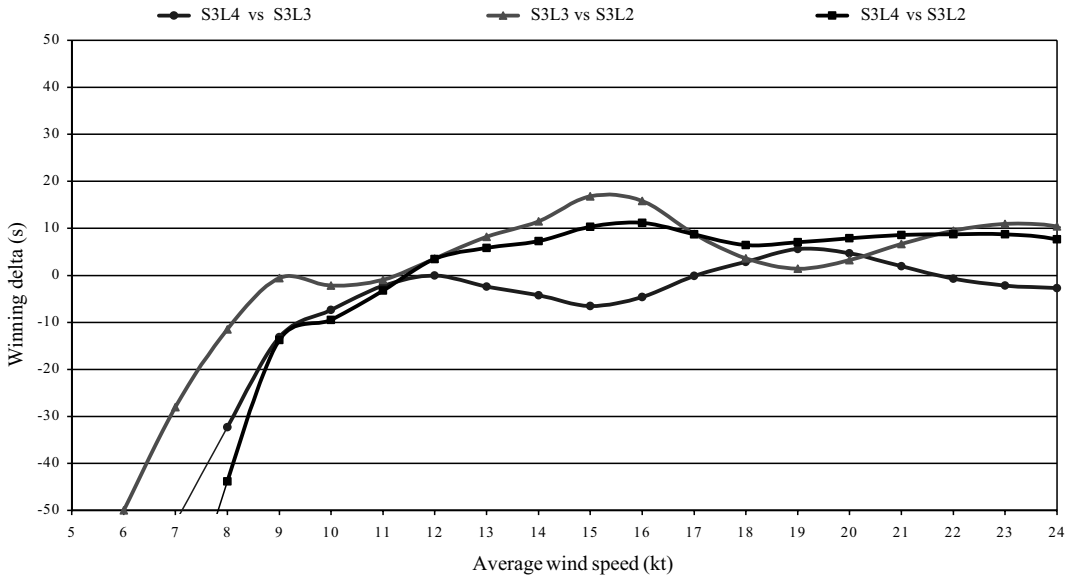

which shows win/loss probabilities estimated by simulating 10,000 races for three candidate designs with different lengths. Here S3L2 is shorter than S3L3, which is shorter than S3L4. The probability estimates in these plots have a 95% confidence interval of $\pm 0.01$.

The figure shows that the longest boat performs well in stronger winds, but is beaten by a short boat in lighter airs. One reason for this is that America's Cup design rules (see www.americascup.com, 2003) penalize long boats by requiring them to have a smaller sail area. Moreover, hydrodynamic drag is a decreasing function of waterline length for high boat speeds, but is less sensitive to waterline length when speeds are small (and frictional drag dominates).

It is instructive to compare the win/loss probabilities shown in Figure 11 with the course deltas shown in Figure 12. From this graph, one would predict that S3L3 would beat S3L4 in 15 knots of wind. However, from Figure 11 we can see that the boats appeared to be evenly matched in this wind speed. In fact, S3L4 is slightly faster upwind than S3L3 in this wind speed, but substantially slower downwind. However, the estimated time around the course for each boat does not capture the tactical advantages of upwind speed in a match race that are reflected to some extent in the RMP result.

## Acknowledgments

The authors acknowledge the contributions of Clay Oliver and Tom Schnackenberg to this work. This work was partially supported by New Zealand Public Good Science Fund Grant UOA 803 and Graduate Research in Industry Fellowship TEA701. The work of the second author was partially supported by National Science Foundation Grant DMI 0230528. The authors thank the editor and referees for their careful reading and constructive suggestions.

## References

Bilger, R. W. 1995. A simplified theory for the tacking of a sailboat. *Proc. Twelfth Australasian Fluid Mech. Conf.* The University of Sydney, Sydney, Australia.

Davies, N. 1990. A real-time yacht simulator. Master of Engineering thesis, School of Engineering, University of Auckland, Auckland, New Zealand.

Gretzky, J. A., J. K. Marshall. 1993. The partnership for America's cup technology: An overview. *Proc. 11th Chesapeake Sailing Yacht Sympos.* Annapolis, MD.

International Sailing Federation (ISAF). 2001. *The Racing Rules of Sailing for 2001–2004*. www.sailing.org/rrs2001/rrs2001to2004.pdf.

Kerwin, J. E. 1978. A velocity prediction program for ocean racing yachts revised to February, 1978. M.I.T. Ocean Engineering Report Number 78-11, MIT, Cambridge, MA.

Kirkman, K. L. 1987. The application of VPPs to practical sailing problems. *Proc. 8th Chesapeake Sailing Yacht Sympos.* Annapolis, MD.

Lawson, T. V. 1980. *Wind Effects on Buildings: Statistics and Meteorology.* Applied Science Publishers Ltd., London, U.K.

Letcher, J. S., Jr. 1974. Handicapping rules and performance of sailing yachts. *Proc. 1st Chesapeake Sailing Yacht Sympos.* Annapolis, MD.

Letcher, J. S., Jr., J. K. Marshall, J. C. Oliver III, N. Salvesen. 1987. Stars and stripes. *Sci. Amer.* **257**(2) 34–40.

Marchaj, C. A. 1996. *Sail Performance: Techniques to Maximize Sail Power.* International Marine, Camden, ME.

Milgram, J. H. 1993. Naval architecture technology used in winning the 1992 America's Cup match. *Proc. Centennial Meeting of SNAME*, September 14–19. New York, 1–36.

Philpott, A. B., P. S. Jackson, R. M. Sullivan. 1993. Yacht velocity prediction using mathematical programming. *Eur. J. Oper. Res.* **67** 13–24.

Schlageter, E. C., J. R. Teeters. 1993. Performance prediction software for IACC yachts. *Proc. 11th Chesapeake Sailing Yacht Sympos.* Annapolis, MD.

Teirney, D. P. 1999. Yacht match race simulation. Master of Engineering thesis, School of Engineering, University of Auckland, Auckland, New Zealand.

Twiname, E. 1983. *Start to Win.* Adlard Coles Ltd., London, U.K.

van Oossanen, P. 1993. Predicting the speed of sailing yachts. *Proc. Centennial Meeting of SNAME.* New York.

Walker, S. H. 1976. *Advanced Racing Tactics.* Angus & Robertson Publishers, London, U.K.

www.americascup.com. 2003.

www.americascup.com/datas/textes/acc5_final_draft.pdf.