



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Development of a routing software for inshore match races

Citation for published version:

Tagliaferri, F & Viola, IM 2016, 'Development of a routing software for inshore match races', Paper presented at 22nd Chesapeake Sailing Yacht Symposium, Annapolis, United States, 18/03/16 - 19/03/16.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.





THE 22nd CHESAPEAKE SAILING YACHT SYMPOSIUM

ANNAPOLIS, MARYLAND, MARCH 2016

DEVELOPMENT OF A ROUTING SOFTWARE FOR INSHORE MATCH RACE

Francesca Tagliaferri, Newcastle University, Newcastle upon Tyne, UK

Ignazio Maria Viola, The University of Edinburgh, Edinburgh, UK

ABSTRACT

Yacht races are won by good sailors racing fast boats. A good skipper takes decisions at key moments of the race based on the anticipated wind behaviour and on his position on the racing area and with respect to the competitors. His aim is generally to complete the race before all his opponents, or, when this is not possible, to perform better than some of them. In the past two decades some methods have been proposed to compute optimal strategies for a yacht race. Those strategies are aimed at minimizing the expected time needed to complete the race and are based on the assumption that the faster a yacht, the higher the number of races that it will win (and opponents that it will defeat). In a match race, however, only two yachts are competing. A skipper's aim is therefore to complete the race before his opponent rather than completing the race in the shortest possible time. This means that being on average faster may not necessarily mean winning the majority of races. This paper presents the development of software to compute a sailing strategy for a match race that can defeat an opponent who is following a fixed strategy that minimises the expected time of completion of the race. The proposed method includes two novel aspects in the strategy computation:

- A short-term wind forecast, based on an Artificial Neural Network (ANN) model, is performed in real time during the race using the wind measurements collected on board.
- Depending on the relative position with respect to the opponent, decisions with different levels of risk aversion are computed. The risk attitude is modeled using Coherent Risk Measures.

The software is tested in a number of simulated races. The results confirm that maximising the probability of winning a match race does not necessarily correspond to minimising the expected time needed to complete the race.

NOMENCLATURE

Acronyms

ANN	artificial neural network
BS	boat speed
DP	dynamic programming
RMP	race modelling program
TWA	true wind angle
VMG	velocity made good

Symbols

\mathbb{E}	expected value
$C(U, \omega)$	cost function
G_x, G_y	grid point matrices
t_k	sailing time at step k
$U = u_0, \dots, u_{N-1}$	policy
U^{opt}	optimal policy
\mathcal{U}	set of admissible policies
(x, y)	position coordinates
(x_0, y_0)	initial position
$(x_{1L}, y_{1L}), (x_{1R}, y_{1R})$	reachable nodes coordinates
w	wind vector
ω, ω_k	random variables

INTRODUCTION

A yacht race is a competition where two or more boats race each other to complete a certain course in the shortest time. Traditionally, the problem that a sailor has to solve is addressed as an optimisation problem consisting in going from point A to point B in the shortest possible time, under certain constraints given by the dynamics of the yacht and racing rules.

This approach however doesn't really capture the competitive aspect of a race. In fact, the real aim of a sailor is not to get to the finish line as fast as possible, but rather to get there before their opponent(s). Moreover, the speed of a sailing yacht is highly dependent on the behaviour of the wind. A sailor doesn't have perfect knowledge of the future wind patterns, and therefore the problem must be addressed as a stochastic problem, based upon probability distributions of the wind behaviour, as done, for instance,

by ?.

In previous studies (Tagliaferri et al., 2015) the authors have shown how the accuracy of a wind forecast can improve the chances of winning a race. The use of artificial neural networks (ANN), compared to other forecasting techniques, was identified as suitable for very-short-term wind prediction (order of seconds/minutes in advance). It was also shown (Tagliaferri et al., 2014) that strategies with different risk tolerance can be computed, but that the strategies that aim at minimising the time needed to complete a race are not necessarily the ones that lead to a higher chances of winning.

This paper focuses on the development of a methodology that allows the computation of a strategy for a sailor, combining the ANN wind forecast and optimal risk modelling. For the first time the presence of a moving opponent is included in the computation of a strategy, and the opponent is not only seen as a moving obstacle, but also as an element of influence in the yacht's speed. The computation of the optimal strategy is based on dynamic programming (DP) over a time-dependent lattice, which is generated according to an ANN-based wind forecast.

Background on yacht racing

Yacht races are held in many different formats and levels: in the case of a *match race* only two boats face each other, while in a *fleet race* the number of participants can be very high. One of the most prestigious sailing competition (and by far the oldest and most expensive) is the America's Cup, which includes match races between various teams fighting for the chance of challenging the Cup defender, i.e. the winner of the previous edition.

Usually, a race course includes several turns around an upwind and a downwind mark, where the marks are aligned with the wind. The course is designed to present some challenges to the skippers. In fact, the speed of a sailing yacht depends on the wind speed and the True Wind Angle (TWA, the supplementary angle between the wind velocity and the boat heading). Figure 1 presents an example of boat speed (BS) as a function of the TWA for a given wind speed in a polar diagram. A polar plot of this kind, which may include different curves associated to different wind speeds, is the conventional way of presenting the boat speed, and although the actual BS can depend on other factors (such as waves and crew), it is considered as a characteristic of a yacht.

As shown in the plot, the highest values for the BS are achieved when sailing at a TWA of approximately 90° (on a *beam reach*). Conversely, when the TWA tends to zero, BS tends to zero. Therefore, when sailing upwind (for instance, from a downwind mark to an upwind mark), the most effective route consists in a zig-zag in the wind direction, sailing at a TWA of 35° - 50° (*close hauled*). In this case, a skipper's aim is to maximise the speed in the

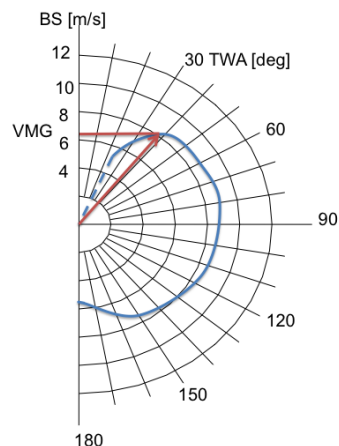


Figure 1: Example of polar diagram.

upwind direction, which means to find the TWA such that the projection of the boat velocity on the upwind direction is a maximum. The corresponding velocity is referred to as Velocity Made Good (VMG) and is shown in red in Figure 1. Similarly, a VMG can be defined for downwind sailing as the projection of the boat velocity on the downwind direction.

The VMG can be defined also for downwind sailing. In

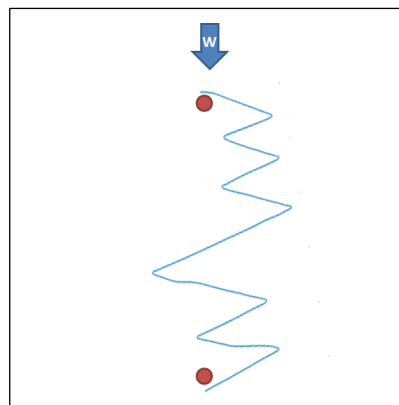


Figure 2: Example of upwind leg.

fact, as shown in Figure 1, even if the velocity is not null when the TWA is 180° , the maximum projection on the downwind direction for this example is obtained at angles of approximately 150° . However, the optimal angle for downwind sailing can have significant variations depending on the yacht geometry.

Yacht racing strategy

Initial research related to competitive sailing was mainly focussed on understanding a yacht's dynamics, in order to improve the design process and create faster and more efficient boats. The outcome of this research was the development of Velocity Prediction Programs (VPPs), computer programs that solve the equations of motion for a sailing yacht, and determine its velocity for given wind conditions. Kerwin and Newman (1979) present one of the earliest studies detailing the development of a VPP. The introduction of such tools determined a big step forward for competitions such as the America's Cup. In fact, The United States had always successfully defended the trophy since its creation, in 1851, but in 1983, for the first time, an Australian team managed to win, and this success was mostly due to the radical innovation in keel design for the yacht *Australia II*, (Oossanen and Joubert, 1986). Americans learned from this defeat, and the following campaign saw a massive effort in applying state-of-the-art technology to the design of the yacht that would challenge the Australians. The resulting yacht, *Stars and Stripes*, succeeded in the mission of bringing the Cup back to the US in 1987, after a campaign that was the first to see a competition not only between sailors but also between the engineering teams of the different countries. This aspect is passionately described in the paper *Stars and Stripes* by Letcher et al. (1987) which focuses on the advantages brought by computer technology.

Since then, the competition has evolved along those lines, and today it is still a fierce battle between engineering teams besides sailing teams, and in sports journalism America's Cup races are often compared to Formula One GP. The development of VPPs allowed design teams to compare different design choices at an early stage, and techniques used for determining forces acting on the boats are determined using a variety of techniques, both experimental and numerical.

The evolution of VPP led to *Race Modelling Programs* (RMP), computer programs aimed at simulating an entire race between two yachts. The *Stars & Stripes* campaign involved one of the very first RMP to analyse the probabilities of win/loss of a yacht.

The subsequent America's Cup saw the development of a RMP which included a statistical weather model based on site-specific environmental data for San Diego. This RMP was developed by the Partnership for America's Cup Technology, and details are described by Gretzky and Marshall (1993).

In those models, the tactical decision process is modeled as a set of fixed decision rules. The tactical and physical interactions between the yachts are not adequately modeled, and this limitation is reflected in the definition of win in Letcher et al. (1987), where a yacht has to win by a certain time margin to be certain of a win.

An important contribution to RMP came from the studies

carried out at The University of Auckland in collaboration with the New Zealand challenger team. Philpott and Mason (2001) investigated the decision-making process, focussing on the development of a strategy. In this work, which constitutes a fundamental basis for the present study, dynamic programming is used to generate a *policy*, that can be computed before the race, and can then be used during the race. Later, Philpott et al. (2004) developed a model to predict the outcome of a match race between two competing designs, still assuming a set of fixed decision rules but taking into account some interactions between yachts (for instance, when crossing).

In these two studies the tactics and strategy modeling was aimed at obtaining a simulation tool that could replicate as closely as possible the situations that can arise during a yacht race, with the ultimate objective of assessing competing designs. Other studies not directly related to the America's Cup have tackled the problem of decision-making for sailors. Ferguson and Elinas (2011) propose a simple Markov decision model, where at all times the sailor has only two options, "do nothing" or "tack". The work developed in this study, focussed on inshore racing, includes a VPP and a model for wind flow around landmasses. The importance of the tacking penalty is investigated by comparing routes produced by assuming different penalty factors associated to tacks.

Recently, the University of Southampton has developed a sailing simulator called "Robo-Race", a tool to model both the physical behaviour of a yacht and the interaction with the crew (Scarponi et al., 2007a,b). The tool is designed so that human sailors can interact with it, racing against a computer in an artificial environment. A VPP using four degrees of freedom is implemented, including the tacking model based on the studies of Masuyama (1995). Improvements on the first implementation are focused on physical interactions between racing yachts (Spenkuch et al., 2008, 2011), and the dynamics of the yacht during manoeuvring (Banks et al., 2010, Spenkuch et al., 2010), both for upwind and downwind sailing. An important contribution of this work is recognising the existence of conflicts between strategy and tactics, for instance when a yacht decides to tack to avoid the blanketing effect from another yacht, but doing so it incurs in an unfavourable wind.

METHOD

Figure 3 shows the boundaries of the race area used in this work. The dimensions used are inspired by the typical length of a 35th America's Cup race area. The distance between the starting point and the upwind mark is of 5000 m, and the width of the area is 3000 m. The course is assumed to be aligned with an average initial wind direction which is kept constant for the entire race. The area as shown in the Figure is delimited by ideal laylines, but in some cases

the actual routes goes beyond those lines. There is a limited tolerance (100 m) on the side boundaries for ease of grid computation.

The DP algorithm is based on a shortest path problem de-




Figure 3: Racing area

finied on a set of nodes connected in a lattice. The set of nodes is not fixed, but their position depends on the wind forecast. Before formally describing the process of grid definition, an example to motivate this choice will be shown. Let us consider the final phase of the upwind leg when the boat is reaching the mark in the case of a gradual wind shift towards the left. Figure 4 shows the optimal route towards the mark with two different underlying grids. In the left grid, the optimal route does not go through the nodes defined by the grid, therefore a certain approximation in the DP algorithm is needed. Conversely, the grid on the right shows an




Figure 4: Comparison between grid with fixed spatial steps (a) and with wind-dependent steps (b).

exact superposition of the route and one of the lines consti-

tuting the grid. Ideally, the nodes defined by the grid should correspond to the reachable points on the racing area. Of course the racing area is a continuum, so every point within the race boundaries is always reachable, but the discretisation should be developed so that, if the yacht is in a given node belonging to the set of nodes defined by the grid, then the neighbour nodes should be reachable from that node. This property is not satisfied by the left grid in Fig. 4, but it is satisfied by the right grid, as shown by the red path followed by the yacht to reach the upwind mark. A curvilinear grid that matches the optimal route can be drawn if the future wind evolution is known.

The grid defining the lattice used in this work is therefore based on the wind forecast, in order to predict the possible reachable points. The grid is then recomputed every time step. The main assumption underlying the construction of the grid is that, in the absence of tactical interactions due to the presence of a competitor, a skipper will always sail at maximum VMG. Figure 5 shows how to build the subse-




Figure 5: Construction of grid points.

quent grid points given an initial node of coordinates (x, y) . If the forecast wind when reaching the point (x, y) is represented by the wind w , then the possible reachable points in a given time step dt have coordinates $(x_{1L}, y_{1L}), (x_{1R}, y_{1R})$ depending on the current tack. Let us assume that the boat is on a port tack. Then in a period of time of $2dt$ the points $(x_{2L}, y_{2L}), (x_{2R}, y_{2R})$ can be reached. w is the wind which is *expected* at the moment when the grid is generated. A subsequent forecast could predict a different wind (e.g. w' in Figure 5), in which case the reachable points become $(x'_{2L}, y'_{2L}), (x'_{2R}, y'_{2R})$. This is why the grid construction is updated at every step.

If every node generated two subsequent nodes, the size of the grid would grow exponentially at each iteration. Rather than building the grid point by point, the grid is therefore built by defining a set of lines and then considering their intersections as the nodes constituting the graph underlying the DP algorithm. Figure 6 shows the construction of the initial grid. A set of M_0 evenly spaced points is defined on the x axis, where M_0 depend on the desired grid resolution. At step one of the computation the following operations are performed:




Figure 6: Example of grid construction.

1. Yacht position: $(x_0, y_0) = (0, 0)$
2. Generate wind speed and direction forecast
3. Compute grid lines
4. Compute lines intersections. These points constitute the DP nodes
5. Store grid points in matrices G_x, G_y

The distance between grid points depends on the chosen time step. The step used for the simulations for this work is $dt = 5s$. This is the time step used in the computation of the optimal strategy and does not necessarily correspond to the time step used for the wind forecast.

The grid is built starting from the current position of the boat. The objective of the boat is to round the upwind mark clockwise. The mark itself is not necessarily a point of the grid. However, by construction the mark will lay between four grid nodes defining a grid cell. The leftmost node is considered the arrival node at each iteration.

At each step k of the computation the grid is re-computed. The current position of the boat, (x_k, y_k) becomes the initial node. The equivalent of the initial points laying on the x axis are now a set of points evenly spaced (according to wind speed as for the first step) laying on the line of equation

$$y = \tan(TWA_k)(x - x_k) + y_k \quad (1)$$

where TWA_k indicates the TWA at time step k .

Dynamic programming algorithm

Let us consider a dynamic system evolving according to the following Equation 2:

$$x_{k+1} = f_k(x_k, u_k, \omega_k), \quad k = 1, \dots, N-1 \quad (2)$$

where k represents a discrete step, x_k and x_{k+1} represent the state of the system at steps k and $k+1$, respectively, u_k represents a decision, also called *control*, and ω_k is a random variable influencing the evolution of the system, characterised by a certain probability function p_k . The step index may refer to an increment over time or space, and the

increment doesn't need to have fixed amplitude. Usually the initial state x_0 is fixed. All the variables defined take values in some determined interval or space; in particular, for a given state of the system x_k , the set of admissible controls $\mathcal{U}_k(x_k)$ is defined as the set containing all the possible decisions that can be taken at that stage. For instance, in financial problems, $\mathcal{U}_k(x_k)$ may be the set of all the possible assets that it is possible to buy or sell. In sailing applications, x_k can represent the state of a yacht on the race area (in this case x_k can be the vector constituted by the yacht's coordinates and the observed wind, assuming values on a limited subset of \mathbb{R}^n), u_k the course followed by the skipper ($u_k \in \mathcal{U}_k \subseteq [0, 360)$), and ω_k the unknown wind evolution between step k and step $k+1$. The position of the yacht at step $k+1$ is then a function of those three variables.

A control, or a *policy*, is a finite sequence $U = u_0, \dots, u_{N-1}$, where $u_k = u_k(x_k)$ is a function of the current state of the system, and all the $u_k \in \mathcal{U}_k(x_k)$ for all x_k . In the following, \mathcal{U} will denote the set of the admissible policies.

The aim of DP is to find an admissible policy $U = u_0, \dots, u_{N-1}$ that minimises a cost function which can assume the generic form as expressed in Equation 3:

$$C(U, \omega) = \sum_{k=0}^N c_k(x_k, u_k(x_k), \omega_k) \quad (3)$$

where $\omega = [\omega_0, \dots, \omega_N]$, subject to the system constraint specified in Equation 2. In sailing, this cost corresponds to time:

$$T(U, \omega) = \sum_{k=0}^N t_k(x_k, u_k(x_k), \omega_k) \quad (4)$$

where t_k represents the time needed to sail from state x_k to state x_{k+1} .

For this class of problems, the cost function is known at every stage. Unfortunately, in practical applications (including sailing) the cost function is only known in terms of a probability distribution, and rather than minimising a cost the aim is to minimise its expected value. In this case, the stochastic version of dynamic programming is used. Going back to the general description, a solution for the problem is then a policy U^{opt} such that

$$\mathbb{E}(C(U^{opt})) = \min_{U \in \mathcal{U}} \mathbb{E}(C(U, \omega)) \quad (5)$$

We assume that the minimum in Equation 5 is well defined. A discussion of this aspect can be found in Bertsekas (2007). According to the principle of optimality, an optimal solution has the property that, considering the subproblem starting at stage M , then the subpolicy $(U^{opt, M} = (u_M^{opt}, u_{M+1}^{opt}, \dots, u_N^{opt}))$ is optimal for that subproblem. The expected values in Equation 5 are computed by using a Markov model for the distribution of wind speed and direction. The Markov model is derived from wind data as detailed in Tagliaferri et al. (2014). This model is also the basis

for the risk model. In fact, by using coherent risk measures, the transition matrix for the Markov process is multiplied by a transformation matrix which has the function of shifting the probabilities of favourable/unfavourable events. A complete description of this procedure can again be found in Tagliaferri et al. (2014). The optimal transformation is found among a set of matrices heuristically selected according to the following principles.

A boat skipper who is losing will seek risk. If she adopts a minimum expected finish time strategy against another skipper who minimises his expected time to finish, then she will tend to make the same decisions (unless the boats see very different winds) and lose the race almost certainly. She will instead seek different wind conditions from the competitor, being optimistic about the possible advantageous wind shifts and assigning a higher probability to these outcomes (i.e. lifting shifts). Being optimistic about random outcomes increases risk, as well as incurring some loss in expected performance.

A sailor who is losing will seek risk. This corresponds to increasing her confidence of a lifting wind shift while discounting the likelihood of a heading wind shift. The transition matrices used to represent the two attitudes are shown in Figure 7. Advantageous shifts (cells below the diagonal when the skipper is to the left of the opposition, and cells above when on the right) happen with higher probability than in the risk-neutral case. The remaining probabilities in each row are reduced to add to one. Following the notation in the aforementioned paper, the transition matrices are represented by using a gray scale, where darker colours represent higher probabilities.

A set of rules aimed at avoiding collisions between the




Figure 7: Modified transition matrices for a risk-seeking skipper. Advantageous wind shifts occur with higher probability than disadvantageous ones. (a) Yacht on the left-hand side of competitor and (b) yacht on the right-hand side of competitor.

boats and at respecting the racing rules are implemented. In particular:

1. If the two boats meet, then the boat on a port tack increases the TWA, passing behind the other boat.
2. A boat cannot tack if this leads to its track crossing the opponent's under a certain fixed safety distance.

The safety distance is defined noting that the boats are modelled as points. The longitudinal safety distance is 10 m, the side distance is 5m.

The computations for the manoeuvre of bearing away and passing behind the opponent's boat is carried out by adding a node to the set of reachable nodes. This temporarily modifies the assumption that a boat always sails at maximum VMG. In the example shown in Figure 8, the red yacht expects to meet the opponent at the node indicated by the red dot. The black node is therefore added to the set of the reachable points. The model for physical interactions

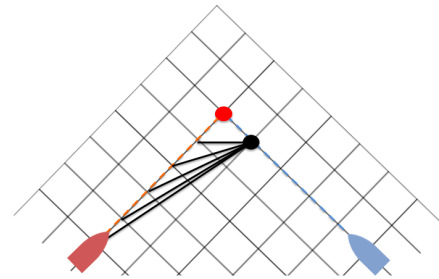


Figure 8: Example of grid modification when two boats meet.

between the two yachts competing in a match race is based on the experimental results of wind tunnel tests presented in Aubin (2013). These results are based on a series of wind tunnel experiments carried out in the Twisted Flow Wind Tunnel facility of Auckland University. The complete set of experiments is described in Aubin (2013).

The opponent is assumed to follow a strategy aimed at minimising the expected time to complete the race. This means that he is expected to follow a reasonable route that depends on the forecast wind. To compute an optimal strategy, it may be possible to forecast the future position of the opponent, to take into account that the two boats might meet further in the future. However, in order to properly take into account such events, the computation of a probability distribution is required. In fact, let's assume that with the wind conditions forecast at the beginning of the race the two boats can compute an optimal strategy which will lead them to meet in proximity of the upwind mark. This event will actually happen with a probability which is equal to the probability that the wind realisation is exactly the one forecast at the beginning and that the boats actually follow the computed strategy. The further this event is in the future, the closer this probability is to zero. In the current software implementation, the future window is set at one minute, which is the time frame at which the wind forecast has an average error lower than 2° (Tagliaferri et al., 2015) and because a yacht is expected to perform not more than one tack in one minute.

An important hypothesis, not necessarily corresponding to reality, is that no yacht will take a decision that leads to a higher expected time with the aim of slowing the other

yacht down.

RESULTS

Routing examples from America’s Cup data

The algorithm is tested using recorded wind scenarios from the past edition of the America’s Cup held in San Francisco. A strategy based on the ANN forecast is compared with a strategy which assumes perfect knowledge of the wind behaviour.

The results of the simulated races are summarised in Table 1. Only 12 upwind legs are simulated using the initial minutes of the last 12 races, as the data relative to the other races was used for training. The average difference be-

Table 1: Simulated races with San Francisco wind dataset.

Race [deg]	1	2	3	4	5	6	
Time (perfect wind knowledge) [s]	603	743	618	743	642	818	
Time (ANN forecast) [s]	608	747	634	781	648	821	
Difference [%]	0.8	0.5	2.5	4.8	0.9	0.4	
Race [deg]	7	8	9	10	11	12	13
Time (perfect wind knowledge) [s]	597	661	654	712	748	684	697
Time (ANN forecast) [s]	608	672	659	715	761	693	712
Difference [%]	1.8	1.6	0.7	0.4	0.3	1.2	2.1




Figure 9: Routes computed using forecast and assuming perfect wind knowledge for Race 2 (a) and Race 4 (b).

tween the different times to completion between a boat with a perfect knowledge of the wind and a boat which uses the ANN forecast is 10.7s. A representative example is given in Figure 9(a), corresponding to Race 2. In this example, the difference between the two strategies is limited to a slight delay of the second tack when the ANN forecast is used.

One of the worst cases is shown in Figure 9(b), corresponding to Race 4. The black trajectory is the one computed by the algorithm having perfect knowledge of the future wind, while the red dashed one is computed by the algorithm using the ANN forecast. The ANN-based algorithm leads to an extra tack at the beginning of the race, due to a wrong forecast for the end of the race. However, the error is soon recovered and in the final part of the race the two strategies become almost indistinguishable.

Optimal risk model

The optimum risk management is investigated considering two boats racing each other. At every step of the simulated race, if A is more than 15 s behind B, she uses the risk-seeking, optimistic matrix for the relevant side of the course. If B is more than 15 s behind A, she uses the risk-averse, pessimistic matrix. For this case, $T_{switch} = 15$ s. The time difference and the matrix transformations are arbitrarily fixed, and the results obtained confirm the results presented in Tagliaferri et al. (2014). Figure 10(a) shows differences between the arrival times of boats A and B. When this time difference is positive, it means that A wins the race. Conversely, if the time difference is negative, B wins the race. This set of results confirms that a risk seeking attitude can constitute an advantage for a skipper who is losing the race. However this advantage can be optimised by changing the amount of risk, i.e. how much the new Markov matrices differ from the original one, and the time at which the attitude is changed. i.e. the time difference between the two boats that triggers the attitude switch.

The amount of risk is investigated by comparing strategies obtained using matrices that have been multiplied for the transformation matrix multiple times. The best outcome, is obtained with the use of the matrix shown in Figure 10(c), and by setting $T_{switch} = 10$ s. The optimised risk model leads to the distribution in Figure 10(b), which corresponds to a win for boat A in 74% of the cases, and is obtained by post-multiplying the risk-neutral matrix for the square of the original transformation. This optimisation was carried out over a limited set of possibilities, and it must be the subject of further research.

Example of upwind leg

In this Section a complete race between two boats that follow the optimum course and have an optimum management of risk is presented. The wind which was measured during the last race of the America’s Cup is used. Figure 11 shows the wind direction and how this is forecast by the two as shown in Tagliaferri et al. (2015). The red bars highlight the critical points where the forecast error is higher than 3° for two consecutive minutes. Differently from the original America’s Cup race where the wind was recorded, in this example the race starts at minute 18, and is made of only one upwind leg. The wind shows a significant shift towards

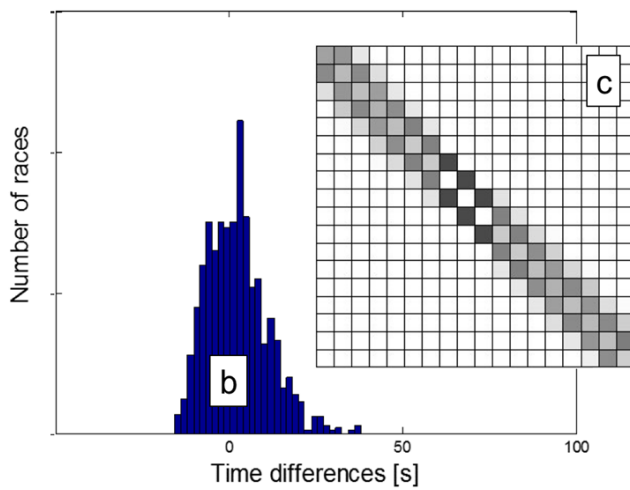



Figure 10: Histograms of time differences for risk-neutral strategy vs optimistic-pessimistic combination (a) and optimal optimistic-pessimistic combination (b) based on an optimal processing matrix (c)

the right during the race, while the wind speed variations are negligible. Figure 12(a) shows the grid corresponding to the wind realization. The grid shown is coarser than the one computed by the algorithm for clarity.

Both boats A (blue in Figure 12) and B (red in Figure 12) start on a starboard tack (sailing towards the left) and boat B is on the left of boat A, at a distance of 25 m (distances between the two boats are magnified in the figures for clarity). Figure 12(a) shows the beginning of the race and the grid represents how the wind was forecast at that time. Both boats begin the race sailing on a starboard tack. In fact, as a significant increasing shift towards the right is forecast, the best strategy consists in approaching the mark on a starboard tack. Boat A chooses to sail towards the left

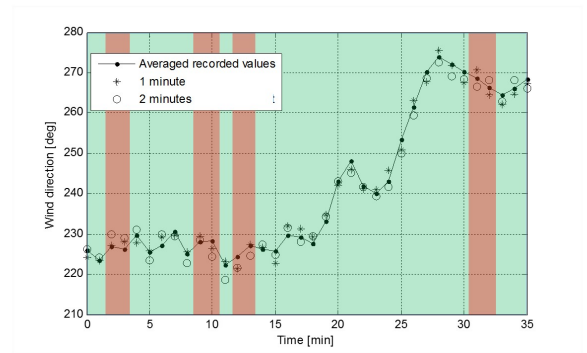


Figure 11: Wind forecast example (Tagliaferri et al., 2015).

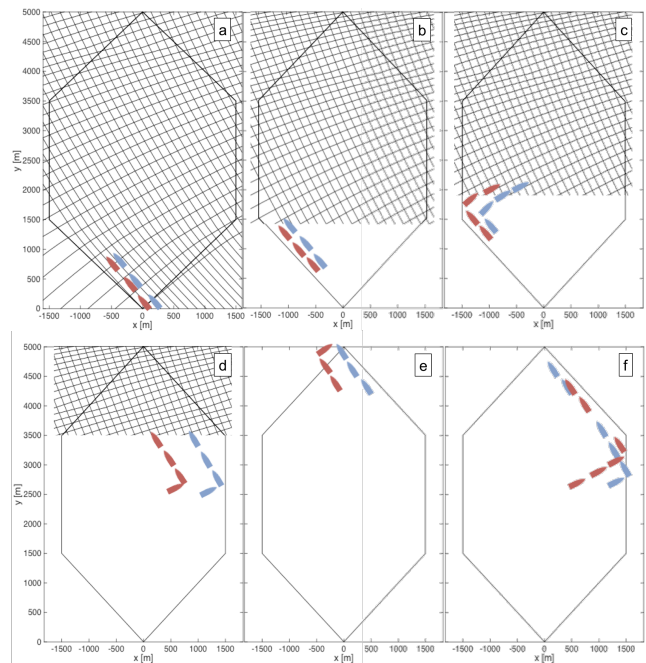


Figure 12: Simulation results for an upwind leg.

of the race area up to where, with only one tack, she can reach the right-hand-side layline. This would be the optimal choice for B as well in the absence of A. Unfortunately the more the wind shifts towards the right, the more B finds herself in her area of unfavourable aerodynamic influence (Figure 12(b)). B cannot tack until A tacks because the two boats are too close.

When eventually A tacks (Figure 12(c)), B is free to tack as well, but she chooses to wait in order to perform the tack outside of the area where she would still be slowed down because of the presence of A.

In Figure 12(d) both boats are initially sailing on a port tack, and A is leading. A reaches the layline and tacks to sail towards the mark. B adopts a risk-seeking behaviour, and instead of waiting to reach the layline as well she tacks hoping in a favourable wind shift.

Figure 12(e) shows the end of the race. Although B has

managed to avoid A to gain more advantage, she is still slightly behind.

Figure 12(f) shows an alternative realization for this example, where the strategy is computed without taking into account the presence of the opponent. In this case, boat B postpones the second tack until she reaches the racing area right boundary, but this then results again in finding herself in an area of negative influence.

CONCLUSIONS

This paper presents a novel methodology for computing an optimal strategy for a sailing match race. This methodology is based on dynamic programming in combination with a forecasting module based on artificial intelligence, which performs a very-short-term forecast of wind speed and direction, and with a risk model that allows a sailor to tune his risk attitude depending on his position with respect to the competitor.

The aim of the method is to compute a strategy that improves the probability of winning the race with respect to strategies aimed at minimising the expected time to complete the race. As an example, for an upwind leg of the 34th America's Cup, the completion time of a boat which uses the proposed forecast is only 10.7 s longer than a boat which has perfect knowledge of the wind. The risk model is based on coherent risk measures in order to investigate whether a change in risk attitude can improve the probabilities of winning a race. An optimistic attitude is associated to a losing skipper, and a pessimistic, conservative one to a winning skipper. The risk model is optimised using three different parameters relative to the distance between boats and the anticipated future wind changes. The results suggest that there is a threshold defining the moment when it is advisable to seek more risk and that not always the risk-seeking and risk-averse behaviours correspond to optimistic and pessimistic anticipations on the wind.

The proposed method is implemented in a computer program with negligible run time and thus capable to compute the optimum route in real-time during a race.

REFERENCES

Nicolas Aubin. Wind tunnel modelling of aerodynamic interference in a fleet race. Master's thesis, Ecole Centrale de Nantes and Universite de Nantes, 2013.

J Banks, A Webb, T Spenkuch, and SR Turnock. Measurement of dynamic forces experienced by an asymmetric yacht during a gybe, for use within sail simulation software. In *Proc. 8th Conf. Int. Sport. Eng. Assoc.*, pages 0–5, Vienna, Austria, 2010. doi: 10.1016/j.proeng.2010.04.024.

Dimitri P Bertsekas. *Dynamic Programming and Optimal Control*, volume I. Athena Scientific Belmont, MA, 2007. ISBN 1886529264. doi: 10.1057/jors.1996.103.

D S S Ferguson and Pantelis Elinas. A markov decision process model for strategic decision making in sailboat racing. *Adv. Artif. Intell.*, 6657 LNAI:110–121, 2011.

J A Gretzky and J K Marshall. The partnership for America's cup technology: An overview. In *Proc. 11th Chesap. Sail. Yacht Symp.*, Annapolis, MD, 1993.

J E Kerwin and J N Newman. A summary of the H. Irving Pratt ocean race handicapping project. In *Proc. 4th Chesap. Sail. Yacht Symp.*, Annapolis, MD, 1979.

John S Letcher, John K Marshall, James C Oliver III, and Nils Salvesen. STARS & STRIPES. *Sci. Am.*, 257(2): 34–40, 1987.

Y Masuyama. Tacking simulation of sailing yachts: numerical integration of equations of motion and application of neural network technique. In *Proc. 12th Chesap. Sail. Yacht Symp.*, pages 117–132, Annapolis, MD, 1995.

P Van Oossanen and PN Joubert. The development of the winged keel for twelve-metre yachts. *J. Fluid Mech.*, 173:55, 1986. ISSN 1469-7645. doi: 10.1017/S0022112086001076.

Andrew B Philpott, Shane G Henderson, and D Teirney. A simulation model for predicting yacht match race outcomes. *Operations Research*, 52(1):1–16, 2004.

Andy Philpott and Andrew Mason. Optimising yacht routes under uncertainty. In *Proc. 15th Chesap. Sail. Yacht Symp.*, Annapolis, MD, 2001.

M Scarponi, P Conti, R A Sheno, and S R Turnock. A combined ship science-behavioural science approach to create a winning yacht-sailor combination. In *Proc. the 18th Chesap. Sail. Yacht Symp.*, pages 143–156, Annapolis, MD, 2007a.

M Scarponi, P Conti, and RA Sheno. Including human performance in the dynamic model of a sailing yacht: a matlab-simulink based tool. In *Proc. Mod. Yacht Conf.*, pages 143–156, Southampton, UK, 2007b.

T Spenkuch, S Turnock, M Scarponi, and A Sheno. Lifting line method for modelling covering and blanketing effects for yacht fleet race simulation. In *Proc. 3rd High Perform. Yacht Des. Conf. 2008*, pages 111–120, Auckland, New Zealand, 2008.

Thomas Spenkuch, Stephen Turnock, Matteo Scarponi, and Ajit Sheno. Real time simulation of tacking yachts: How best to counter the advantage of an upwind yacht. In *Procedia Eng.*, volume 2, pages 3305–3310, Vienna, Austria, 2010. doi: 10.1016/j.proeng.2010.04.149.

Thomas Spenkuch, Stephen Turnock, Matteo Scarponi, and R Sheno. Modelling multiple yacht sailing interactions between upwind sailing yachts. *J. Mar. Sci. Technol.*, 16(2):115–128, 2011. ISSN 0948-4280. doi: 10.1007/s00773-010-0115-9.

F Tagliaferri, A B Philpott, I M Viola, and R G J Flay. On risk attitude and optimal yacht racing tactics. *Ocean Eng.*, 90:149–154, 2014. ISSN 0029-8018. doi: 10.1016/j.oceaneng.2014.07.020.

F Tagliaferri, I.M. Viola, and R.G.J. Flay. Wind direction forecasting with artificial neural networks and support vector machines. *Ocean Eng.*, 97:65–73, 2015. ISSN 00298018. doi: 10.1016/j.oceaneng.2014.12.026.