OPTIMAL YACHT ROUTING TACTICS

F. Tagliaferri, Yacht and Superyacht Research Group, School of Marine Science and Technology, Newcastle University, UK, f.tagliaferri@ncl.ac.uk (corresponding author)
A.B. Philpott, Yacht Research Unit, The University of Auckland, Auckland, NZ, a.philpott@auckland.ac.nz
I.M. Viola, Yacht and Superyacht Research Group, School of Marine Science and Technology, Newcastle University, UK, ignazio.viola@ncl.ac.uk
R.G.J. Flay, Yacht Research Unit, The University of Auckland, Auckland, NZ, r.flay@auckland.ac.nz

When the future wind direction is uncertain, the tactical decisions of a yacht skipper involve a stochastic routing problem. The objective of this problem is to maximise the probability of reaching the next mark ahead of all the other competitors. This paper describes a system that models this problem. The tidal current at any location is assumed to be predictable, while the wind forecast is based on current observations. Boat performance in different wind conditions is defined by the output of a velocity prediction program, and we assume a known speed loss for tacking and gybing. The resulting computer program can be used during a yacht race to choose the optimum course, or it can be used for design purposes to simulate yacht races between different design candidates. As an example of application, we compare strategies that minimise the average time to sail the leg, as opposed to those that maximise the probability of winning, and show how optimal routing strategies are different for leading and trailing boats.

NOMENCLATURE

Scalars

d	Distance between two competitors
$D(s)$	Delay in finishing under strategy s versus the perfect strategy
i_k	Values that a discrete-time stochastic process can assume at the k^{th} time step
$S(s)$	Time to finish under the strategy s
T	Time to finish under the perfect strategy
X_k	Discrete-time stochastic process, e.g. wind direction at the k^{th} time step
Y_k	k^{th} random variable uniform in (0,1)

Matrices

$\boldsymbol{M}^{k,l}$	Policy matrix at the cross-section k on the tack l, where l is starboard or port
$\boldsymbol{P}$	Transition matrix

Set of matrices

s	Strategy, i.e. set of policy matrices

Operators

$\mathbb{E}$(A)	Expected value of (A)
$\mathbb{P}$(A|B)	Probability density function of A conditioned on B

Abbreviations

BS	Boat speed
RMP	Race modelling program
SPP	Shortest path problem
VMG	Velocity made good

1. INTRODUCTION

Finding a minimum cost route on a set of points is a shortest path problem (SPP) [1]. Specifically the aim is to find a path between two vertices of a graph such that the sum of its constituent edges, often representing a cost, is minimised. When cost depends on random quantities it becomes a stochastic problem, and the objective is to minimise expected costs (where costs include time) [2].

Many problems fall into the category of SPP and involve routing for emergency response (both civil [3] and military [4]) and applications in logistics [5] and transport [6].

When minimising expected costs there is always a risk factor that must be taken into account. Is the best route the one that allows the average shortest time with a small probability for a disaster or the one that has a higher average time but without the risk for disasters, or even the one with an even higher average time but with a positive (even if small) probability of a particularly high gain?

Decisions taken by a sailor during a race can be seen as a stochastic SPP. The speed of a sailing boat depends on the wind speed and on the angle between boat heading and wind direction. It is usually expressed as a polar diagram like the one shown in Figure 1. The numbers around the semicircle represent different true wind angles, while the radial ones represent the boat speed. The red line correspond to the plot of boat speed corresponding for a particular true wind speed. While no direct course is possible straight into the wind, it is possible to sail upwind with an angle between wind direction and sailed course which is usually between 30° and 50°. Sailing closer to the wind direction (lower angle) makes the course shorter, but when sailing at higher angles a boat is faster. The compromise is given using the concept of velocity made good (VMG), which is the maximum velocity into the wind direction and is usually around 40°- 45° (as in this example).

In a polar diagram like the one in Figure 1, it is possible to find the maximum VMG for a given wind speed by finding the intersection between the polar corresponding to the wind speed and the line perpendicular to the upwind direction.

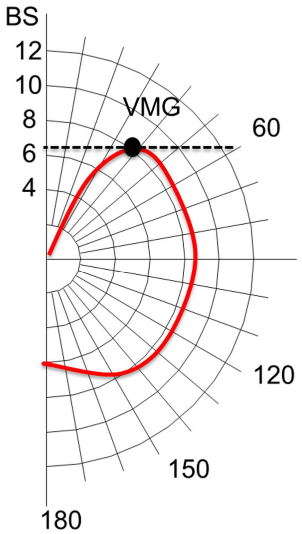

Figure 1. Example of a polar diagram (velocities in m/s and angles in deg).

For this reason the common route towards an upwind mark, or in general towards the direction from which the wind blows, is a zigzag route. Such a route requires changes of direction which are called tacks. When manoeuvring for a tack, a boat points for a few seconds directly into the wind, therefore causing a temporarily decrease in boat speed. If the wind is constant during the race and all over the racing area, trying to do the minimum number of tacks is the best choice. Figure 1(a) shows two possible routes, and the one on the left is the faster because it involves just one tack. However, this situation is very unlikely, and wind can change in many different ways. Figure 2(b) shows a situation in which the wind shifts constantly towards the left. The best choice in this case is to go to the left of the course, and then tack and point towards the mark, while beginning a race going to the right, after what can seem an initial advantage, results in being the wrong choice.

In real races the evolution of the wind can be much more complicated than this example, with temporary shifts or gusts that a sailor should take advantage of. While racing it can be difficult to know how the wind is behaving at another point, and to foresee how it will behave once that point is reached. Some sailors would prefer to be conservative and stay safely at the centre of the course, or in general close to the competitor, while others might prefer to take the risk and explore the corners hoping for a favourable wind shift.

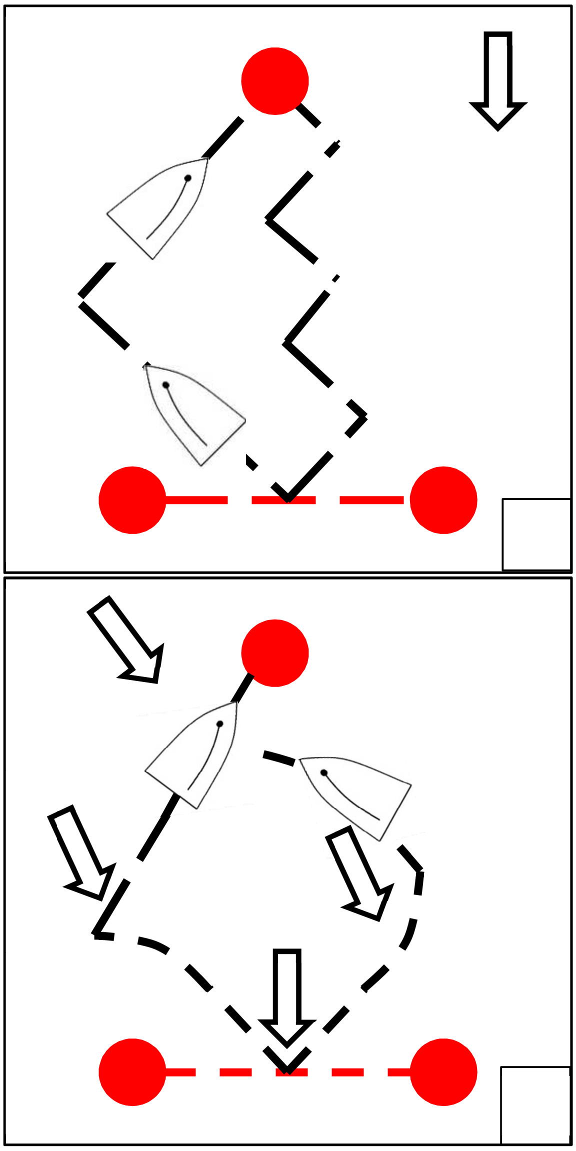

Figure 2. Example of upwind routes with constant wind (a) and with a consistent left wind shift (b)

One way to compute an exact solution to the problem of finding the optimum route between two points is an exhaustive search which is computationally impossible. In fact, even assuming that a boat travels always at its maximum VMG, there are 2^n possible routes, where n is the number of possible tacks. Considering that it is virtually possible to tack at every point of the race this is a continuum of possibilities, and for each one we have to compute the actual racing time.

Dynamic Programming is a popular way of overcoming this problem. This technique divides the problem into smaller sub-problems that are solved in stages. Moreover, instead of finding a solution for a single specific wind pattern, it allows computing a solution according to a certain probability distribution. We model the wind using a Markov process, where the state of the wind at a point depends on the state at the previous point. In addition to this, we want to include the attitude towards risk in order to consider more realistic behaviour of a sailor during a race.

The risk that a skipper is willing to take is usually influenced by his position with respect to the opponent. A common behaviour pattern is to be conservative, or risk averse, when in a leading position, while being risk seeking when losing. When looking for a solution for our SPP, we want to also take the risk factor into account.

We developed a race modelling program (RMP) for simulating races between two boats that can be used, for instance, to assess different designs, or as in our case, to compare different tactical decisions.

The first RMP was developed in 1987 for the America's Cup syndicate Stars and Stripes and is described in [8]. Since then, RMP have been used mainly in America's Cup applications, and mostly to compare different designs. In our case as we are interested in comparing tactical choices, we model two identical boats (i.e. they have the same polar diagram).

2. METHOD

2.1 MARKOV MODEL FOR THE WIND AND SOLUTION METHOD

A Markov process is a stochastic process used to describe the evolution of a dynamic system in which the state at the discrete time k depends on the state of the system at time $k-1$. More specifically, the Markov property states that for a discrete-time-stochastic process X_k, the probability distribution for the variable X_k, conditioned on all the previous values, is equal to the distribution for the variable X_k conditioned on the previous event:

$$\mathbb{P}(X_k = i_k | \{X_{k-1} = i_{k-1}, X_{k-2} = i_{k-1} \dots X_0 = i_0\}) = \mathbb{P}(X_k = i_k | X_{k-1} = i_{k-1})$$

for every $k > 0$, and for every i_k in the state space.

Therefore, the current scenario depends only on the previous one; this is the reason why the Markov property is also referred to as "loss of memory" for random processes.

For a system with a finite number of states the stochastic process is uniquely defined with an initial distribution for X_0 and a transition matrix $\boldsymbol{P}$. The matrix elements P_{ij} represent the probability that the system at time step k is in state j conditioned on the fact that it was in state i at the previous time step $k-1$:

$$P_{ij} = \mathbb{P}(X_k = j | X_{k-1} = i)$$

Interested readers can find more details on Markov processes in Norris [9].

For tactical purposes we are interested in changes in wind direction that significantly affect the racing time. We therefore define the state space to be $-45°, -40°, \dots, 0°, +5°, \dots, +45°$, where $0°$ represents the wind direction at which the upwind mark is set, and the other states represent shifts of $\pm 5°$ from that direction. In order to obtain a realistic transition matrix we considered a time series of wind measurements from a weather station installed on the Newcastle University research vessel, and then built the matrix $\boldsymbol{P}$ using a maximum likelihood estimator. As we use for the model a grid with 15m resolution in the upwind direction and the decision of tacking is taken every time the boat crosses the grid line in the y direction, the Markov model is built assuming a time step of three seconds. The wind direction signal was sampled every three seconds, and the corresponding wind directions were placed in bins of amplitude 5°. The number of jumps from bin i to bin j divided by the total number of jumps out of bin i corresponds to the value P_{ij} in the transition matrix. We consider an upwind leg of 6000m (corresponding to 3.24 nautical miles, which is a realistic length for the 2013 America's Cup course), divided by 400 lines perpendicular to the upwind direction. We refer to those lines as "cross sections". Each one of those lines is divided in a linear grid of 19 segments. It should be noted that the resolution of 5° for the wind shifts and 15m (6000m/400) for the racecourse can be increased by employing more powerful computational resources. Figure 3 shows a schematic diagram of the course with the axis orientation that is used throughout this paper.

The solution method is based on the algorithms described by Philpott and Mason [7] and Philpott. [10]. It is implemented in a highly modular code written in Matlab with some specific subroutines in C. This program computes the policy that gives the minimum expected time for the completion of the leg of the race. The output of the algorithm is a policy, expressed as a set of 19x19 matrices, $\boldsymbol{M}^{1,s}, \boldsymbol{M}^{2,s}, \dots, \boldsymbol{M}^{400,s}$, $\boldsymbol{M}^{1,p}$, $\dots$, $\boldsymbol{M}^{400,p}$ one for each of the cross sections on the course. The element $M_{ij}^{k,l}$ represents the optimum angle at which the yacht should sail when it reaches the k^{th} cross section, if it is on the i^{th} sub-segment of the cross section, and observing a wind in state j. The index l can assume the value p or s, corresponding respectively to a port or starboard tack. If a boat is on a port tack it means that its windward side is the left side, while it is on a starboard tack if the windward side is on the right one.

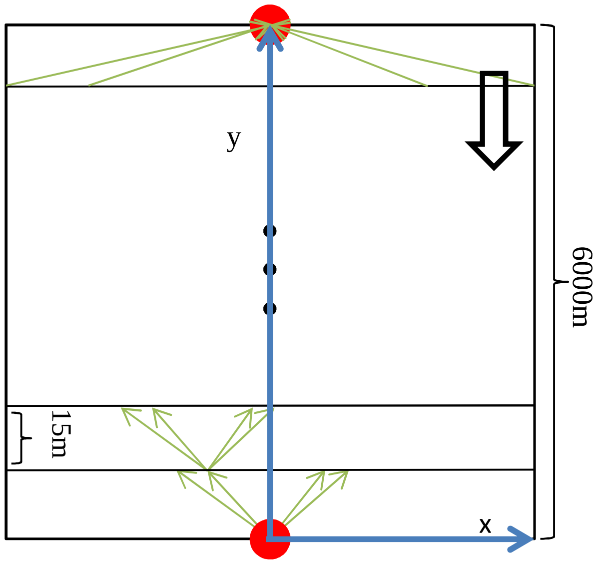

Figure 3. Schematic representation of the course. The y axis is oriented in the upwind direction, and the course is divided into 400 stripes by lines at 15m spacing.

The computation of the optimum angles begins from the top line and then proceeds iteratively with a backward procedure according to the description by Philpott and Mason [7].

2.2 RISK MODELLING

We modify the transition probabilities used to compute the solution in order to have conservative or risk seeking decisions. In a conservative case we want to model the behaviour of a boat skipper who is winning. She will try to behave safely, trying to stay ahead and to minimise her losses in bad wind outcomes. A probabilistic interpretation of this attitude is to assume that at the next step the wind will transition to a bad state with a higher probability than we have estimated from the data. In other words the skipper is pessimistic about the next transition. We implement this by adding a transformation in the solver, post multiplying the 19×19 transition matrix $\boldsymbol{P}$ by another 19×19 matrix which redistributes the probabilities. The resulting matrix has to be normalised in order to represent again a probability distribution. Figure 4 shows a graphical representation of an example of transformation that can be applied to a transition matrix in order to obtain a more decentred distribution. With a notation that is used throughout this paper, we use a grey scale to represent values in the interval $[0,1]$ where white represents 0 and black represents 1. The effect of this transformation on the transition matrix is to increase the volatility of the wind process.

2.3 RACE SIMULATIONS

A race simulator based on a simple SPP was developed in order to compare different policies. The *y*-axis is oriented in the upwind direction, positive upwind. The starting position of a single boat is the origin. The lines of equation $y = 15k,\ k = 1, \dots, 400$ determine the positions at which the decision is taken.

Wind patterns are generated according to the Markov model on states $i = 1,2, \dots 19$. Using a standard procedure for simulating Markov processes we set $X_0 = 1$. We then generate a series of random variables Y_k with uniform distribution in the interval $(0,1)$ using a random number generator.

Given $X_k = i$ we set $X_{k+1} = j$ if both the following conditions are satisfied

$$\begin{cases} \sum_{n=1}^{j-1} P_{in} < Y_{k+1} \\ \sum_{n=1}^{j} P_{in} \geq Y_{k+1} \end{cases}$$

The course of the sailing boat starts from the point (0,0), on a starboard tack. It follows a course corresponding to the angle $M^{1,s}_{10,X_0}$, until the second cross section is reached. The time needed to go from a position to the next is computed according to a polar diagram like the one shown in Figure 1.

2.4 MODEL VALIDATION

In order to assess the effectiveness of the model in finding an optimum solution, we use the algorithm to generate a policy by giving as input to the software the actual wind realisation. The expected values are then computed at each step by assigning a probability of one to the actual realisation. In this way we simulate the behaviour of the perfect tactician, who takes her decision knowing exactly how the wind is going to behave. In a real situation this is obviously not possible, but assuming that a very experienced sailor is able to fairly accurately predict what is going to happen in a race according to her experience, we want to show that our model still allows a good result against this ideal sailor.

3. SIMULATION RESULTS

Figure 5 shows a graphical representation of the transition matrix for the Markov model obtained with the maximum likelihood estimator as described in the previous section. It can be noticed that the diagonal is dominant, meaning that, in general, if the wind is in state i, the most probable state for the next step is to remain in state i. Moreover, when the wind has deviated from the mean, the event of a shift back towards the mean value is more likely than one in the same direction.

The wind for the simulations was generated as described in the previous section. When Markov chains are used, it is common practice to add a noise component to the generated output in order to avoid a step signal. However as in our case we are interested only in wind shifts of at least 5˚, we clustered the measured wind in steps of 5˚, and we found that the behaviour of the simulated wind signal, achieved with no additional noise component, was fairly similar to the original one, as can be seen in Figure 6, with close values of mean and variance on different sub-intervals. A wind history of 400 values was generated for each of the 4000 simulated races.

Figure 7 shows a histogram of the time needed by a yacht following the policy generated to minimise the expected time of arrival, according to the wind distribution previously modelled. The distribution is asymmetric, and this is due to the fact that even with a very favourable evolution of the wind there is a minimum time needed to complete the course. On the other hand, even with a policy which is effective in the majority of the cases, it is possible to be very "unlucky" and need a much higher time.

This policy was generated according to the wind distribution pictured in Figure 4. This policy was then compared with another one, generated according to a new transition matrix obtained from a transformation of the previous one. As mentioned in the previous paragraph, in order to model the attitude of a sailor who is not in a winning position, we use a transformation aimed at giving a higher volatility to the wind process, therefore giving a higher probability to unlikely future wind directions.

Simulations were carried out in order to verify the differences between a risk-neutral policy that minimises expected arrival time at the top mark, and a policy generated assuming a more volatile wind evolution. Results showed that following this second policy gives an overall worse performance with respect to the risk-neutral one. The risk neutral policy led to a win in 80% of the cases with an average difference of 92s. Those values were obtained by simulating races independently for each boat, but using the same 4000 wind patterns for all of them. These results confirmed the optimality of the policy previously computed.

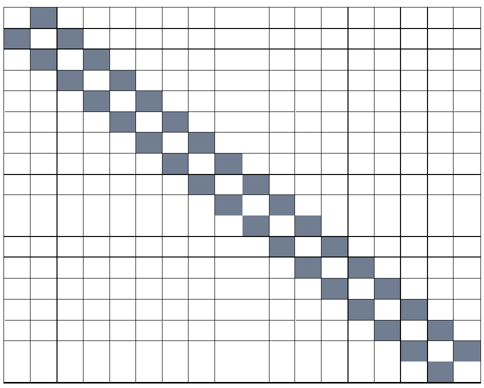

Figure 4. Example of a transformation used to modify the transition probability of the Markov model.

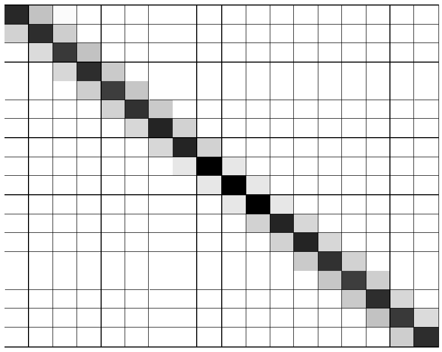

Figure 5. Representation of the transition matrix ***P*** obtained for the wind model.

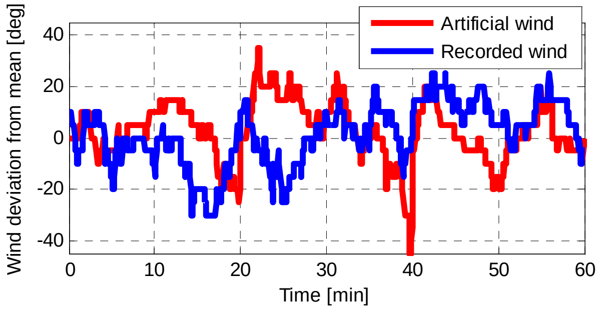

Figure 6. Sixty-minute example of artificially generated wind and sixty-minute example of recorded wind.

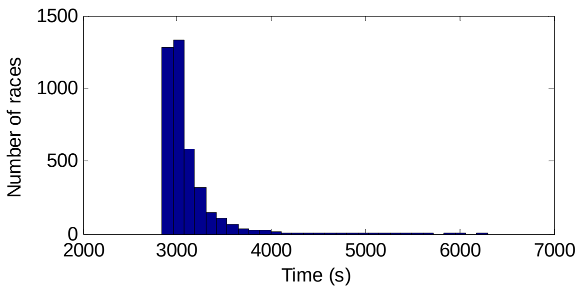

Figure 7. Distribution of time of arrival needed by a boat following the optimum policy.

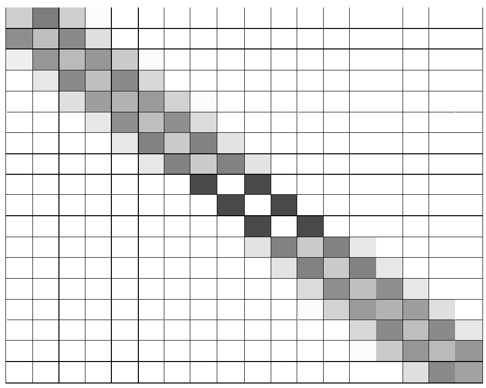

Figure 8. Transition matrix ***P*** with increased volatility

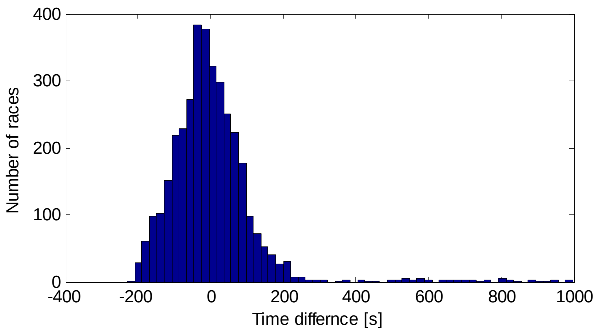

Figure 9 Distribution of time of arrival differences between boats using the two different policies

However, combining together the two strategies can lead to a significant improvement in the chances of winning. We simulated races between two boats that are denoted as boat A and boat B. Both boats start the race at the same time, on two different (random) points along the starting line. Boat A experiences the same wind as in the previously simulated races and always follows the optimum policy. Boat B experiences the same wind as A if their distance is less than $d_{min} = 10$m, an independent wind if their distance is greater than $d_{max} = 100$m, and a linear combination of the d_{min} and d_{max} wind if their distance is between d_{min} and d_{max}. At every spatial step, if B is more than 15s behind A, she uses the risk-seeking policy, while she uses the optimum risk-neutral policy otherwise. Results of those simulated races are shown in Figure 9.

The x-axis shows the arrival time of boat B minus the arrival time of boat A at the top mark. The average time difference is positive (actually 39s in this plot). This means that B arrives 39s later on average than A, as one would expect, since A is using the optimum policy to minimise the average time. However about 57% of the race outcomes are to the left of the vertical axis, meaning that B wins 57% of the time (always by a small margin). Of course sometimes B is hopelessly outclassed, losing by 2000 seconds (just around 0.01% of the times, and those are extremely unfavourable events) but this is because B takes high risks when behind. If we consider $p = 0.5$ win probability as a null hypothesis, then the probability of winning more than 57% of 4000 races by chance is the probability that a binomial random variable with mean $4000p$ and variance $4000p(1-p)$ exceeds 2280, which is about 10^{-14}.

The standard error of the value 0.57 can be estimated using the central limit theorem to be approximately 0.008. So we can be 97.5% confident that the hybrid policy will win at least 55.4% of the races (i.e. 2 standard errors less than 0.57).

In order to quantify the tactical improvement on the policy we compare the results obtained by boat A and boat B with a third boat C that has perfect knowledge of the future behaviour of the wind. In this case we simulated 1000 races. Obviously the boat with perfect knowledge of the wind scenario always wins and the differences in arrival time are always positive. The sample average difference in time of arrival is 133s for boat A while for boat B the sample average difference is 114s. The difference is not significant because of high variance and low sample size. Indeed we show in the appendix that the expected time difference for boat A relative to C must be lower than the expected time difference for boat B relative to C.

4. CONCLUSIONS AND FUTURE WORK

In this paper we have presented a method for approximating a solution of a stochastic shortest path problem with applications to yacht racing. We showed that with an adequate subdivision of the problem it is possible to find a solution that minimises the expected time needed to reach an upwind mark during a race.

Moreover, we introduce for the first time a model of the risk attitude of the sailor. We showed that if a skipper of a trailing boat has a risk-seeking attitude it enhances the chance to win the race. An important result of the simulations run to simulate races was that aiming at minimising the expected time to finish is not always the best approach: being on average slower might allow a bigger probability of winning against an opponent following a fixed policy.

The results presented in this paper underline that, when trying to optimise a policy in order to win a competition, looking at average values is rarely the best approach, and accounting for differing risk attitudes might give policies that perform significantly better. Further work is being carried out in order to validate the model with data registered during America's Cup races, and we are developing methodologies for learning risk parameters that yield maximum win probabilities.

ACKNOWLEDGEMENTS

This research has been performed within the SAILING FLUIDS project, which is funded by the European Commission under the 7th Framework Programme through the Marie Curie Actions, People, International Research Staff Exchange Scheme.

REFERENCES

1. Cherkassky, B.V., A.V. Goldberg, and T. Radzik, *Shortest paths algorithms: Theory and experimental evaluation.* Mathematical Programming, Series B, 1996. **73**(2): p. 129-174.
2. Bertsekas, D.P. and J.N. Tsitsikilis, *An Analysis of Stochastic Shortest Path Problems.*

Mathematics of Operations Research, 1991. **16**(3): p. 580-595.
3. Yamada, T., *A network flow approach to a city emergency evacuation planning.* International Journal of Systems Science, 1996. **27**(10): p. 931-936.
4. Resch, C., et al. *Path Planning for Mine Countermeasures*. in *SPIE - The International Society for Optical Engineering*. 2003. Orlando, FL.
5. Fleischmann, B., S. Gnutzmann, and E. Sandvoß, *Dynamic vehicle routing based on online traffic information.* Transportation Science, 2004. **38**(4): p. 420-433.
6. Li, S. *Study on routing optimization problem of the logistics center*. in *World Automation Congress*. 2012. Puerto Vallarta, Mexico.
7. Philpott, A. and A. Mason, *Optimising Yacht Routes Under Uncertainty*, in *Proc. of the 15th Chesapeake Sailing Yacht Symposium*2001.
8. Letcher Jr, J.S., et al., *STARS & STRIPES.* Scientific American, 1987. **257**(2): p. 34-40.
9. Norris, J.R., *Markov chains.* 1st pbk. ed. Cambridge series on statistical and probabilistic mathematics1998, New York: Cambridge University Press. xvi, 237.
10. Philpott, A. Stochastic optimization in yacht racing, in *Applications of Stochastic Programming*, W. Ziemba and S. Wallace (ed.), SIAM, 2005.

APPENDIX

Proposition: Minimising the expected arrival time over all strategies will give a policy that is slower than a perfect skipper by the least amount on average.

Proof: Suppose a perfect skipper sails races in wind that she predicts perfectly. Each race is a random sample of wind and so her time to finish is an independent identically distributed random variable T.

Suppose she now sails a strategy s that is not clairvoyant in each of these same wind conditions. The time to finish under this strategy is an independent identically distributed random variable $S(s)$.

Now the delay in finishing under strategy s versus the perfect strategy is also an independent identically distributed random variable $D(s) = S(s) - T$. The expected delay from sailing s is then

$$\mathbb{E}[D(s)] = \mathbb{E}[S(s)] - \mathbb{E}[T].$$

To minimise this we should minimise $\mathbb{E}[S(s)]$ as $\mathbb{E}[T]$ is a constant. So the strategy that minimises expected delay after a clairvoyant skipper is the one that minimises expected arrival time.

AUTHORS BIOGRAPHY

F Tagliaferri is a PhD student at Newcastle University and member of the Yacht and Superyacht Research Group. She holds a Masters degree with Honours in mathematics and her PhD project aims at developing navigation software for yacht races under uncertain weather conditions.

AB Philpott, PhD, is a Professor in the Department of Engineering Science and Director of the Electric Power Optimization Centre at the University of Auckland, New Zealand. His research interests encompass optimisation under uncertainty and game theory with particular application to electricity markets. He has also applied these technologies to yacht routing and race modeling in several America's Cup campaigns.

IM Viola, PhD, is Lecturer in Naval Architecture at the School of Marine Science and Technology of Newcastle University, UK. He has a background in applied fluid dynamics and a specialist expertise in yacht engineering. His previous experience includes a Post Doctoral Fellowship at the Yacht Research Unit (The University of Auckland), which is the Scientific Advisor of the America's Cup team Emirates Team New Zealand, and a PhD at the Politecnico di Milano, sponsored by the America's Cup team Luna Rossa, on experimental and numerical modelling of the aerodynamics of sailing yachts.

RGJ Flay, PhD, is Professor of Mechanical Engineering and Director of the Yacht Research Unit in the Department of Mechanical Engineering at the University of Auckland. He has had a longstanding research interest in the wind and sailing. His PhD degree was awarded for a study of wind structure based on full scale wind data. His Postdoctoral research as a National Research Council Visiting Fellow in Canada was focused on carrying out wind tunnel studies over topographic models to compare with full-scale measurements, and for wind energy prospecting. He then spent four years as an Aerodynamic Design Engineer in a Consulting Engineering company in Toronto where he worked on the design of several wind tunnels and environmental test facilities. Since 1984 he has worked at the University of Auckland, and in 1994 designed the World's first Twisted Flow Wind Tunnel.