Optimal Weather Routing Using Ensemble Weather Forecasts

Asher Treby
Department of Engineering Science
University of Auckland
New Zealand

Abstract

In the United States and the United Kingdom it is commonplace now for forecasters to issue ensemble forecasts, which are a collection of individual forecasts generated from slightly varying initial conditions. This paper looked at using ensemble forecasts to optimise the choice of sailing directions in a yacht race.

Each individual forecast of the ensemble (called a member) will always have an associated realisation probability. This is the likelihood that the member forecast correctly predicts the weather. As time progresses it is possible to compare what the weather has actually done with what each member predicted and adjust the realisation probabilities accordingly.

When selecting a sailing direction there is a trade-off between changing course to catch strong winds and staying on course for the finish line but in weaker winds. The factors affecting this decision are where in the ocean the yacht is, what time it is and what the weather is going to do. This paper took these three factors to define a state space for a stochastic dynamic program. The weathers observed behaviour is captured in the state space using sets of realisation probabilities on the ensemble members.

The dynamic program developed will fit easily with different yacht movement models but is limited to forward sailing directions. Pilot problems were solved in reasonable time giving logical solutions demonstrating the trade off decisions made, showing this new concept to be plausible.

---

1 Introduction

Yachting happens world wide in many different capacities and large amounts of time and money are put into research for yacht racing. When racing, the aim is to get your yacht to the finish faster than all the other yachts in the race. Design and construction has a major impact on a yacht's performance in any conditions. However, where and how the yacht is sailed is paramount to winning a race.

Because a yacht's movement is powered by the wind, navigating the yacht into the best winds is one of the most important and difficult jobs of a skipper. When deciding which sailing route to use, there are two conflicting ideas to take into account. Finding the best winds for fast sailing speeds and keeping the yacht on the correct course.

So even though there may be stronger winds a little way off it could be better in terms of the race to stay in the current winds because sailing to the stronger winds might set the yacht too far off-course and result in a time loss instead of a gain. These route decisions are often done using just gut-feelings based on past experience. Developing more accurate and rigorous approaches to this navigation should prove very beneficial to yachting.

Any approach to finding an optimal sailing route will make use of at least one weather forecasting technique. This paper aimed to investigate methods for finding optimal sailing routes using ensemble weather forecasts. Ensemble weather forecasts are actually a collection of different weather forecasts computed using weather models starting from perturbed initial conditions. These ensemble forecasts have become commonplace in North America and so this investigation is in anticipation of the growing popularity and issuing of these forecasts.

Figure 1 shows the basic physical situation of the problem where the aim of yacht racing is to cross the finish line, after leaving the start point, in the shortest time possible. The solution approach taken for this paper is stochastic dynamic programming.

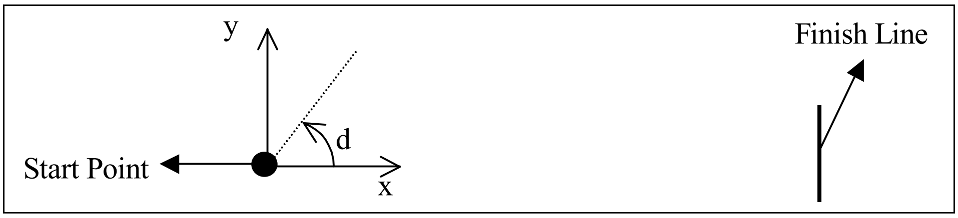

Figure 1 Reference geometry for problem.

2 Ensemble Forecasting

The material in this section is drawn mainly from reference [1]. Weather forecasts are computed by running a mathematical model using some initial weather conditions. However there is no method currently available to sample weather perfectly. This means any weather forecast generated will be computed using initial conditions containing some error. When the mathematical model is run it is likely that these errors will be magnified, as the forecasting time stretches further into the future.

Ensemble forecasting addresses this problem by computing a collection of forecasts all from slightly different initial conditions. It is hoped that by using a distribution of initial conditions the actual weather situation is more likely to be represented. The collection of forecasts is known as the ensemble and each forecast in the ensemble is known as a member. The method of perturbing the initial conditions is not arbitrary and reference [1] gives outlines on two different methods.

Member Realisation Probabilities

Within the ensemble it is possible that some members are more likely to be accurate than others. For this reason each member will have a realisation probability associated with it. This probability describes how likely it is that the weather happens as the member forecast predicts. This realisation probability will be denoted P(M) where P(2) is the probability that member 2 is realised.

To demonstrate how these probabilities are important consider a simplified situation where there are three members of the ensemble and each member is a series of forecasted wind strengths. For now we will simplify the situation so that it is only possible to have low, medium and high winds. These winds will be called observations and each will have an associated realisation probability P(O) where P(Low), P(Med), P(High) are the probabilities that the winds will be low, medium or high respectively.

Table 1 gives the member forecasts that would be generated at time zero and the actual wind strengths that will happen.

	Time 0	**Time T**	**Time 2T**	**Time 3T**
Member A	Medium	Low	Low	Medium
Member B	Medium	High	Medium	High
Member C	Medium	Medium	Medium	Low
Actual Weather	Medium	High	Medium	Low

Table 1 Example ensemble forecast and actual weather comparison

Usually at the start of the forecasted period the members will be similar and so will have equal realisation probabilities so for this example P(A) = 1/3, P(B) = 1/3 and P(C) = 1/3. However at time T we would see that the wind strength has in fact increased and so we would adjust the realisation probabilities to reflect this.

Because member B predicted an increase in wind strength its realisation probability could be increased. What it should be increased to is dependant on another set of probabilities that should be defined at the start of the forecasted period. These probabilities are P(O|M) which are the conditional probabilities that observation O happens given that we know member M is occurring. Table 2 shows the conditional probabilities that will be used here.

	Time T			Time 2T		
	M=A	**M=B**	**M=C**	**M=A**	**M=B**	**M=C**
P(Low|M)	0.850	0.050	0.075	0.850	0.075	0.075
P(Med|M)	0.100	0.100	0.850	0.100	0.850	0.850
P(High|M)	0.050	0.850	0.075	0.050	0.075	0.075

Table 2 Example conditional probabilities.

When adjusting the realisation probabilities what we really want to know is what is the probability that member m is happening given that high winds are being observed at time T. This will be denoted P(M|High) and is calculated using Bayes' rule (1).

$$P(M_i \mid O) = \frac{P(O \mid M_i)P(M_i)}{\sum_{j=1}^{N} P(O \mid M_j)P(M_j)} \qquad (1)$$

So the updated P(M=B) will be:

$$P(B \mid High) = \frac{P(High \mid B)P(B)}{P(High \mid A)P(A) + P(High \mid B)P(B) + P(High \mid C)P(C)}$$

$$= \frac{0.85 \times 1/3}{0.05 \times 1/3 + 0.85 \times 1/3 + 0.075 \times 1/3}$$

$$= 0.872$$

Similarly P(M=A) = 0.051 and P(M=C) = 0.077 at time T.

Current Availability and Uses of Ensemble Forecasts

In North America and the United Kingdom ensemble forecasts are readily available and there is a variety of companies that provide ensemble-forecasting services at a cost. The United States Navy regularly issue a variety of free marine ensemble forecasts, which provide information such as gale warnings for a range of different areas. There are currently no readily available ensemble forecasts issued in New Zealand.

3 Optimal Yacht Routes for Deterministic Weather

To make it easier to understand the developed model the dynamic program will be first explained within a deterministic framework and then extended to a stochastic framework. The best direction to sail at any point in time will depend on where in the ocean the yacht is and what the weather will be doing everywhere in the surrounding area, up until the yacht arrives at its destination. If the skipper had access to perfect weather forecasts this problem would be considerably easier to solve and the solution would be deterministic.

An optimal weather route is going to consist of a series of directions in which to sail. The direction to sail, from whatever point the yacht is at, to finish the race in the shortest time, is the optimal direction. If we know exactly what the weather is going to do over the course of the race it is possible to use a dynamic program to find a yachting route before starting the race.

Dynamic Program State Space

In order to use dynamic programming the state space must first be defined. For the deterministic problem the states are all the possible combinations of position on the ocean and time that the yacht could experience. Unfortunately both time and space are continuous parameters and so there will be infinitely many states to calculate for. To decrease the state space, time and ocean positions need to be discretised.

For time, solutions are only found for a set of discrete times that vary by a set period length, T. So for the dynamic program it is assumed the yacht travels in one direction for one period of time and then a new sailing direction is taken.

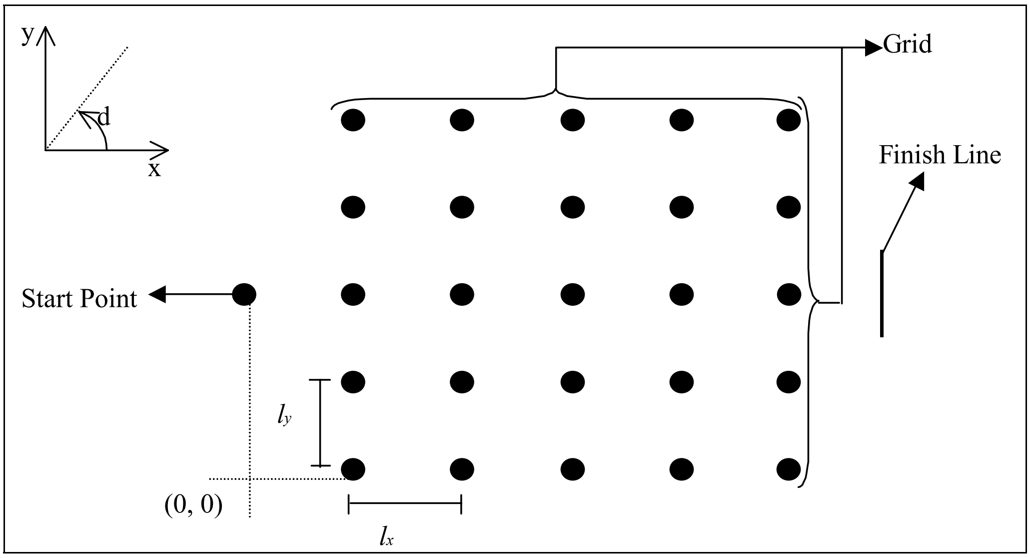

Figure 2 Problem layout using a 5x5 grid example.

For position, a grid of evenly spaced reference nodes is placed between the start point and the finish line. Figure 2 shows the addition of a 5x5 grid to the reference geometry introduced figure 1. It should be noted that the centre line of this grid is the line that runs between the start point and the finish line while being perpendicular to the finish line. Also the origin of the vertical axis is set to be in line with the bottom row of the grid and the origin of the horizontal axis is set to be in line with the start point. This grid is important for decreasing the size of the state space by placing a physical boundary on the area in which the yacht can go and by defining nodes from which the optimal directions need to be found.

While it may be possible for a yacht to be at any physical location in the ocean it has been assumed that for a yacht race the yacht's horizontal location x will stay between the start point ($x = 0$) and the finish line ($x = l_x(\beta+1)$], where β = the number of columns in the grid). The vertical location, y has also been restricted, as it seems logical that if the yacht strays too far from the finish line it will take too long to complete the race no matter how strong the winds are. Where these lines are placed is dependent on the specific race.

The nodes on the grid define a set of discrete locations and thus discrete states the yacht can be in. The yacht of course can still be at other locations within the grid boundary. Bilinear interpolation is used to cope with locations between nodes and is explained later.

Solving the Dynamic Program

The optimal direction from any state (combination of time and location) is the direction that minimises the finish time of the race i.e. the direction that if sailed in would mean the yacht would get to the finish at the earliest time possible.

To solve the dynamic program find the optimal direction for every possible state in the state space, starting from the last column and working backwards. This means start at one node in the last column of the grid and for the first period the yacht could be at that node, calculate the time the yacht would finish the race for all forward sailing directions. The earliest of these finish times is stored as the optimal finish time for that period at that node. The direction that would produce this finish time should also be stored and is the optimal direction for that state. This should be repeated for all possible periods. Then this should be repeated for every node in that column. In turn repeat this for each column in the grid working from the last column to the first.

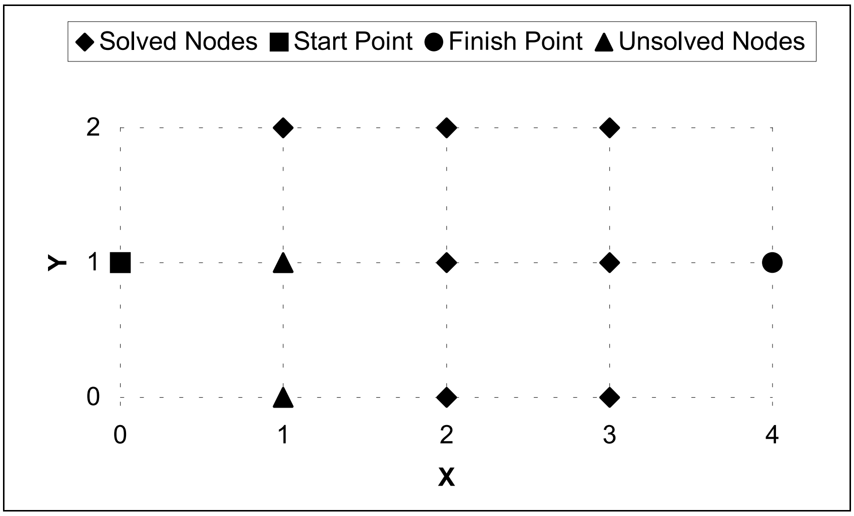

Figure 3 Dynamic programming example layout.

X	Y	Period	Optimal Finish Time	X	Y	Period	Optimal Finish Time
3	2	0	0.3	2	2	0	1.3
		1	1.4			1	2.4
		2	2.3			2	3.2
		3	3.6	2	1	0	1.1
3	1	0	0.1			1	2.3
		1	1.3			2	3.2
		2	2.6	2	0	0	1.7
		3	3.2			1	2.4
3	0	0	0.6			2	3.1
		1	1.8	1	2	0	2.5
		2	2.4			1	3.4
		3	3.9				

Table 3 Optimal Finish times for solved nodes for example deterministic problem

Calculating what the finish time will be for a given direction from a specific state is easiest to understand by working through an example. Here we are going to use the physical layout shown in figure 3. We will assume that the diamond nodes have already been solved, where a node is solved if all the optimal finish times for states including that node have been found. The optimal finish times for the states involving solved nodes that will be used for this example are listed in table 3. We will assume a period of 1 for simplicity.

The next node we will solve will be x=1, y=1 and we will find the optimal finish time for period 0. This means we want to find the optimal direction and thus finish time if the boat was at [1, 1, 0] (node x=1, y=1 at t=0). We shall assume for now that the only three directions the yacht can travel in are 45° above the horizontal, along the horizontal and 45° below the horizontal. To stay within the discrete time states it is assumed for the model that the yacht will travel in the optimal direction for one periods length and then change direction to the optimal direction for the state that it is then in. Because we know what the weather is going to do we can calculate where the yacht would end up after one periods sailing.

The focus of this paper was on the stochastic nature of a problem involving ensemble forecasts, not on the modelling of yacht movement. A simple model has been used for testing the stochastic dynamic program, however it would be very easy to slot a more sophisticated model such as one based on a velocity prediction program (VPP) into the framework developed. Reference [2], would provide a detailed explanation if of VPP if desired.

So say the yacht would arrive at [2, 2, 1] after travelling 45° above the horizontal from [1, 1, 0] for one period. From table 3 we know that if we are at [2, 2, 1] that we will finish the race at t=2.4. This means that if we travel from [1, 1, 0] at an angle of 45° above the horizontal we will finish the race at t=2.4. However if travel from [1, 1, 0] along the horizontal for one period and we arrive at [2, 1, 1] table 3 shows we will finish the race at t=2.3. This is a better finish time and so we would no longer consider sailing 45° above the horizontal from [1, 1, 0].

If we assume sailing from [1, 1, 0] at 45° below the horizontal for one period takes us to [2, 0, 1] then sailing in this direction from [1, 1, 0] will mean we finish the race at t=2.4. This is worse than the currently best finish time and so we would not use this direction. At this point we realise we have investigated all the directions we can sail from [1, 1, 0] and so the currently best finish time (t=2.3) and corresponding direction (d=0°) will become the optimal finish time and optimal direction for [1, 1, 0].

We would then move on to investigate [1, 1, 1]. So far we have assumed that travelling for one period in a given direction will take us exactly to a solved node. This is of course is not very realistic and as mentioned earlier, bilinear interpolation is used to find the optimal finish time for a location between nodes. This bilinear interpolation is shown in figure 4.

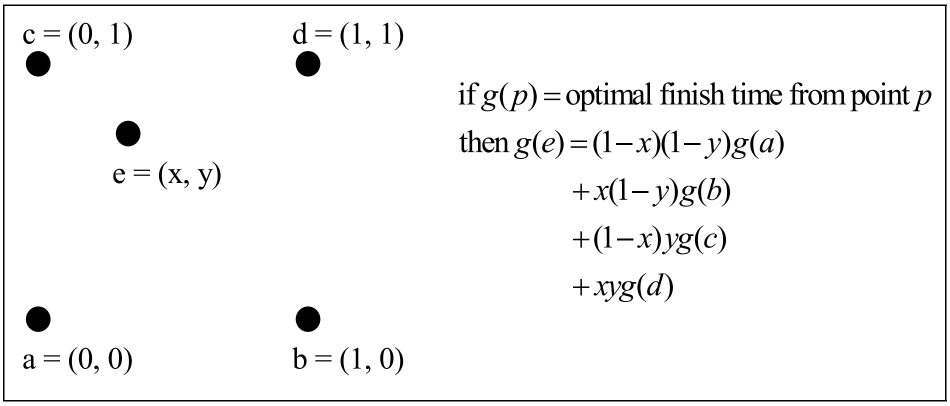

Figure 4 Bilinear interpolation

To demonstrate how this works we will say the yacht can travel from [1, 1, 1] at 30° above the horizontal, along the horizontal and at 45° below the horizontal. If travelling at 45° below the horizontal for one period brings us to state [2, 0, 2] then by looking at table 3 as before we will see this direction will result in a finish time of 3.1. If however travelling for one period at 30° above the horizontal from [1, 1, 1] would bring the yacht to [2.25, 1.72, 2], then the finish time will be:

$$\begin{aligned} g([2.25,1.72,2]) &= (1-0.25)(1-0.72)g([2,1,2]) \\ &+0.25(1-0.72)g([3,1,2]) \\ &+(1-0.25)0.72g([2,2,2]) \\ &+0.25\times 0.72g([3,2,2]) \\ &= 0.21\times 3.2+0.07\times 2.6+0.54\times 3.2+0.18\times 2.3 \\ &= 2.996 \end{aligned}$$

If the yacht was to travel from [1, 1, 1] along the horizontal for one period and arrive at [2.75, 1, 2] then using a similar calculation as above the finish time would be t=2.75. Travelling along the horizontal would clearly give the best finish time and so the optimal finish time for [1, 1, 1] would be 2.75 and the optimal direction would be 0^{o}.

4 Optimal Yacht Routes for Stochastic Weather

Section 3 developed a dynamic programming solution for finding the best directions to sail, from any time and place. This was based on perfect weather forecasts. So it was possible to calculate what state the boat would arrive at after sailing in one direction for one period of time. However it is not possible to have perfect forecasts and the purpose of this project is to look at the stochastic nature of ensemble forecasts and their affect on finding optimal yachting routes.

Because we can no longer say for certain what is going to happen with the weather we cannot calculate the finish time from a given state. So instead we have to calculate the expected finish time. This change also extends to the definition of an optimal direction. The optimal direction will now be the direction most likely to minimise the finish time or alternatively the direction that minimises the expected finish time.

Realisation Probability Sets

In a stochastic framework the time and space dimensions will be the same as those in the deterministic framework. However there is the added dimension of sets of realisation probabilities for the ensemble members. Like time and space the realisation probabilities are continuous so a method of choosing lists of discrete sets of realisation probabilities was required. Fortunately using a Bayesian framework makes this relatively simple.

To use Baye's rule (1) the set of all possible weather patterns needs to be broken down into a set of discrete possible weather observations. This was done in the example in section 2 where there were only three different possible types of wind, high, medium and low. How these discrete weather patterns are selected is dependant on the specific race.

At the start of the race when the ensemble is first generated there is only the one set of realisation probabilities, in which the probabilities are even. For any realisation probability set a new set of realisation probabilities can be reached for each discrete weather observation using Baye's rule (1). So if there was three discrete weather observations there would be one set of realisation probabilities at t=0, three sets at t=T, nine sets at t=2T etc. This is used to decrease the number of yacht states that need to be solved for, by finding all the possible realisation probability sets for each period before running the dynamic program.

Extending to a Stochastic Dynamic Program

The stochastic dynamic program would be solved in much the same way as the deterministic dynamic program. However the two main differences come in the addition of a third dimension to the state space i.e. the sets of realisation probabilities and in the use of expected finish times instead of known finish times.

The addition of the possible sets of realisation probabilities to the state space mean that each state that would be solved for in a deterministic framework needs to be further spilt into possible states of realisation probabilities for the stochastic problem. This means that each state [x, y, t, set of realisation probabilities] must be solved for. In terms of the computation, where one optimal direction would be found for each period at a node, now one optimal direction must be found for each possible set of realisation probabilities within that period at a node.

The most complicated extension to the deterministic framework is calculating the expected finish time for each state. To explain this we will look at the problem partially represented in figure 5 using a three-member ensemble with the simple discrete weather observations of low and high winds.

Now the important aspect of this model is that it takes into account what weather has actually been observed. The observed weather is captured in the state space by sets of realisation probabilities. So say we want to find the expected finish time from the heart node in direction d at time t under realisation probability set (1/3, 1/3, 1/3) i.e. [x_{heart}, y_{heart}, t, (1/3, 1/3, 1/3)] and we have already solved for the triangle and circle nodes. We will assume the circle, diamond and square points are the locations that would be reached if travelling from the heart node at time t in direction d for one period if ensemble members 1, 2 and 3 respectively are occurring.

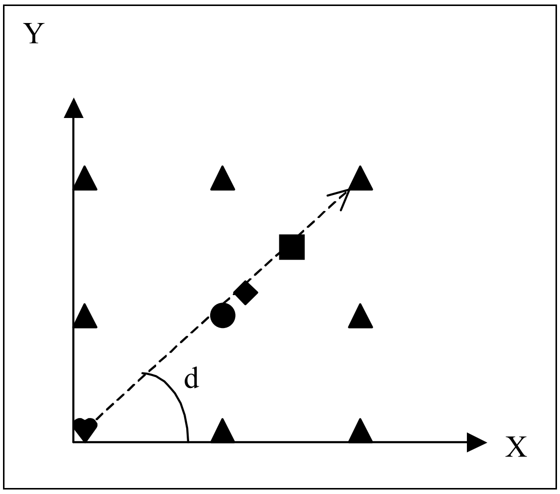

Figure 5 Finding the expected finish time for a three member ensemble.

No matter what weather is observed the destinations under the members do not change. However the realisation probabilities will be different. So for this case we will say that there is a 70% chance that low winds will be observed and 30% chance that high winds will be observed. If high winds are observed, using Bayes' rule will return realisation probabilities of say (0.1, 0.6, 0.3) and for a low winds observation say (0.7, 0.2, 0.1).

The bilinear interpolation is done the same way as for the deterministic framework just using expected finish times instead of known finish times. Each observation would change the set of realisation probabilities differently under the Bayesian framework and so the states that would be looked up when taking the expected finish points for bilinear interpolation would be the four locations surrounding the point reached under the observation, the time t+T and the realisation probability set that would be reached if that observation was realised.

So the expected finish time for [x_{heart}, y_{heart}, t, (1/3, 1/3, 1/3)] =

1/3{ 0.3(expected finish time for [x_{circle}, y_{circle}, t+T, (0.1, 0.6, 0.3)])
+ 0.7(expected finish time for [x_{circle}, y_{circle}, t+T, (0.7, 0.2, 0.1)])}

+ 1/3{ 0.3(expected finish time for [$x_{diamond}$, $y_{diamond}$, t+T, (0.1, 0.6, 0.3)])
+ 0.7(expected finish time for [$x_{diamond}$, $y_{diamond}$, t+T, (0.7, 0.2, 0.1)])}

+ 1/3{0.3(expected finish time for [x_{square}, y_{square}, t+T, (0.1, 0.6, 0.3)])
+ 0.7(expected finish time for [x_{square}, y_{square}, t+T, (0.1, 0.6, 0.3)])}

5 Results

In order to test the approach an example ensemble using two members was examined. A 5x5 grid was used with vertical and horizontal spacing of 1.5 and the finish line lay between the two points (9, 2.25) and (9, 3.75). The following five discrete weather observations were used:

1. Strong winds everywhere.
2. Medium winds everywhere.
3. Weak winds everywhere.
4. Strong winds around the Southern rows and low winds elsewhere.
5. Strong winds around the Northern rows and low winds elsewhere.

The first ensemble member had strong winds around the Southern rows and low winds elsewhere and the second ensemble member was the reflection of the first with strong winds around the Northern rows and low winds elsewhere.

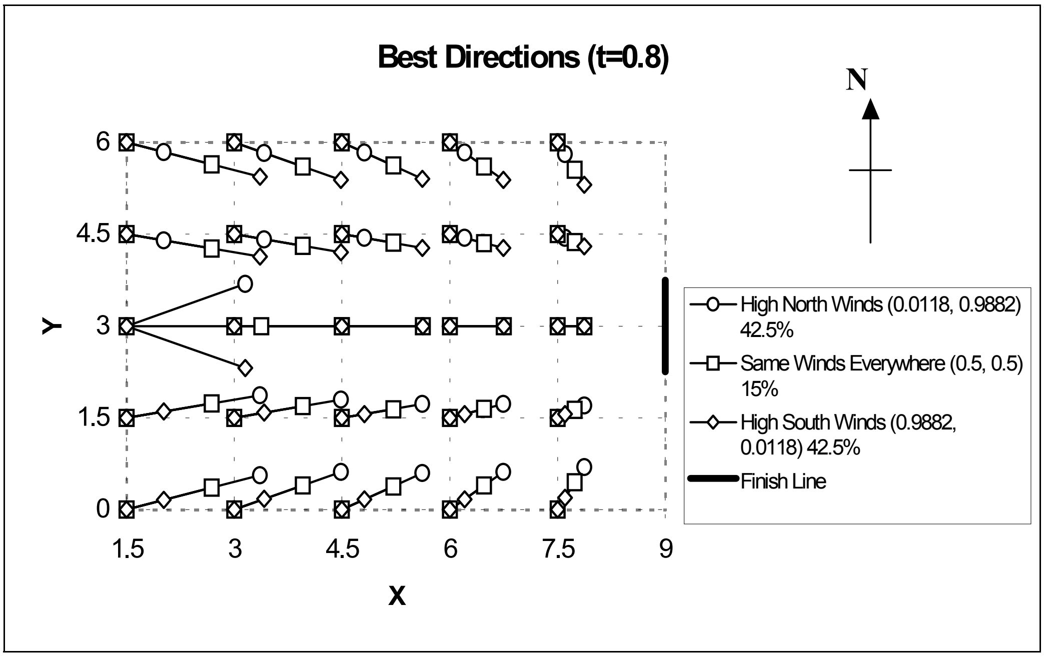

Figure 6 Optimal directions at the beginning of the second period for a two-member ensemble with one member of strong North winds and one member with strong South winds.

So the results of the model could be visualised the data was plotted as a series of lines from the grid nodes showing the optimal directions. The length of each line is proportional to the expected time to finish from that node. Figure 6 shows the results for this problem at the beginning of the second period The pairs of bracketed numbers to the right of the wind patterns in the key are the realisation probability sets that would result from the observation and the percentage is the probability of that observation occurring.

To allow results simple enough to visualise in this paper the ensemble members are completely different from the beginning of the forecasted period. This is not realistic as members should generally start around the same values and slowly diverge. However this does show the important features of results from the stochastic dynamic program.

The first thing to note is that there were five discrete observations applied to this problem but there is only three in the results. This is because of the conditional probabilities used. Basically $P(O|1) = P(O|2)$ for $O=1,2,3$ so if observation 1, 2 or 3 was actually realised then the model would say both members are equally likely to occur i.e. $P(1) = P(2) = 0.5$.

Figure 6 shows that the optimal directions are the same for most nodes on the grid. However the length of the lines and thus the expected finish times vary. If at the beginning of period two when a weather map is received, the last observation appears to be happening we could be 98.8% sure that the second member of the ensemble is happening and we would use the circle directions on the plot. So for this weather the expected finish times in the North are much smaller than in the South because member two predicts strong winds in the North. The opposite is true for a realised weather pattern of observation four.

At $x=0$, $y=3$ the optimal directions for the different probability states varies. This is because if observation four occurred we would be 98.8% sure that we would experience strong South winds and it is desirable to move into the south to finish the race in better time. However the opposite is true for observation five. Of course if the winds turn out to be either observation one, two or three i.e. the same everywhere then we think both members are equally so to avoid getting stuck in low winds later we would stay on the straight ahead course. This is because the members disagree on where the strong winds will be.

6 Conclusions

This problem is a new concept that has been shown plausible in this paper and pilot problems have been solved in reasonable time using the developed model.

Although this model is currently using very simple yacht movement calculations the model will run just as easily using more sophisticated calculations such as calculations based on a velocity prediction program.

The developed model is slightly limited in the fact that it can only deal with forwards yacht movement due to the backwards recursion in the dynamic program.

Further work could be done to develop a model that allows backwards movement for comparison.

Acknowledgements

The author acknowledges project supervisors Professor Andy Philpott and Dr. Andrew Mason for their patience, guidance and encouragement during the project period. Their enthusiasm for the topic and lively discussions made the project process a lot more interesting than first anticipated.

References

1. NOAA national weather service NCEP ensemble training page (2001 – last update) [Online] Available: http://www.hpc.ncep.noaa.gov/ensembletraining/ensembletraining.htm [2002, March 23]

2. Philpott, A. (2002) 'Stochastic optimisation in yacht racing'. Forthcoming in: *Applications of stochastic programming.* Eds: Ziemba, W. and Wallace, S.W. Springer Verlag.