Dagstuhl Seminar Proceedings 05031
Algorithms for Optimization with Incomplete Information
http://drops.dagstuhl.de/opus/volltexte/2005/65

Rowing to Barbados

Geoff Leyland* Andy Philpott†

March 4, 2005

Abstract

In October 2003, sixteen boats set off from La Gomera in the Canary Islands headed for Barbados 4800 km distant. Each boat was manned by two oarsmen who were competing in the Transatlantic Challenge, an ocean rowing endurance event. This paper describes an optimization model developed for route planning in this event. It was used successfully by the Holiday Shoppe team to win the race in world record time. We describe the tool, its history, and the way it was used in the race.

1 Introduction

In December 2002, Kevin Biggar walked into Andy Philpott's office at the University of Auckland with a proposal. He had heard of Andy's work with Team New Zealand on the optimal routing of America's Cup yachts (see [2],[3]), and wanted to use similar techniques to determine an optimal route to row across the Atlantic in the 2003 Woodvale Transatlantic Challenge (see [5]). Kevin had resigned from a high paying consultant's job in Auckland to pursue his dream, and had teamed up with Scott Donaldson and team manager Rob Hamill to put together a successful campaign for the race to begin in October 2003.

Since Kevin had a limited budget, Andy suggested that he run an undergraduate student project in the Department of Engineering Science at the University of Auckland - this would entail no cost but would carry no guarantees. Kevin was enthusiastic and the project began in March 2003 under the supervision of Andy Philpott and Andrew Mason (see [4] for the student's report). In July 2003 it became clear that more manpower was required, and so Geoff Leyland from Stochastic Optimization Ltd joined the team to construct and complete an implementation of some of the findings from the student project.

2 Solution Approach

The big challenge in providing a routing solution for the rowers was to offer something that could be used at sea, and that complied with the race rules - no outside strategic or routing assistance is allowed. Ocean racing yachts have on-board computers, weather faxes, radios and high-tech navigational equipment. The rowers did

*Stochastic Optimization Limited (www.sol.co.nz)
†Department of Engineering Science, University of Auckland, NZ

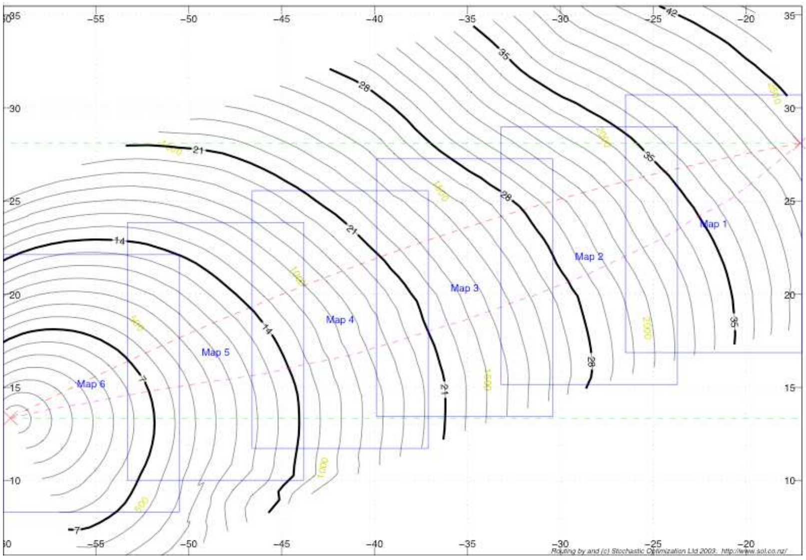

Figure 1: Isochrone map provided to rowers

in fact have some of this equipment, but so little power was available on board that none of it could be used effectively, and the environment so harsh that it couldn't be relied upon. Furthermore, there was so little room in the cabin (enough space for just one of the two rowers) that it was difficult to work on navigation.

The solution was to turn away from electronic solutions and keep things simple. We provided the rowers with a map, printed on six pieces of laminated A4 paper that they could use each day to determine, given the weather, the most effective direction in which to row. The maps are shown in Figure 1.

The contours on this map are known as *isochrones*, a term first coined by Francis Galton in a paper [1] that he submitted to the Proceedings of the Royal Society of London in 1873. They represent the contours of equal minimum expected time sailed by a vessel that follows an optimal course from their location to a given destination. The contours can be computed using dynamic programming. It is remarkable that Galton devised a methodology (and a machine!) for computing isochrones that predates Bellman's work by 80 years. Alas Galton's machine was (to our knowledge) never constructed.

The dynamic programming recursion used to compute the contours differs depending on assumptions on the information available to the rowers. The most convenient assumption is to assume that the random variables defining the weather and sea state are independent from the observations of these quantities on the previous day. We call this the *climatology* assumption. It is then possible to determine the best route from any point by computing the expected speed in any given direction, and then using the expected time to the destination at closer points to give an optimum. This is the approach adopted by Galton, and employed by most isochronal yacht-routing computer packages (that typically assume deterministic weather and sea state).

Because of the race rules, and the difficulty of getting live weather data onto the boat, the routing was based entirely on past weather data taken from routing charts. These charts are based on observations taken over 20 years by ships crossing the atlantic, and show what wind direction and strength each ship observed.

Since ships tend to take the same routes across the ocean, some areas have many thousands of observations, while others have very few. Ships also tend to cross the oceans more at some times of the year than others, but luckily, a chart is published for every month, so any bias is manageable. The chart for October is shown in Figure 2[1].

One of the circular diagrams (called a *wind rose*) on this chart is shown blown up in Figure 3. Each wind rose gives the total number of weather observations corresponding to its location, as well as the proportion of observations at each wind strength in each direction.

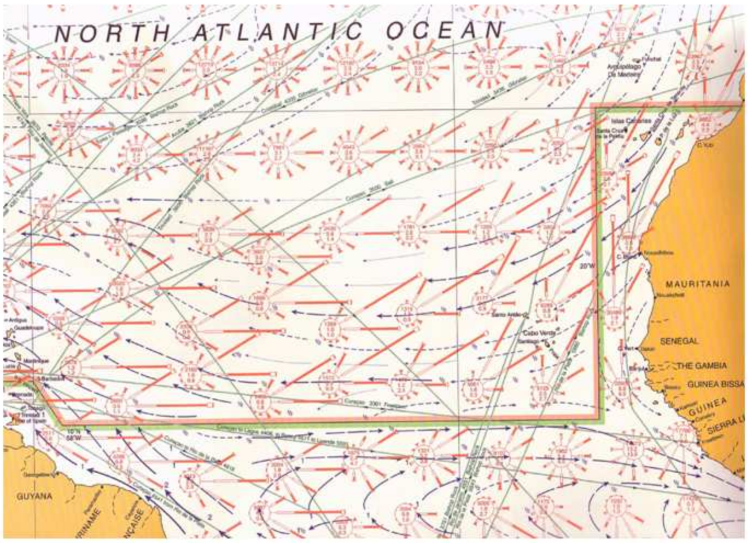

Figure 2: October climatology chart for the North Atlantic

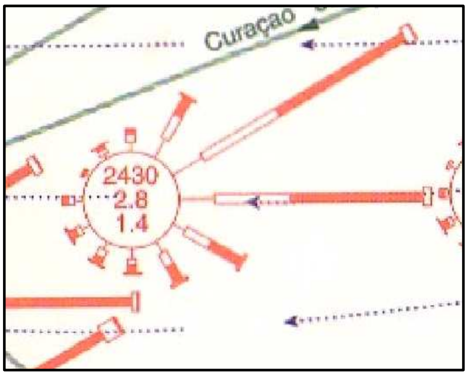

Figure 3: A wind rose (from Figure 2) for 2430 observations. The lengths of each segment gives the proportion of observations in each wind range 0-5kt, 5-10 kt, etc.

Our approach was to compute an optimal *policy* rather than an optimal route. The policy is conditioned on the particular observation of weather and sea state

[1]Figure 2 and Figure 3 are reproduced by permission of the Controller of Her Majesty's Stationery Office and the UK Hydrographic Office (www.ukho.gov.uk).

made by the vessel at its location, and so the expected minimum time to the destination is computed by taking expectations after an optimal heading is chosen. This is often called a *wait-and-see* policy as opposed to the *here-and-now* policy that chooses the optimal heading after computing expected speeds. Though familiar to mathematicians, the use of a policy rather than an optimal route was a novel idea for the rowers and it took some persuading to convince them that this provided them with much greater flexibility. Indeed during the race they used this to advantage at a critical time to allow them to overtake the race leaders and break the world record.

To compute the isochrones one needs a method for determining the speed of a rowing boat in a given set of weather and sea-state conditions. This was difficult to compute before the race as the crew in training before the race will not encounter enough weather conditions to be able to estimate a useful "polar" for the rowing boat giving its speed in any direction to a wind of given strength. Polars are typically computed by designers of ocean-racing yachts; for rowing boats, expert estimates of rowing efficiency were highly variable so an analytical model was likely to be unsuitable. We settled on estimates of velocity obtained from position and meterological data recorded in the previous race.

3 Using the Map

The reverse isochrone map is actually quite simple to understand. Each isochrone can be thought of as a "finish line" for a much shorter race than the one actually being rowed – getting to the next finish line first means we'll be in the best position to get to the subsequent line, and so on until we finish the race.

The hard problem of how to row from La Gomera to Barbados is reduced to a series of simpler problems of how to row over a day, this reduction being possible because we have worked out in advance how long (on average, if we follow the right strategy) it will take to row from each point in the ocean to Barbados.

The "daily race" isn't quite as simple as it could be – each day, instead of racing to a point (any real finish line looks like a point from far enough away) one has to race to cross a line. A first cut at the problem is to work out which point on the line is the closest to the boat's current position. However, because wind, currents and waves mean that the boat's speed is not the same in every direction, the real problem is to work out which point on the line can be reached the quickest.

To do this at sea, the crew plot their position on the map and then guess roughly where the nearest (in time) point on the next isochrone is, and take a bearing on it. They then row as best they can in that direction, and after 5 minutes use the GPS to see their actual course, and speed over the five minutes. They then change their course a little to the left and repeat the process, and then to the right, and try once again. It is important to note that it is practically impossible to row a fixed course in the middle of an ocean with strong winds, large waves and currents, so it is not feasible to, say row at 260, 265, 270, 275 and 280 degrees – the best one can do is try a rough course, and see how it goes.

Given this (rough) information about speeds on different courses, it is possible to work out what course to head to get to the next isochrone the fastest. If the course picked is the leftmost or the rightmost of the two tried, one could imagine repeating the above process with a new set of directions. In practice, however, even doing one iteration of the above algorithm proved difficult.

The transatlantic rowing race is one of the most gruelling sporting events in the world. Energy output of the top contestants (and input – how much they are eating) is comparable to cyclists in the Tour de France. However cyclists in the Tour de France cycle for about 100 hours, and have large support teams. The contestants in the Transatlantic race row for more like 600 hours each, are alone in the middle of the ocean, and eat dried food. Consequently, in times of high stress, even what seem like simple tasks become unusually difficult.

After the race, Kevin Biggar declared the isochrone map the "most useful piece of technology on the boat" – it gave them constant guidance as to where to head, how they were doing in the race (it proved an accurate predictor of time and distance to go, as well as a way of estimating leads), and in two cases inspired them to keep rowing when the rest of the fleet was on sea anchor. By adhering to the policy recommended by the map the Holiday Shoppe team won the race in record time: 40 days, 5 hours and 31 minutes, nearly 22 hours faster than the previous world record.

Back in New Zealand, we used the map to analyse the race. We were able to provide information on who was leading the race (a surprisingly subtle problem) and estimates of when the boats would finish, and the total distance they would row. These estimates were far more accurate than those published by the race organisers - our estimated distances to go were within 5km from about four days into the race. Estimates of the race finish time were less accurate - we didn't have a good model of how fast the boats would go, but were still significantly better than other estimates.

4 Aftermath: the protest

Holiday Shoppe were so successful in their campaign that the second placegetters, Team CRC, formally protested shortly after the race on the grounds that Holiday Shoppe had gone faster than physically possible. They offered little evidence to support their protest - some speculation that Holiday Shoppe may have used their solar panel as a sail, and a report from a naval architect explaining that the energy required to row was proportional to the cube of the boat's speed. In particular, they claimed that in the last two weeks of the race, Holiday Shoppe had inexplicably sped up, and that this could only be the result of cheating.

Because we had been following and analysing the race very closely, we had proof that the claims were untrue. Our analysis of the race was used to argue Holiday Shoppe's innocence when the protest went before the race committee in London. We showed that:

1. With respect to allegations of use of a sail, that Holiday Shoppe made their biggest gains into headwinds - they won because they did not go on a sea anchor when CRC did. CRC made their large gains when winds were favourable.

2. With respect to the naval architect's claims that it was very difficult to row as fast as Holiday Shoppe had gone, we showed that it was in fact CRC that had had the highest peak speed of the two - and that CRC had at one stage had the highest speed advantage over Holiday Shoppe when they were rowing in roughly the same part of the ocean.

3. With respect to the claim that Holiday Shoppe had sped up inexplicably in the last two weeks of the race, we showed that both boats had, on average, slowed down. CRC had slowed more.

The race committee dismissed all claims in CRC's protest, and Holiday Shoppe were finally declared official winners and world record holders more than two months after they crossed the finish line.

5 Conclusions

This project was a very successful application of stochastic optimization techniques. Our "client" was happy with the results and, even better, won the race. Part of that success came about because we spent a lot of time thinking of how to present the results of a reasonably complex calculation in a tangible and usable manner.

For the rowing race, given the race rules and the onboard equipment (or lack thereof) the climatalogical assumption was reasonable. In yacht races, more real-time weather information is available, and the races are often short enough that good weather forecasts are available for a significant part of a race. Planning based around forecast weather, rather that historical weather, becomes important. One approach that handles correlation in the weather uses branching scenarios as described in [3]. The boundary conditions for such a model are typically defined by climatology, but a statistically and meteorologically sound methodology for defining scenarions and blending these together with climatology remains to be worked out.

References

[1] Galton, F. On the employment of meteorological statistics in determining the best course for a ship whose sailing Qualities are known. *Proc Royal Society*, 21, 1873.

[2] Philpott, A.B. Stochastic optimization in yacht racing, to appear in *Applications of Stochastic Programming*, W. Ziemba and S. Wallace (ed.), SIAM, 2005.

[3] Philpott A.B. and Mason, A.J. Optimizing yacht routes under uncertainty, in *Proceedings of the Chesapeake Bay Yacht Symposium*, 2001, 89-98.

[4] Stankovich, I. Routing for the trans-Atlantic rowing race, *Year 4 Project Report*, Department of Engineering Science, 2003.

[5] http://www.woodvale-events.com/default.asp?ct_id=50