

Testing various backtracking algorithms in airborne maritime surveillance modelling

D. O. Marlow^a and J. E. Murphy^a

^a Joint and Operations Analysis Division, Defence Science and Technology Organisation, Fishermans Bend, Victoria

Email: david.marlow@dsto.defence.gov.au

Abstract: Maritime surveillance is an important task for the Australian Defence Force. Australia has a vast coastline, many offshore territories and extensive offshore resources in fisheries, oil and gas. There is a need to monitor these areas for illegal activities a need met by ships from the Royal Australian Navy and aircraft from the Royal Australian Air Force (RAAF) and Border Protection Command (BPC).

In a typical airborne maritime surveillance scenario, an aircraft will fly in a specified Area of Interest (AI) – in this case a corridor – and seek to detect and classify all the ships (i.e. a ‘classify-all’ search) that move up and down the corridor. In this particular instance a corridor search of a shipping lane is considered, where an aircraft flies up the corridor as ships move up and down. Security of shipping lanes is of vital interest to nations, especially maritime nations such as Australia. Major shipping lanes occur to Australia’s north, especially around Singapore. Merchant and cargo ships carry trade from the Suez Canal and Persian Gulf via the Straits of Malacca and on to North Asian countries such as China, Japan and Korea. Much of Australia’s exports such as iron ore and coal pass through this region.

In previous maritime surveillance modelling work, it was found that an aircraft may struggle to classify all targets in an AI in the available time, due to the high ship density or the lack of endurance of the aircraft. Either the aircraft does not complete the search of the whole AI or not all targets are classified. This work considers ways to address this scenario through testing various ‘backtracking’ algorithms. These backtracking algorithms allow the aircraft to consider classifying ships that are detected behind the aircraft (such as ships advancing through the shipping lane) while still ensuring that the whole area is searched.

The backtracking algorithms tested here can be grouped into four categories:

- *On/off.* Backtracking is either unconstrained (where the aircraft classifies the next ship regardless of position) or ignored (so the aircraft only ever classifies ships ahead of it).
- *Time-based.* The aircraft backtracks depending on its *pro rata* progress through the AI. For example, if the aircraft has reached half of its maximum endurance, the aircraft may be excluded from searching the first half of the corridor.
- *Classification-based.* Here the aircraft is allowed to classify a maximum number of ships behind it, before being compelled to classify at least a minimum number ahead of it.
- *Dynamic.* In this case, the decision to classify a ship behind the aircraft is determined using probabilities. If the aircraft is ahead of where it needs to be in the AI based on the time remaining, there is a greater probability that it will be permitted to classify a ship behind it. The probability depends how far ‘ahead of schedule’ it is. If it is behind, it will be forced to move forwards. The probability decreases to zero by the end of the mission to force the aircraft to fully cover the area.

A simulation model is used to conduct the analysis for a generic helicopter searching a corridor 250 nautical miles long and of varying widths. Ships travel at speeds of 15 knots or 25 knots up or down the corridor and various densities are considered. A Nearest Neighbour search is used for simplicity to test the various backtracking options. The Measures of Effectiveness (MOEs) are the percentage of ships classified and the percentage of the area searched. It is found that for corridors of narrow width or low numbers of ships, techniques favouring greater backtracking perform best. Here the classify-all search is achievable, as there is time available to explore the area given the lower ship density. In higher-density and wider corridors, where the need to divert is greater, there is less time available for such diversions. In these cases, the dynamic backtracking and linear time-based backtracking options are best, in that they ensure that the search is completed and achieve the highest classifications given that pre-condition. Both these techniques track the *pro rata* progress of the aircraft through the corridor, thus helping to ensure that the search is completed. Other options may achieve more classifications but not complete the search.

Keywords: *Maritime surveillance modelling, simulation, heuristics, backtracking, analysis*

1. INTRODUCTION

1.1. Problem description

Airborne maritime surveillance remains an important role for Australia. Given its vast coastline, the Australian Government seeks to protect its approaches and offshore territories from illegal activities such as fishing and people smuggling. The recent 2013 Defence White Paper (Commonwealth of Australia, 2013) asserts that “Defence will continue to provide effective intelligence, surveillance and reconnaissance capability across Australia’s vast maritime area of interest”. To this end, the Government intends to acquire new P-8A Poseidon aircraft to be operated by the Royal Australian Air Force (RAAF). Unmanned aircraft may also feature in a future force mix. Border Protection Command (BPC) also conducts airborne maritime surveillance, operating a fleet of fixed-wing and rotary-wing aircraft.

In a maritime surveillance scenario, an aircraft searches an Area of Interest (AI) for either a particular contact (ship) of interest, or all contacts – i.e. a “classify-all” search. The size of the area to be searched will depend on the speed and endurance of the aircraft, so a fixed-wing aircraft such as the P-8A travelling at 300 knots with an 8 hour endurance will cover a larger area than a helicopter with travelling at 100 knots with a 3 hour endurance. Generally an aircraft will move through the AI by following waypoints, which are fixed geographical areas designed to guide the search. Aircraft will use their sensors such as radars or electro-optics to look for ships. When a detection is made, the aircraft will divert from the direct line between the waypoints to classify the ship again using appropriate aircraft sensors or by visually identifying the contact. Sensor ranges are influenced by the environmental conditions, such as weather and sea state, as well as aircraft altitude and ship type.




Figure 1. Searching a shipping lane. In this instance a helicopter moves through the corridor between waypoints (red crosses) seeking to classify all ships (black dots) moving through the lane. As ships come within detection range of the aircraft sensors (blue dashed line), they are detected (blue dots) and classified (red dots).

The problem is illustrated in Figure 1 with a helicopter searching a shipping lane. The aircraft will move between the two waypoints in a direct line. When a ship is detected the aircraft will divert from this line in order to attempt to classify it.

In this work, a classify-all search of a shipping lane is considered. Shipping lanes are areas of high-density traffic, generally

merchant or cargo ships, where ships move in a two-way direction. As the shape of a shipping lane is analogous to a road, the search can be likened to a corridor search.

A map of the world’s shipping lanes is shown in Figure 2. Major shipping lanes occur in the South China Sea and the Straits of Malacca, carrying cargo to and from the Suez Canal and Persian Gulf to Asian ports in China and Japan via Singapore. Much of Australia’s trade passes through the North Asian region around the South China Sea. Being a maritime nation, Australia is heavily dependent upon such trade, and ensuring its security is paramount.




Figure 2. Global shipping routes (Figure 1(a) from Kaluza et al., 2010. Reproduced by permission)

This work examines algorithms for ‘backtracking’ used during a search – that is, if an aircraft should turn back in order to classify a ship that has been detected behind it. This ship may have recently entered the AI, or not been previously detected. In previous work (e.g. Marlow et al., 2009), the aircraft has not been constrained in this way. The corollary of this is that aircraft may not search the whole AI in the time available. Consequently, either the search area must be reduced in an *ad hoc* manner, or not all ships in the original AI will be classified. The aim of this work is to find a backtracking algorithm that ensures that the aircraft keeps moving in a generally forwards direction through the corridor and thus covers more of the area.

1.2. Previous work

Although aspects of maritime surveillance are well studied (e.g. Koopman, 1980), literature around classify-all maritime surveillance modelling remains quite sparse. One paper (Reiss, 1980) examined a similar problem in optimising search techniques including barrier patrols for ship movement down a corridor. This paper accounted for an aircraft conducting a classify-all search, but the ships are only moving in one direction towards the approaching aircraft. Another paper (Grob, 2006) is similar to this work and describes a different model, where the Measures of Effectiveness (MOEs) include to detect and classify as many ships as possible. That work examined varying aircraft altitudes and the variation of classification range with ship type, but did not include endurance limits for aircraft and a finite radar detection range.

Further papers on maritime surveillance modelling treated the problem as a variation of a Travelling Salesman Problem (Kilby et al., 2008). The focus of these papers were to test various search techniques (Marlow et al., 2007), use novel techniques such as ‘ghost ships’ to attempt to maximise classifications (Marlow et al., 2009) and examine the impact of varying the aircraft detection and classification range on the number of ship classifications (Mercer et al., 2008). That work looked at a square-shaped AI and ships moving with random velocities, with aircraft undertaking generally a ‘ladder’ or ‘lawnmower’ search through the AI. Other examples of work used a genetic algorithm approach to solve for stationary ships (Sancı and İşler, 2011), and a branch-and-bound approach to solve for moving ships (Öner and Kayaligil, 1999), but in the latter case only for problems with a very small number of ships. In this work, the corridor search can be considered as a single leg of the square-shaped AI search, with ships moving up and down the corridor.

In earlier work (e.g. Marlow et al., 2009), issues arose with ‘budgeting’ – i.e., determining which ships should be visited at which stage of the mission given the endurance constraint on the aircraft. This paper is the initial attempt in examining potential budgeting options.

2. DESCRIPTION OF BACKTRACKING ALGORITHMS

The backtracking algorithms considered in this work have been extended from work (Murphy et al., 2013) which used back-of-the-envelope calculations to gain insights into maritime surveillance problems. The options tested in this paper are explained below.

2.1. No backtracking and unconstrained backtracking.

Two simple options are no backtracking and unconstrained backtracking. No backtracking simply means that the aircraft is always travelling forwards in the corridor. The likely impact of this is that the aircraft will bypass many ships moving through the corridor. For unconstrained backtracking, as the name suggests, the consequence is likely to be that the aircraft may not be able to search the entire AI in the available time.

2.2. Time-based backtracking

Time-based backtracking forces the aircraft to keep moving through the AI in accordance with the time remaining for the mission. Two cases are tested here, being:

- *‘Linear’ backtracking.* Here the aircraft is allowed to backtrack in accordance with the length of the AI and the available time in the AI. For available time T hours and corridor length L nautical miles, the aircraft is excluded from backtracking into an area that is Lt/T nautical miles into the corridor at any time t . Thus at $t = \frac{1}{2}T$, the aircraft can no longer enter the first half of the AI, and at $t = \frac{3}{4}T$, it can only stay in the last quarter of the AI, etc.
- *‘Quadratic’ backtracking.* The difference with quadratic backtracking is that the aircraft can now no longer search an area defined by $L(t/T)^2$ nautical miles into the corridor. Thus at $t = \frac{1}{2}T$, the aircraft can no longer enter the first quarter of the AI. This allows more scope for an aircraft to search the portion of the AI closest to the initial waypoint, but forces it to move more quickly towards the end of the search.

2.3. Classification-based backtracking

Classification-based backtracking forces the aircraft to keep moving generally forwards, but does not ensure that the aircraft will search the entire AI in the time available. Here the classification-based backtracking algorithms tested here are of the X-Y form, where X represents a *minimum* number of *forwards* movements, and Y represents a *maximum* number of *backwards* movements.

Therefore, ‘one-two’ backtracking means that, if the previous two aircraft movements to classify a ship have been backwards, then the next aircraft movement to classify a ship must be forwards. The aircraft may then

classify a ship behind it again, although it is also permitted to move forwards. By contrast, ‘three-one’ backtracking means that, if the most recent classification was the result of a backward movement, then the next three ships to be classified must each be ahead of the aircraft in the corridor. Once this is achieved, the aircraft is allowed to classify a ship behind it again, although it can still classify a ship ahead of it in the corridor if it is closer.

Figure 3 illustrates examples of both two-one and one-two backtracking. In two-one backtracking, only one consecutive backwards movement is allowed at a time. For one-two backtracking, two consecutive backwards movements are allowed. In this particular example, all the ships in this particular part of the corridor are allowed to be classified as a result of this policy.

2.4. Dynamic backtracking

In dynamic backtracking, whether or not to backtrack is dependent on the progress through the AI compared with what would be expected at that stage of the mission. This amount of backtracking depends on the current status of the search.

When each new event occurs (such as a new detection or classification), the dynamic backtracking algorithm is called. Firstly, bounds for the probability of backtracking are set depending on the time remaining. Initially this is between 10% and 90%. As the mission continues, the probability of backtracking is reduced in order to ensure forward motion. Until the half-way stage of the mission, the bounds reduce linearly to between 25% and 75%. From that point, the bounds reduce to ensure that by the end of the mission, the probability of travelling backwards is zero. Two methods are used to change the bounds:

- a linear method (‘dynamic-lin’), where both the upper and lower bounds decrease linearly to zero from the half-way point to the end; and
- a quadratic method (‘dynamic-quad’), where the upper and lower bounds decrease as a quadratic function from the half-way point to the end. This means that the probability will be higher initially but decrease rapidly towards the end.

These bounding procedures are similar to those of the non-dynamic ‘linear’ and ‘quadratic’ backtracking algorithms provided in Section 2.2. The variation of the probability bounds with time is shown in Figure 4.

In dynamic backtracking, when the mission begins, the probability of backtracking will be 50%. As the search continues, the probability of backtracking will be determined based on

the aircraft’s relative position in the corridor. If the aircraft is ahead of where it needs to be on a *pro rata* basis, it has scope in its search to explore backwards in the corridor. The probability of backtracking is then determined based on how far ahead of the *pro rata* position it is, bounded by the corridor length. If X_a denotes the aircraft’s position in the corridor relative to length and X_p the *pro rata* position, then the ‘probability differential’ P_{diff} is denoted by:

$$P_{diff} = P_b(X_a - X_p) / (\min(L, V_s t) - X_p).$$




Figure 3. Examples of (above) two-one backtracking and (below) one-two backtracking. Ships are shown as black dots and the aircraft path as red arrows

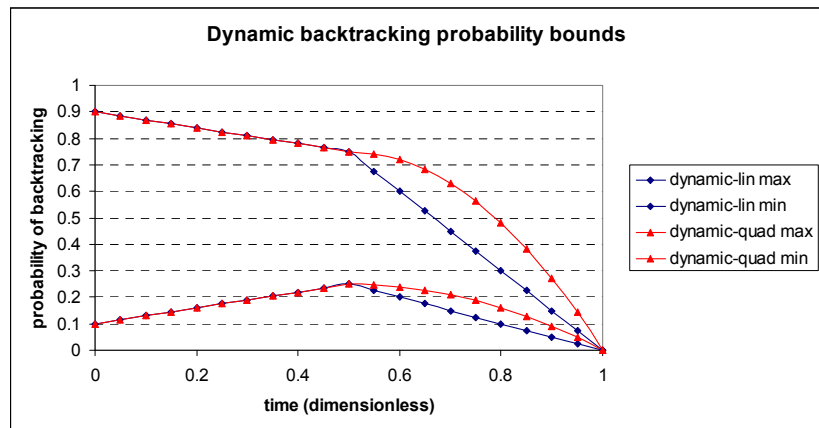


Figure 4. Dynamic backtracking probability bounds

Here P_b is the maximum extra probability (i.e. the difference between the maximum bound and the average) and V_{st} is the maximum distance the aircraft could have travelled (along the wayline) to that point. This value is then added to the average probability at the time to generate a probability of backtracking.

Conversely if the aircraft is in a position such that it is behind the *pro rata* position, the probability of going backwards is set to zero. The aircraft is therefore forced to fly forwards in order to ‘catch up’ its position in order to complete the search in the required time.

3. RESULTS

The backtracking options were tested in an adapted version of a simulation model used in previous work (Marlow et al., 2009). Ships are initially randomly placed within the AI with random velocities. The model ensures that the number of ships in the AI is maintained on average over multiple runs.

Table 1 shows the input data used in the model for a generic helicopter. The ship speeds are chosen as 15 knots or 25 knots, befitting the speed of merchant or cargo ships in a shipping lane. All ships are assumed to travel at the chosen speed in the corridor. Ships will travel either up or down the corridor, and the random movement of other ships of different types through or across the corridor is not considered. It is important to note that corridor widths are deliberately chosen so that the detection range will always exceed the corridor width. This allows the aircraft to ‘see’ across the width of the corridor, at least when in the middle of the corridor. Classification ranges have been chosen to be very small in this work, equating to a requirement for visual classification.

Table 1. Input data for maritime surveillance model for generic helicopter

Name	Units	Number
Corridor width	Nautical miles	10, 20, 40
Corridor length	Nautical miles	250
Aircraft speed	Knots	100
Aircraft endurance	Hours	4
Detection range	Nautical miles	50
Classification range	Nautical miles	1
Number of ships	Number	10, 50
Ship speed	Knots	15, 25

For simplicity and ease of comparison a Nearest Neighbour search technique is used. The aircraft will seek to intercept each ship, and thus will not chase a ship that it calculates will exit the corridor before it can intercept it. If a new contact will be closer than the ship currently being ‘chased’ to classify, it will break its current route and divert to the new contact. Results are provided for 100 runs for statistical significance.

The MOEs for this analysis, given the dual assumed aims of a particular maritime surveillance mission, are:

- The percentage of ships classified (where this is measured as the number classified over the average number of ships that were in the AI during the mission); and
- The final aircraft position, expressed as a percentage based on the distance from the final waypoint.

Therefore, the ideal algorithm will have a high level of classifications and will finish the search close to or at the final waypoint. The last measure indicates that the aircraft has fully searched the AI. Two additional parameters are included: the mission end time (i.e. when the aircraft reached the waypoint) and the distance travelled per ship classified.

Results for the case of a 250x10 and 250x20 corridor with 10 ships are not shown as there is no significant difference between the methods. As the ship density is low and corridor narrow, the helicopter can complete the search well within the required time. For a 250x10 corridor, the classification rates vary from 95.6% to 99.6% (15 kn ship speed) and from 94.7% to 99.7% (25 kn). The percentage of ships classified is relatively unchanged with ship speed (Murphy et al., 2013).

Table 2 shows the results for all backtracking algorithms tested in this work against the MOEs. This shows the results for a 250x40 corridor with a ship speed of 15 knots for 10 ships (equating to a density of one ship every 1000 square nautical miles). In this example some differences between the methods begin to emerge. In each case the aircraft reaches the end point with time to spare. However, if there is no backtracking the percentage of ships classified is much lower. Methods that favour backtracking such as one-two and one-three, as well as unconstrained backtracking, perform best in these instances, as they allow more time for the corridor to be searched. Results are similar for ship speeds of 25 knots.

Table 2. Mean results and 95% confidence intervals for 250x40 corridor, ship speed 15 kn, 10 ships

Backtracking option	% classified	% final position	Mission end (hr)	Distance / ship classified
None	80.9±2.2	99.7±0.0	2.91±0.05	36.5±1.9
Unconstrained	93.2±2.3	99.7±0.8	3.17±0.07	35.4±1.8
Linear	83.8±2.2	99.7±0.0	3.03±0.05	37.1±2.2
Quadratic	88.8±2.4	99.6±0.1	3.11±0.07	35.5±1.8
One-two	90.6±2.4	99.3±0.4	3.16±0.07	35.5±1.8
One-three	93.1±2.3	99.3±0.4	3.17±0.08	35.5±1.8
Two-one	88.4±2.5	99.4±0.3	3.12±0.07	36.0±1.9
Three-one	84.6±2.5	99.6±0.1	3.05±0.06	36.0±1.9
Dynamic-lin	85.8±2.4	99.6±0.1	3.08±0.07	36.9±1.9
Dynamic-quad	86.3±2.5	99.7±0.0	3.09±0.07	36.8±1.9

illustrate the differences between the algorithms. If there is no backtracking, the aircraft moves through the area and reaches the final waypoint in the quickest average time but with a lower number of classifications. Conversely, if backtracking is unconstrained, the aircraft will on average finish less than two-thirds of the way through the corridor. Similarly, pre-determined classification-based backtracking algorithms such as one-two and two-one, while achieving generally a higher percentage of classifications, also have the aircraft finishing well before the final waypoint. Those with a more forwards-leaning emphasis such as two-one and three-one finish closer to the end point.

The time-based linear and quadratic backtracking algorithms perform slightly differently. The linear algorithm is more likely to ensure the aircraft completes the search in the required time, while the quadratic algorithm gives the aircraft too much time to search the initial part of the corridor, resulting in an inability to reach the end point in time. The dynamic algorithms both reach the waypoint within the required time and achieve a percentage of classified ships result that is on average higher than the time-based backtracking methods, but lower than the classification-based methods.

The distance per ship classified does not vary greatly between the options, with the quadratic search being the most inefficient. Results are similar for ship speeds of 15 knots.

For a 250x20 corridor with 50 ships (not shown), some of the trends in the Table 3 results are found, with time-based and dynamic procedures ensuring that the aircraft completes its search, while the others do not. However, the reduced search area means that some classification-based options (two-one and three-one) can complete the search (in about 25% of runs). They also continue to provide superior classification rates: by ~5% over the dynamic options and ~2% over the linear option. These results suggest that, depending on the problem, different classification-based heuristics may provide the best overall results against both MOEs.

Table 3 shows the results for a 250x40 nautical mile corridor with a ship speed of 25 knots and 50 ships. This may be likened to a corridor in the Straits of Malacca, which carries around 50,000 vessels per year (Evers and Gerke, 2006). Assuming a constant rate of flux, this means that 5.7 ships per hour pass through the straits (both ways). Therefore for a 4 hour mission in this corridor for these ship speeds, the average number of ships is around 57, compared to the 50 used here.

These results more clearly

Table 3. Mean results and 95% confidence intervals for 250x40 corridor, ship speed 25 kn, 50 ships

Backtracking option	% classified	% final position	Mission end (hr)	Distance per ship classified
None	59.5±1.4	99.5±0.2	3.55±0.05	12.1±0.3
Unconstrained	65.4±1.8	62.8±4.7	4.01±0.01	12.7±0.4
Linear	62.0±1.6	99.6±0.1	3.84±0.03	12.7±0.3
Quadratic	60.3±2.3	93.7±0.6	4.01±0.01	14.0±0.5
One-two	66.8±1.8	77.1±3.1	4.00±0.01	12.4±0.3
One-three	66.1±1.8	70.1±3.9	4.00±0.01	12.6±0.4
Two-one	67.7±1.6	86.9±2.0	3.98±0.02	12.1±0.3
Three-one	68.1±1.4	89.3±1.8	3.98±0.02	12.0±0.3
Dynamic-lin	63.6±1.3	98.8±0.4	3.83±0.03	12.4±0.4
Dynamic-quad	64.0±1.4	98.7±0.5	3.85±0.03	12.4±0.4

4. SUMMARY

This paper tests various algorithms where aircraft may have to backtrack in order to conduct a classify-all maritime surveillance search in an area of interest (in this case a shipping lane). For narrow corridors or low density environments, a method which takes advantage of the additional time to exploit backtracking may be advantageous. In this case, there is time available to explore the AI, so unlimited backtracking or methods such as one-two or one-three can exploit this circumstance. When the time available to search becomes more constrained due to the greater area to be covered and the number of ships to classify, algorithms which track the progress of the aircraft through the AI perform best. These algorithms, such as 'linear' or dynamic options, allow limited scope for additional exploration depending on the corridor width and aircraft endurance. Overall, these findings suggest that the corridor area should be set such that the aircraft has time to broadly explore it, in order to maximise classifications while completing the search. Combining these algorithms with an heuristic search technique such as a genetic algorithm is a potential area for further work.

ACKNOWLEDGMENTS

The authors would like to thank David Cox from Joint and Operations Analysis for his helpful comments, and Helen Pottage from the DSTO Library for their assistance with this paper and sourcing references.

REFERENCES

- Commonwealth of Australia (2013). *Defence White Paper 2013*, Australian Government Department of Defence, 132pp.
- Evers, H.-D. and Gerke, S. (2006). The strategic importance of the Straits of Malacca for world trade and regional development. ZEF Working Paper Series No. 17, Department of Political and Cultural Change, Center for Development Research, University of Bonn.
- Grob, M.J.H.B. (2006). Routing of platforms in a maritime surface surveillance operation, *European Journal of Operational Research*, 170, 613-628.
- Kaluza, P., Kölzsch, A., Gastner, M.T. and Blasius, B. (2010). The complex network of global cargo ship movements. *Journal of the Royal Society: Interface*, 7(48), 1093-1103.
- Kilby, P., Tobin, P., Luscombe, R., Barry, S. and Hickson, R. (2008). The maritime surveillance problem. In T.R. Marchant, M. Edwards and G.N. Mercer (eds.), *Proceedings of the 2007 Mathematics-in-Industry Study Group*, 32-56.
- Koopman, B. O. (1980). *Search and Screening*. Pergamon Press, New York.
- Marlow, D.O., Kilby, P. and Mercer, G.N. (2007). The travelling salesman problem in maritime surveillance – techniques, algorithms and analysis. In Oxley, L. and Kulasiri, D. (eds), *MODSIM 2007 International Congress on Modelling and Simulation*. Christchurch, New Zealand, 10-13 December 2007, 684-690.
- Marlow, D.O., Kilby, P. and Mercer, G.N. (2009). Examining methods for maximising ship classifications in maritime surveillance. In Anderssen, R.S., R.D. Braddock and L.T.H. Newham (eds.) *18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation*, July 2009, 1629-1635.
- Mercer, G.N., Barry, S.I., Marlow, D.O. and Kilby, P. (2008). Investigating the effect of detection and classification range and aircraft dynamics on a simplified maritime surveillance scenario, *ANZIAM Journal*, 49, C475-C492, Proceedings of the 8th Biennial EMAC 2007 Conference, G.N. Mercer and A.J. Roberts (eds.)
- Murphy, J. E., Marlow, D. O., Cox, D. and Tu, Z. (2013). A Back of the Envelope Calculation of a Maritime Surveillance Problem, DSTO Technical Report in preparation, Defence Science and Technology Organisation.
- Öner, A. and Kayaligil, S. (1999). Flight Optimization in Reconnaissance Missions. In RTO SAS Symposium, *Modelling and Analysis of Command and Control*, Issy les Mouleaux, France, 12-14 January 1999, 22:1-15.
- Reiss, M. (1980). On the optimization of a search pattern for maritime surveillance. *Search Theory and Applications*, NATO Conference Series Volume 8, Springer-Verlag, 141-153
- Sancı, S. and İşler, V. (2011). A Parallel Algorithm for UAV Flight Route Planning on GPU. *International Journal of Parallel Programming*, 39(6), 809-837.