18th World IMACS / MODSIM Congress, Cairns, Australia 13-17 July 2009
http://mssanz.org.au/modsim09

Examining methods for maximising ship classifications in maritime surveillance

Marlow, D.O. [1], **P. Kilby** [2] **and G. N. Mercer** [3]

[1] *Defence Science and Technology Organisation, Melbourne, Victoria*
[2] *National ICT Australia, Canberra, ACT*
[3] *University of New South Wales at the Australian Defence Force Academy, Campbell, ACT*
Email: david.marlow@dsto.defence.gov.au

Abstract: Both Royal Australian Air Force (RAAF) and Border Protection Command (BPC) aircraft fly maritime surveillance missions for the Australian Government on a regular basis. These missions involve searching particular areas of interest (AIs) for illegal fishing or people smuggling activities. In this work a square AI is considered, with an aircraft tasked to search for ships using its sensors (e.g., radar, electro-optical). Waypoints (points that must be visited) are included in the AI to ensure that the entire AI is covered. The aim for a particular search is to detect and classify as many ships as possible, while doing so in the shortest time. Depending on the ship density, the aircraft may not have time to search the entire AI.

In this paper an augmentation of the traditional Travelling Salesman Problem (TSP) is considered, where the ships (cities) move with random velocities (dynamic TSP), have different start and end points (open TSP) and there is incomplete *a priori* knowledge of the problem space (on-line TSP), making this problem much more challenging.

Earlier work (Marlow et al, 2007) considered a “baseline” case of an S-shaped search pattern and a default heading (the route flown when there are no ships currently detected) direct to the next waypoint. This paper considers three extensions with an aim to increase the level of ship classifications: these are 1) alternative initial flight paths, 2) alternative default headings and 3) including “ghost ships” in the search.

The principle behind the S-shaped pattern is that the aircraft will cover the entire AI with its sensors, giving aircrew the best chance of detecting all ships in the AI. This paper considers alternative spiral waypoint patterns, both an “inspiral” (from one corner of the AI, spiraling towards the centre) and an “outspiral” (the reverse). These approaches are theoretically more likely to detect ships that enter the AI during the mission.

The direct-to-waypoint default heading will minimise the travel time, but it also may potentially result in the aircraft not covering the entire AI with its sensors, particularly if it has already been diverted significantly from the “wayline” (the direct line between waypoints). In this work, a perpendicular return to the wayline is considered, which increases the distance travelled but is also likely to increase the probability of detecting more ships. A third option is also considered in which the aircraft continually aims for the midpoint on the wayline of the perpendicular intercept point and the waypoint.

The object of ghost ships is to direct the aircraft to fly to areas of the AI that it may not otherwise visit. This may particularly be the case in low-density environments where the aircraft has to substantially divert from the wayline to classify ships and, in returning, inadequately cover other areas where ships may be present. Ghost ships remain in the current tour until they are “detected”, whereupon they are removed.

Results suggest that, at lower densities, the perpendicular default heading and including ghost ships provide an overall improvement in classifications of the order of a few percentage points, which in real terms translates to an extra 1-2 ships on average. Significantly it also translates to an increase in the percentage of cases where 100% classifications are achieved. These improvements generally come at a cost of increased distance travelled by the aircraft and thus greater fuel consumption. In the case of ghost ships, there is an additional cost in computational time due to the requirement to include them in the tour. Beyond the critical ship density (where classifications cannot physically reach 100%), these variations offer no real advantage and in some cases are counter-productive, so the baseline pattern is more appropriate in these circumstances.

Keywords: *maritime surveillance, ship classifications, mission time, flight path, default heading, ghost ships*

1. INTRODUCTION

1.1. Problem description

Both Royal Australian Air Force (RAAF) and Border Protection Command (BPC) aircraft fly maritime surveillance missions for the Australian Government on a regular basis. These aircraft include AP-3C Orion and de Havilland Dash 8 fixed-wing aircraft and Eurocopter Squirrel rotary-wing aircraft. Maritime surveillance is a vital role for the Australian Government in seeking to defend Australia's northern approaches from criminal activities such as illegal fishing or people smuggling.

The problem is illustrated in Figure 1. Aircrew are tasked with searching a particular Area of Interest (AI). In some cases, the mission aim may be to search for a single target; in other cases, the aim may be to classify every ship in the area. The size of the area depends on the type of aircraft. A fixed-wing aircraft can travel faster, carries more fuel and has better sensors than a rotary-wing aircraft, so its allocated AI will generally be larger. Here a simple square AI is shown, although AIs can be a more general region. For the S-shaped search pattern shown, the waypoints (points that must be visited) are determined in such a way that if the aircraft flies directly between each waypoint (along the "wayline") without diverting from that path, the aircraft's sensors will cover the entire AI in the process. The AI can be divided into segments, the sizes of which correspond to the area covered by the aircraft's sensors if it follows the wayline between two sequential waypoints.

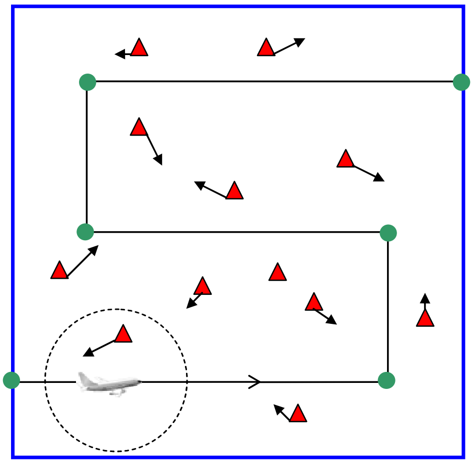

Figure 1. The maritime surveillance scenario. An aircraft moves through the AI following the direct line (wayline) between waypoints (green circles). As ships (red triangles) come within detection range of its sensors (dashed line), it diverts to classify them.

When an aircraft detects a ship with its sensors (varying from long-range radars to visual), the aircrew will divert the aircraft from the current flight path to classify it. Both detection and classification range vary with sensor type and are influenced by environmental conditions, such as cloud or sea state. The aim in this particular problem is to classify as many ships as possible before completing the search of the AI or before the aircraft reaches its maximum flight time (e.g., fuel limit). Having accomplished this, the secondary aim is to perform this as efficiently as possible to minimise fuel usage – i.e., in the shortest possible time.

If the ships are treated as stationary, and the aircraft starts and finishes at the same point and has an infinite detection range, then this problem is a version of the Travelling Salesman Problem (TSP). Given the nature of this particular problem – i.e., maximising ship classifications as the primary objective, with minimising tour length as the secondary objective or as a constraint when considering the maximum flight time – this work can be sub-classified as an orienteering problem (Feillet et al, 2005) in the TSP domain. However, the added difficulties of moving ships (dynamic TSP), detecting ships as the aircraft flies (on-line TSP) and different start and end points (open TSP) make this problem more challenging.

1.2. Previous work

The maritime surveillance modelling literature specific to this problem is sparse. Many papers exist that study elements of the maritime surveillance modelling problem, but other than that undertaken by the authors and their associates, only one (Grob, 2006) is currently published that directly relates to the problem considered here. Other aspects of the problem have been addressed, such as the best way to fuse information from different sensors to form maritime tracks (St-Hilaire et al, 2008) and how best to search a sequence of sub-regions in order to maximise the probability of detection of a single target (Ng and Ghanmi, 2002). One paper (John et al, 2001) looks at determining the shortest surveillance path in a single area of interest, with various scan and flight options available to minimise the distance travelled. Another paper (Ng and Sancho, 2007) examines an aircraft surveying multiple AIs in a single mission. An aircraft searches for a single target of interest in each area, with the entry and exit points to each area varied in order to minimise the overall distance travelled. These papers generally apply solution methods that draw on the abundant TSP literature.

A previous MODSIM paper (Marlow et al, 2007) on maritime surveillance modelling compared the simple Nearest Neighbour search with variations on the 2-opt method (Croes, 1958) for a range of ship speeds and ship densities. The 2-opt methods were found to be superior in maximising ship classifications. Other work

(Mercer et al, 2008) studied the impact of varying detection and classification ranges on ship classifications and the impact of aircraft turning circle radius on tour length. Another paper (Looker, 2008) has developed a robust means of determining the intercept point (and if one exists) when incorporating the turning circle.

2. FLIGHT PATH OPTIONS

This section describes three variations from previous work (Marlow et al, 2007) tested in the current paper.

2.1. Waypoint position

Waypoint positions are chosen based on the size of the AI and the sensor range of the aircraft, so as to maximise the probability that the entire AI is searched during the mission. Given that the AI is simply an allocated sub-set of the ocean, ships can be moving in and out of the AI while the aircraft is on task. The current approach (an S-shaped pattern) progresses slowly "upwards" through the AI in a way that resembles conventional barrier patrols (Wagner et al, 1999). Given the size of the AI, it is possible that ships at the "top" of the AI may escape before the aircraft is able to detect and classify them. Alternatively, ships may enter from the "bottom" of the AI once the aircraft has moved towards the top.

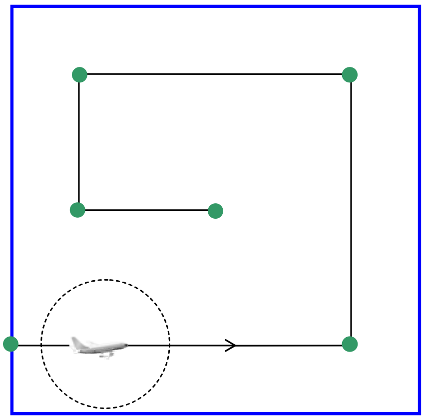

Figure 2. A inwards-spiral pattern

This paper tests some alternative initial waypoint positions involving spiral patterns, both an "inwards" spiral and an "outwards" spiral, as illustrated in Figure 2. The premise for the inwards spiral is to effectively "bound" the AI initially and "trap" the ships within it, giving less chance for a ship that starts the mission within the AI to escape and be undetected or unclassified. The outwards spiral is the reverse, with a high probability of classifying all targets near the centre of the AI, and a good probability of classifying targets that enter the AI as the mission progresses.

2.2. Default heading

When an aircraft commences its mission, if no ships are immediately detected, it heads along the wayline to the next waypoint. When it detects a ship, it diverts from the wayline to classify it. Once classified, if the aircraft again does not detect any ships, the aircraft's default heading is directly to the next waypoint.

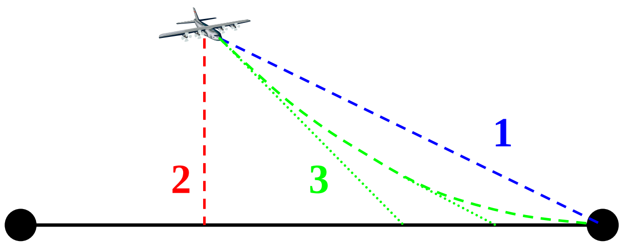

Figure 3. Examples of various default heading options when no contacts are visible: 1) direct to the next waypoint; 2) perpendicular to the wayline; and 3) to the midpoint of 1) and 2), which is constantly updated as the aircraft moves.

This default heading will reduce the travel time of the aircraft, but may also reduce classifications. An aircraft heading directly to the next waypoint could miss a significant section of the AI which may contain undetected ships. This may be a particular issue in low-density environments. In a higher-density environment, aircrew will divert the aircraft more regularly, thus expecting more detections and a greater coverage area.

In this work, two other options are considered for this default heading:

- A perpendicular return. In this instance, the aircraft heads to the nearest direct intercept point with the wayline. This will increase tour length, but also the likelihood that more ships will be detected.
- A "midway" return. As the name suggests, this heading is a compromise between the direct-to-waypoint heading used thus far and the perpendicular return described above.

These options are shown in Figure 3. Note that including the aircraft's turning circle may adversely affect tour length depending on aircraft type (Mercer et al, 2008). The greatest effect would be in the perpendicular case: an aircraft heading away from a wayline to classify a ship would require an obtuse-angle turn to return perpendicularly to the wayline.

2.3. Ghost ships

In a previous paper (Marlow et al, 2007), the percentage of ships classified was found to increase with ship density up to a certain critical point and then decrease beyond that, due to the aircraft having too many ships to classify in the available time. Before this point, ship classifications should be close to 100%. In these instances, the aircraft has enough time to sufficiently cover the area and divert from the wayline to classify the ships it detects. This is evident by the fact that the mission time is always below the maximum time available, meaning that the aircraft finishes sweeping the area with time to spare.

One reason that classifications are below 100% in these circumstances is that in diverting the aircraft from the wayline to classify a detected ship, the aircrew may miss other parts of the AI that contain undetected ships. Once it returns to the wayline, it may have flown past this unexplored area.

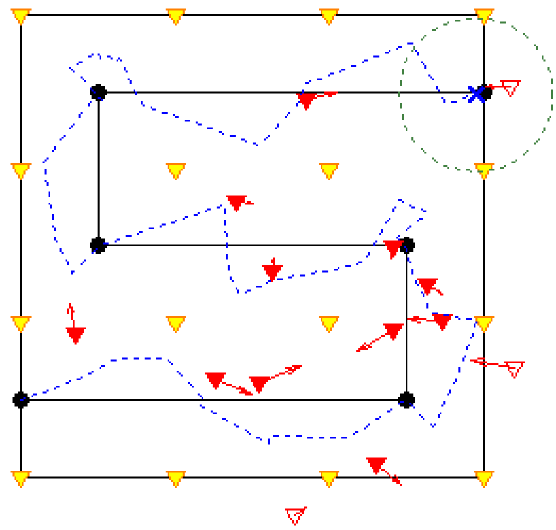

Figure 4. Sample flight path (blue dashed line) when ghost ships (yellow triangles) are placed in an AI. Here they are placed in the "step" pattern.

The concept of creating "ghost ships" is to reduce the chance of these circumstances arising. Ghost ships are stationary and placed at known locations in the AI. For each segment, the positions of the ghost ships are included in the current tour, forcing the aircraft to fly towards them. Once "detected" by the aircraft's sensors, a ghost ship is removed from the tour. Including ghost ships will increase the distance travelled by the aircraft in seeking to classify extra ships. They are also more computationally expensive, as they are included in the tour and thus considered by the tour improvement algorithm.

Two ghost-ship patterns are included in this model for comparison, with ghost ships placed on the AI and segment boundaries. Figure 4 shows the ghost ship placement and sample run for the "step" pattern. The other is a "zigzag" pattern, which encourages a more wave-like path through each segment. In both cases the ghost ships are spaced at distances corresponding to the diameter of the aircraft's radar detection range.

3. MODEL DESCRIPTION

The model considers a square AI, 300 n mile by 300 n mile, with aircraft speed of 300 kn and detection range of 50 n mile. (Knots and nautical miles are conventional in military aviation and marine navigation.) The maximum time is 8 hrs: if this time is exceeded and the aircraft has not covered the entire AI, the model stops and no more ships may be classified. Classification range is 5 n mile. Ship headings are randomly sampled from a uniform distribution over 360°. Ship speeds are randomly sampled from a uniform distribution between zero and a maximum speed set as a model parameter. Only ships within the AI may be classified.

Given these and other parameters, the "critical ship density" beyond which classifications cannot physically reach 100% can be estimated (Bocquet, 2008) at around 8-9 ships per 10,000 n mile^2. This is because TSP heuristics give tour lengths proportional to the square root of the ship density (Supowit et al, 1983).

The model makes some simplifying assumptions. A "cookie-cutter" radar model is used, whereby a ship is detected if it is within the radar range, and undetected if it is not. All ships are assumed to be of the same type for purposes of radar detection range: in reality a fishing boat would have a substantially smaller radar signature than a large merchant vessel, for example. The model is 2-D, with no consideration given to variations in aircraft altitude. The speed and heading of each ship remains constant during a simulation run.

The maritime surveillance model is coded in MATLAB® with a one minute time step. The tour improvement technique is a genetic algorithm (Holland, 1975) to solve open TSPs obtained from the MATLAB® Central file exchange. The population size is set to four times the number of ships in the current tour (including ghost ships when chosen) with the number of iterations set to five times that amount. These parameters were found to be a reasonable balance between accuracy and computation time. The GA results matched well with an exhaustive tour search technique at lower densities. A total of 200 runs were completed for each case.

One interesting feature observed during analysis was the model occasionally oscillating between alternate tours. The improvement method chooses a particular route for a given tour: once executed, the new tour may include a new ship detection, or the removal from the tour of a ship now out of detection range. When applying the improvement method to the new situation, it may return a route which, when executed, results in the previous tour arising, and so on. To escape this, a rule was implemented forcing the aircraft to head to the closest ship until oscillations ceased. Allowing the model to "remember" all ships previously detected should overcome this issue and is plausible in low-density environments, but less so for higher densities.

In addition to the tested enhancements, a "nearest wayline" rule is applied throughout constraining the aircraft to only classify ships within the boundaries of the current segment. An exception is a ship just outside those boundaries, with a heading that will cause it to enter the current segment.

4. RESULTS

Results (averages and 95% confidence intervals) are shown for the three options considered in section 2. The "baseline" configuration is an S-shaped search pattern, a default heading of direct-to-waypoint and no ghost ships. Average ship densities vary from around 1-14 ships per 10,000 n mile^2: i.e., ~10-120 ships in the AI.

Two cases for ship speeds are considered – stationary ships; and speeds between 0 and 20 knots. Cases of ship speeds to 10 knots were also run. Note that the primary Measure of Effectiveness (MOE) of percentage of ships classified considers the number of ships classified divided by the average number of ships in the AI over the duration of the mission. Earlier work (Marlow et al, 2007) considered the total number of ships that ever entered the AI, so MOE values in this paper will be higher. The average number that enter the AI per hour can be calculated (Koopman, 1980): for the values considered here, it is 0.17ρ (to 10 kn) and 0.34ρ (to 20 kn) where ρ is the ship density per 10,000 n mile^2.

4.1. Waypoint variation

Figure 5 shows the average percentage of ships classified against the average ship density in the AI for the different waypoint patterns. The results for all three search patterns are effectively identical across the considered range of ship densities, with the largest difference at very low densities for the 20 knot case. At lower densities, the aircraft still completes the search of the AI with considerable time to spare.

Each case tested theoretically has the aircraft covering an area of the AI only once. Extra runs were performed where the aircraft could overlap areas previously covered, e.g., using extra spirals or a double-S shape. These gave classifications that were up to several percentage points better than those in Figure 5 at lower densities (as the aircraft spent more time covering a sparse AI), but were worse beyond the critical density (as the aircraft did not have enough time to cover a dense AI).

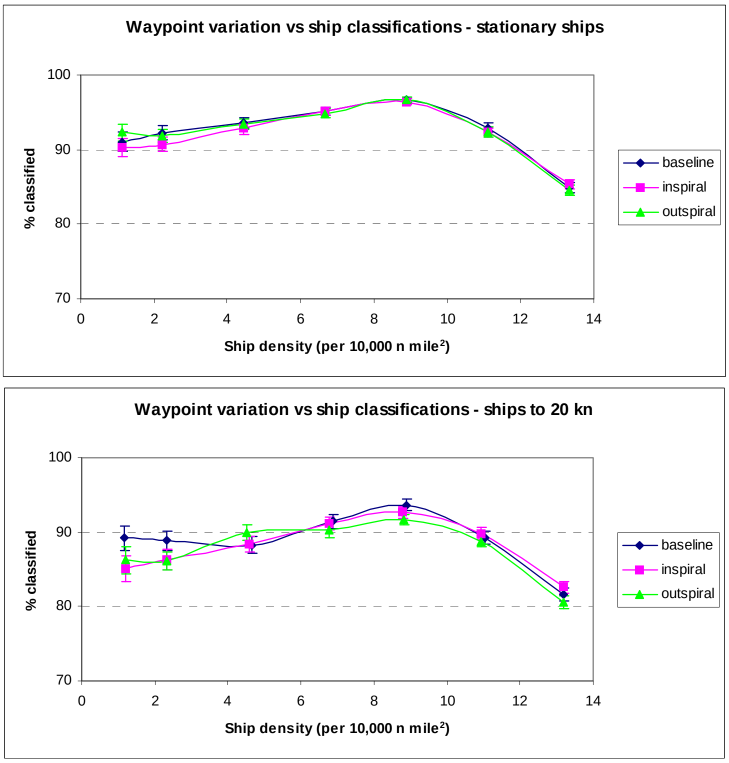

Figure 5. Ships classified versus density for various initial search paths: (top) stationary ships; (bottom) ship speeds to 20 kn

4.2. Changing the default heading

Figure 6 shows the results for ships classified against ship density for various default headings. The perpendicular return-to-wayline rule performs the best of the three at lower densities, as expected, as this method enables the aircraft to cover a greater area of the AI. The largest difference is seen in the case where ship speeds can range up to 20 knots. The percentage difference in classifications between the perpendicular return-to-wayline and the default direct-to-waypoint heading is around 5.5 percentage points at the largest point (at ship density around 4.5). This equates to around 1.5 extra ships on average: the average extra time taken to achieve these additional classifications slightly exceeds half an hour.

It is also instructive to note the number of times that 100% classifications are achieved in the stationary ships case. The perpendicular case achieves 100% classifications 65% of the time with 10 ships in the AI (ship density 1.1) and 48% of the time for the midpoint case, compared with 39% of the time for the baseline heading. For 40 ships (density 4.4), 100% classifications are achieved in 41% of runs (perpendicular), compared with 16% (midpoint) and 13% (baseline).

In all graphs it is evident that once the critical ship density is reached, the perpendicular return method is inferior to the others. The advantage gained by the perpendicular return in using the extra time to search the AI is a disadvantage beyond this critical point, as it allows less coverage of the AI in the available time.

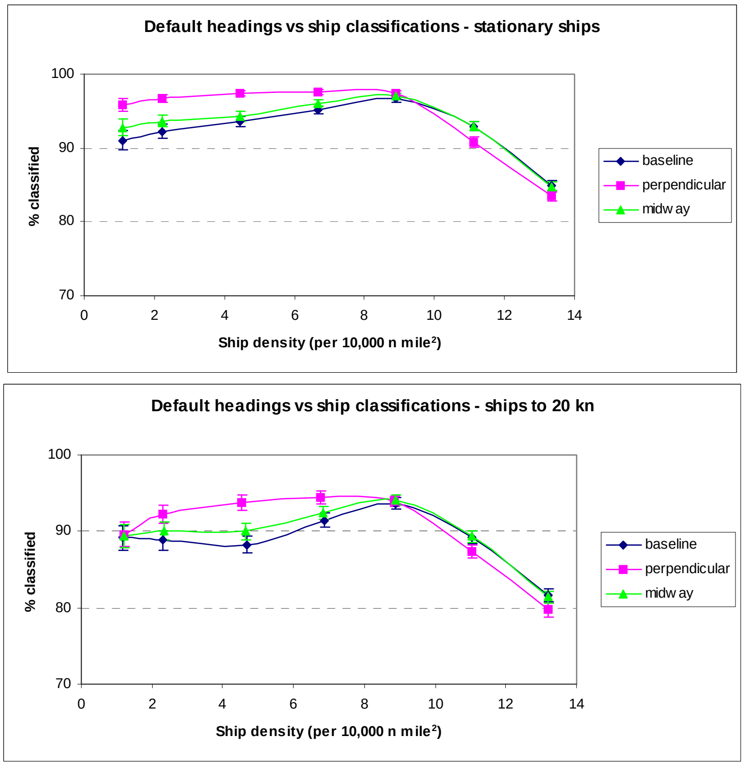

Figure 6. Ships classified versus density for various default headings: (top) stationary ships; (bottom) ships to 20 knots

4.3. Including ghost ships

Figure 7 show the results for the two ghost ship patterns against the baseline results. At lower densities, the presence of the ghost ships succeeds in drawing the modelled aircraft to classify a larger percentage of ships. As for the default headings, this equates to around 1-2 ships more on average at lower densities.

Again, it is worth noting the improvement in achieving 100% classifications by including ghost ships. For stationary ships, a classification rate of 100% is achieved 62% of the time for 10 ships in the AI (ship density of 1.1 per 10,000 n mile^2) for both ghost ship patterns, while it is achieved only 39% of the time without ghost ships. For 20 ships (density 2.2), the 100% level is achieved 50% of the time (ghosts) compared with 23% (no ghosts), and for 60 ships (density 6.7), it is 31% (ghosts) compared with 10% (no ghosts). Including more ghost ships in appropriate locations would probably increase classifications further, but at the cost of increased computational time.

At higher densities, using ghost ships is ineffective, and in fact counter-productive to the objective. Ghost ships are superfluous at such densities, as the AI is already sufficiently dense to enable the

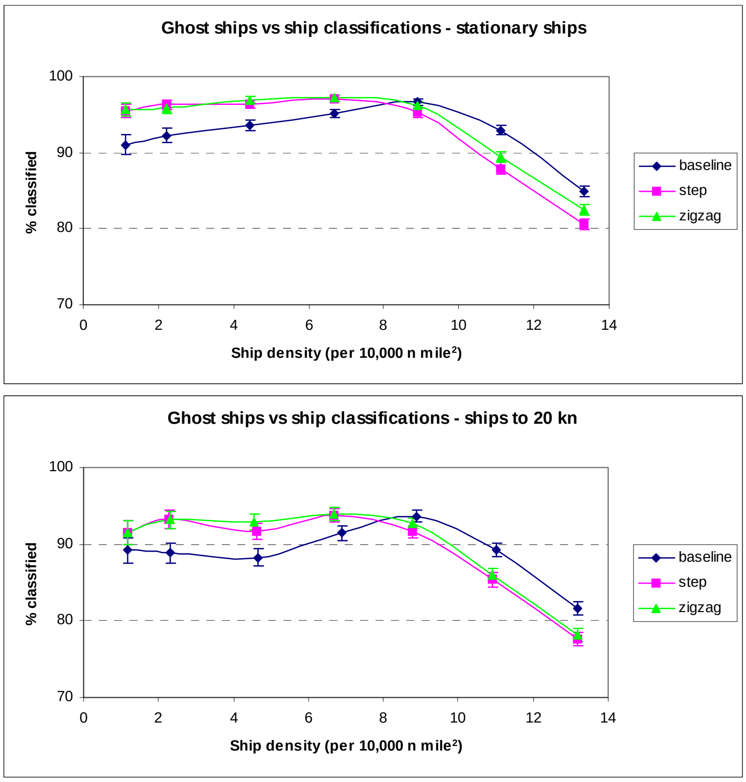

Figure 7. Ships classified versus density for ghost ship patterns: (top) stationary ships; (bottom) ships to 20 knots

aircraft to regularly divert to classify ships. The extra diversions only add to the length of the mission and thus negatively affect the ability of the aircraft to traverse the AI in the time available, resulting in less coverage and therefore fewer classifications. Including ghost ships in the AI in these instances also adds to computation time, due to their inclusion in the tour until their position is covered by the aircraft's radar. Computational runs on average take over twice as long as for the baseline case at these higher densities.

5. DISCUSSION AND CONCLUSIONS

The perpendicular default heading and the inclusion of ghost ships have been successful in improving ship classifications at lower densities, and in increasing the frequency of achieving 100% classifications. At higher densities, results indicate that the baseline pattern is most suitable, as the variations generally fail for the same reasons that they succeed in the lower-density cases – i.e., they take more time to search the AI.

Tests were employed to run the perpendicular heading and ghost ships options together, but no further improvement was found – in fact the results were very similar to those for the ghost ships options only. This is probably due to the same basic principle behind the two methods, requiring the aircraft to divert either from the ship to the wayline (for perpendicular return) or from one part of a segment to a "ship" (for ghosts).

It is proposed to extend this work to examine methods for maximising classifications for densities beyond the critical density. One possible method is the implementation of "budgeting" routines. Given the constraint requiring the aircraft to cover the entire AI in the available time, such routines would enable aircrew to choose which ships to visit, and which to omit from the tour, in order to maximise classifications.

ACKNOWLEDGMENTS

The authors would like to thank Simon Goss, Jason Looker and Stephen Bocquet (DSTO) for their input and assistance, Joseph Kirk for providing the open TSP genetic algorithm function on the MATLAB® central file exchange, and Li Zhou for initial investigations into the effectiveness of spiral waypoint patterns. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

REFERENCES

Bocquet, S., (2008), Empirical model of a search and classify mission, private correspondence.

Croes, G.A. (1958), A method for solving traveling-salesman problems, *Operations Research*, 6, 791-812.

Feillet, D., P. Dejax and M. Gendreau (2005), Traveling salesman problem with profits, *Transportation Science*, 29, 188-205.

Grob, M.J.H.B. (2006), Routing of platforms in a maritime surface surveillance operation, *European Journal of Operational Research*, 170, 613-628.

Holland, J.H. (1975), *Adaptation in Natural and Artificial Systems*, Univ. of Michigan Press, Ann Arbor.

John, M., D. Panton and K. White (2001), Mission planning for regional surveillance, *Annals of Operations Research*, 108, 157-173.

Koopman, B.O. (1980), *Search and Screening*, Pergamon Press, New York.

Looker, J.R. (2008), Minimum paths to interception of a moving target when constrained by turning radius, DSTO-TR-2227, Defence Science and Technology Organisation, Commonwealth of Australia.

Marlow, D.O., P. Kilby and G.N. Mercer (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.

Mercer, G.N., S.I. Barry, D.O. Marlow and P. Kilby (2008), Investigating the effect of detection and classification range and aircraft dynamics on a simplified maritime surveillance scenario, *ANZIAM J.*, 49, C475-C492, *Proceedings of the 8th Biennial EMAC 2007 Conference*, G.N. Mercer and A.J. Roberts (eds.)

Ng, K.Y.K. and A. Ghanmi (2002), An automated surface surveillance system, *Journal of the Operational Research Society*, 53, 697-708.

Ng, K.Y.K. and N.G.F. Sancho (2009), Regional surveillance of disjoint rectangles: a travelling salesman formulation, *Journal of the Operational Research Society*, 60, 215-220.

St-Hilaire, M-O., E. Lefebvre and C. Helleur (2008), Track modeling for maritime surveillance, 11th International Conference on Information Fusion, Cologne, Germany, 30 June-3 July 2008, 1-8.

Supowit, K.J., E.M. Reingold and D.A. Plaisted (1983), The traveling salesman problem and minimum matching in the unit square, *SIAM Journal of Computing*, 12, 144-156.

Wagner, D.H., W.C. Mylander and T.J. Sanders (editors) (1999), *Naval Operational Analysis, Third Edition*, Naval Institute Press, Annapolis.