

**Some Consequences of a Computer Model to Simulate
the Performance of a Land SAR Searcher**

by Dave Perkins

The Centre for Search Research, Ashington, Northumberland, UK

July 2011

Abstract

This report discusses the outcomes of a study that used a computer simulation model to investigate the distances at which a searcher can detect a target in different environments.

The model works by placing targets in random locations in an empty search area. It simulates a variety of different types of terrain by randomly distributing obstacles throughout the search area; an obstacle is something that obstructs the searcher's line of sight, and is the model's version of those objects in the real world that restrict the searcher's ability to detect a target, for example vegetation or uneven terrain.

The model was provided with no guidance as to whether or not a particular target would be detected, other than the simple rule that said that if there were no obstacles between the searcher and a target then the searcher would detect it.

The model produced lateral range curves, and from them was able to calculate effective sweep width. Since the model was able to determine critical separation for each type of terrain, the relationship between sweep width and critical separation could be explored.

Other results from the model were used to produce a coverage / POD curve for a field team, and this led to a new version of the graph showing how POD varied with searcher spacing.

Another outcome was an analysis of long range detections, i.e. the detection of targets at greater distances than might otherwise have been expected. The model was able to quantify this phenomenon (target detections at distances in excess of three times critical separation were not uncommon) and offer an explanation.

The model was able to investigate the outcome of searchers using purposeful wandering, and demonstrated the effect of purposeful wandering on effective sweep width.

The report concludes with a comparison of the results from this study with those from the two previous related studies.

Contents

Introduction

Obstacles, detections and targets page 1

The structure of the simulation model

The search area page 2

Obstacles and obstacle density page 2

How it works page 3

Running the model

The size of the search area page 5

First results page 5

Lateral range curves and effective sweep widths page 6

Detections, non-detections and the crossover curve page 7

Effective sweep width and obstacle density page 8

Critical distance and critical separation page 10

Coverage, percentage detections and the coverage / POD curve page 12

The two searcher model page 13

Long range detections page 16

Purposeful wandering page 17

Comparison with previous work page 19

Limitations – what the model doesn't do page 20

Conclusions page 20

References page 21

Acknowledgements page 21

Figures and tables in the text

Fig. 1: the search area, 1,000 rows x 100 columns	page 2
Fig. 2: the search area, with targets T and obstacles O	page 4
Fig. 3: graph of percentage of targets detected against distance from the searcher for an obstacle density of 5%	page 6
Fig. 4: lateral range curves for the standard set of experimental obstacle densities	page 7
Fig. 5: crossover curves (detection and non-detection) for 5% obstacle density	page 8
Fig. 6: graph of ESW against obstacle density, and the hyperbola $xy = 103$	page 9
Fig. 7: graph of n against percentage obstacle density for the standard set of experimental obstacle densities	page 11
Fig. 8: graph of coverage against percentage detection from the model, with the random search curve for comparison	page 13
Fig. 9: the model with two searchers, S1 and S2	page 14
Fig. 10: graph of coverage against percentage detection from the two - searcher model, with the random search curve for comparison	page 15
Fig. 11: graph of % POD against searcher spacing measured in multiples of critical separation, with fixed spacing and no purposeful wandering	page 15
Table 1: effective sweep widths in metres calculated as area under the lateral range curve and crossover distance for the standard set of experimental obstacle densities	page 8
Table 2: $n = \text{ESW} / \text{CS}$ for the standard set of experimental obstacle densities	page 11
Table 3: long range detections for the standard set of experimental obstacle densities	page 16
Table 4: sweep widths for the standard set of experimental obstacle densities with and without purposeful wandering (PW)	page 18

Introduction

Obstacles, detections and targets

What is it that allows us to see an object in one environment but not in another? Provided that there is sufficient daylight, we would expect to see a person walking on a beach from a considerable distance, whereas it might be difficult to see the same person dressed in the same clothing and walking in the woods, even if we were close to them. What is the difference? One answer is that a beach is usually flat and level with not a lot to obstruct our view, but in the woods there are things that can get in the way.

There are other reasons why we might not see them; some of these will be discussed later.

The “things that can get in the way” have a central role in this study. We will refer to them as **obstacles**. As far as this study is concerned, it is the number of obstacles that distinguishes one type of terrain from another; the example of the beach and the woods illustrates this. The fundamental property of an obstacle is that it prevents you from seeing something which is on the opposite side of it from where you are. Clearly scale is important; a pebble would not prevent you from seeing an unconscious person lying on the ground, but a large boulder would. The pebble, however, might prevent you from seeing a clue.

There are many different obstacles in the natural environment. These could be trees, bushes, long grass, the bumps and hollows of uneven terrain, or rocks and boulders; if we add to this list all the man-made obstacles we might encounter in an urban environment then it is clear that the real world – the world in which we search – is full of things that can get in the way and prevent us from seeing objects that we might be searching for.

From now on we will use the word ‘**detecting**’ instead of ‘seeing’; ‘seeing’ carries with it a suggestion that for some reason the person doing the seeing might not be aware of what it is that they have seen, whereas ‘detecting’ implies that the person is aware of the object and possibly recognises it. In the computer simulation model that this report describes this distinction does not exist; under the right circumstances, the ‘searcher’ in the model will detect a target – it is not an option for them to see it but not be aware of it.

In this study we will imagine that we are searching for an object (referred to from now on as a **target**) which, with no obstacles in the way, is just as likely to be seen from a long distance, say 500 metres, as it is from 50 metres. In other words, as far as our study is concerned, the likelihood of detecting the target will depend only on whether or not there are any obstacles in the way, and not on the distance from the searcher to the target. This is a valid assumption over the range of distances and size of targets that we will consider, and although at first sight this may seem unrealistic, consider again the example of the likely distances at which we can detect a person on a beach or in the woods, and the difference between those two types of terrain.

The structure of the simulation model

The search area

Imagine a rectangular search area, 1km in length and 100 metres wide; that gives us a grid with 1000 x 100 squares of 1 metre x 1 metre arranged in 1000 rows each containing 100 grid squares (fig. 1). In the model, a searcher, **S**, is going to walk along the edge of this search area (using column 0) and look to their right along each of the 1000 rows in turn, recording each target that they see. Detections will therefore be made in a direction which is perpendicular to the searcher's direction of travel.

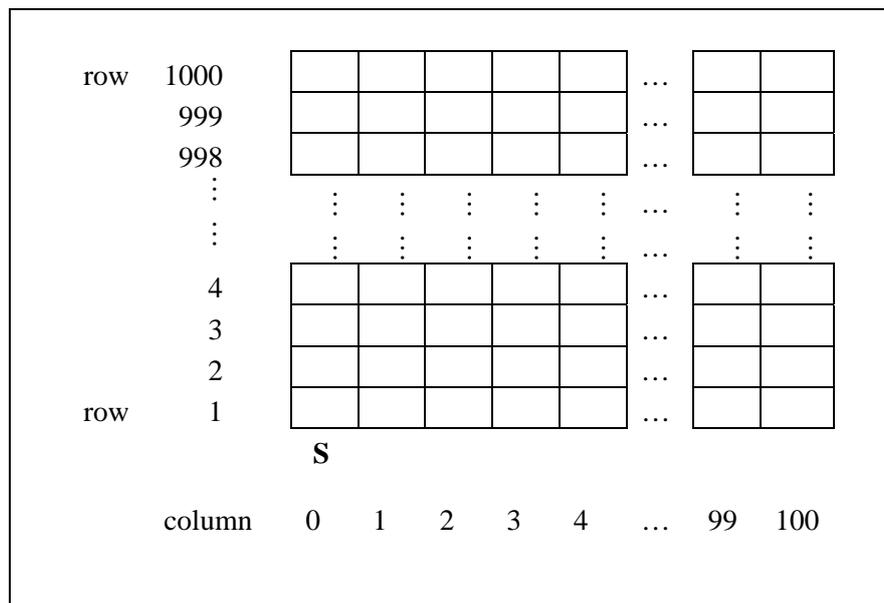


Fig. 1: the search area, 1000 rows x 100 columns

When the searcher reaches the end of the search area they will have a record of the number of targets detected.

Distances throughout the paper are reported in metres; this is somewhat arbitrary, and reflects only the author's personal preference. They could just as well have been reported in feet or yards, as long as their use is consistent.

Obstacles and obstacle density

We saw that the difference between how successful we are likely to be at detecting a person on a beach compared with detecting the same person in the woods was due to the relative number of obstacles in each of these environments. This helps us clarify exactly what we

mean by an obstacle: an obstacle is something which, when placed between the searcher and the target, obscures the target and prevents the searcher from detecting it. If we are searching for a missing person in the woods then a bush or large clump of vegetation would constitute an obstacle. In practice we would need to consider the orientation of the target – was the person standing or lying down – but for the purpose of our model that is not significant. In practice we would also need to consider the height of the searcher, and whether or not they would see over the obstacle and still detect the target. Again, this is not part of the model – by our definition of what constitutes an obstacle, if the searcher can see over it then it is not an obstacle.

There are other issues that, for the moment, we are not considering: for example, a person in the woods wearing a high visibility orange coverall is likely to be easier to detect than a person wearing a green jacket; this will be discussed towards the end of the paper. The effects of searcher fatigue or expectation are related to the searcher and not the target or environment, and are therefore outside the scope of the current model.

The difference between the beach and the woods is really down to the number of obstacles that you would encounter in each of them. On the beach there would be very few; in the woods there would be many. We can interpret this as the number of obstacles that you would expect to find in the same sized portion of each of them; we will refer to this as **obstacle density**, and in our model we will measure it as a percentage. An obstacle density of 5%, for example, means that on average there are five obstacles in every hundred grid squares of the search area in fig. 1.

A simple example will help to visualise all this. Our search area is analogous to the rows and columns of squares on a chess board; our castle can take an opponent's piece on the same row provided that there is nothing in the way ... the searcher can detect the target provided there are no obstacles in the way. Notice the units and scale involved: the pieces on a chessboard just about fill the squares; in our model the targets and obstacles can be thought of as just about filling the grid squares. If the targets and obstacles change in size (for example if we are considering clues) then the units will change to reflect that, and instead of metre grid squares we may have grid squares of 10 cm by 10 cm.

How it works

The model begins by placing **one target in a random position in each row**. This notion of one target on every row is an important feature of the calculations that take place in the model; the search area in fig. 2 has 1,000 rows and therefore contains 1,000 targets. We say that this search area is providing the searcher with 1,000 **detection opportunities**. In fig. 2, targets are denoted by **T**.

Suppose that we are investigating the outcome of searching an area with an obstacle density of 5%. The search area in fig. 2 consists of 100,000 grid squares, and therefore the model

will **randomly distribute 5,000 obstacles (5% of 100,000) throughout the search area.** Only one object (target or obstacle) can occupy a grid square, so if the random location selected for an obstacle already contains a target or another obstacle then another random location is chosen. In our example, there will be 5 obstacles per row on average, but since they are randomly distributed some rows will have less than 5 and some will have more. This is similar to our view of the real world: in general, obstacles are not uniformly distributed throughout our field of view. In fig. 2, obstacles are denoted by **O**.




Fig. 2: the search area, with targets **T** and obstacles **O**

The model then carries out the following procedure as the searcher **S** travels up column 0 from row 1 to row 1000 (fig. 2):

row 1: there are no obstacles between the searcher and the target, and so the target is detected; the model adds 1 to the count of targets detected so far, and records the fact that the detected target was in column 3

row 2: the obstacle in column 1 prevents the searcher from seeing the target; target not detected

row 3: target not detected; the obstacle in column 2 is in the way

row 4: target detected; the model adds 1 to the count of targets detected (two so far) and records the fact that this target was located in column 1

and so on

By the time that the searcher reaches row 1000 we will have a count of the total number of targets detected out of 1000, and a record of the number of detections at distances of 1 metre, 2 metres, 3 metres and so on up to 100 metres, assuming that the grid in fig. 2 consists of metre squares.

Running the model

The size of the search area

In practice, the dimensions of the search area used in model were very different from the 1,000 x 100 grid discussed so far. As we have seen, the model places a target on every row, and therefore there are as many detection opportunities as there are rows. We wanted as many detection opportunities as could reasonably be accommodated within our model, since this would give stability to the results; after all, these are random trials, with randomly placed targets and obstacles. If we had only a small number of rows and columns then we would expect to get a different result each time we ran the model; if we roll a dice three times it does not tell us much about the dice, whereas if we roll it three thousand times then it does. Therefore we want to simulate what happens with a large number of detection opportunities to even out any random variations.

The model was constructed with 250,000 rows, which equates to a search area that is 250 km long. **One run of the model gave 250,000 detection opportunities**, which is equivalent to 6,250 searchers participating in a field trial in which there are 40 targets.¹

Early results from the model suggested that targets were often detected at greater distances from the searcher than had been anticipated. In order to avoid restricting the range of possible detections, the model was given 500 columns. Our search area therefore consists of a 250,000 x 500 grid. Remember that we said earlier that our target was as likely to be seen at a distance of 500 metres as it was at a distance of 50 metres provided that there were no obstacles in the way.

First results

The search area contains 250,000 rows and therefore 250,000 targets, with one target placed in a random position on each row. The search area contains 500 columns, and therefore on average each column contains 500 targets, with some having more and others less than 500. The model will report the percentage of targets detected in each column.

Fig. 3 shows how the percentage of detections varied as the distance from the searcher increased for an obstacle density of 5% for one run of the model. Remember that one run of the model provides 250,000 detection opportunities. Fig. 3 shows that the percentage of detections is high for targets that are close to the searcher, but as the distance from the searcher increases it falls away rapidly for the first 10 metres or so; thereafter the rate of fall reduces, so that for targets at around 45 metres from the searcher the percentage of detections is down to around 10%, with a shallow decline after that; the furthest detection distance on this particular run was of a target at 178 metres from the searcher.

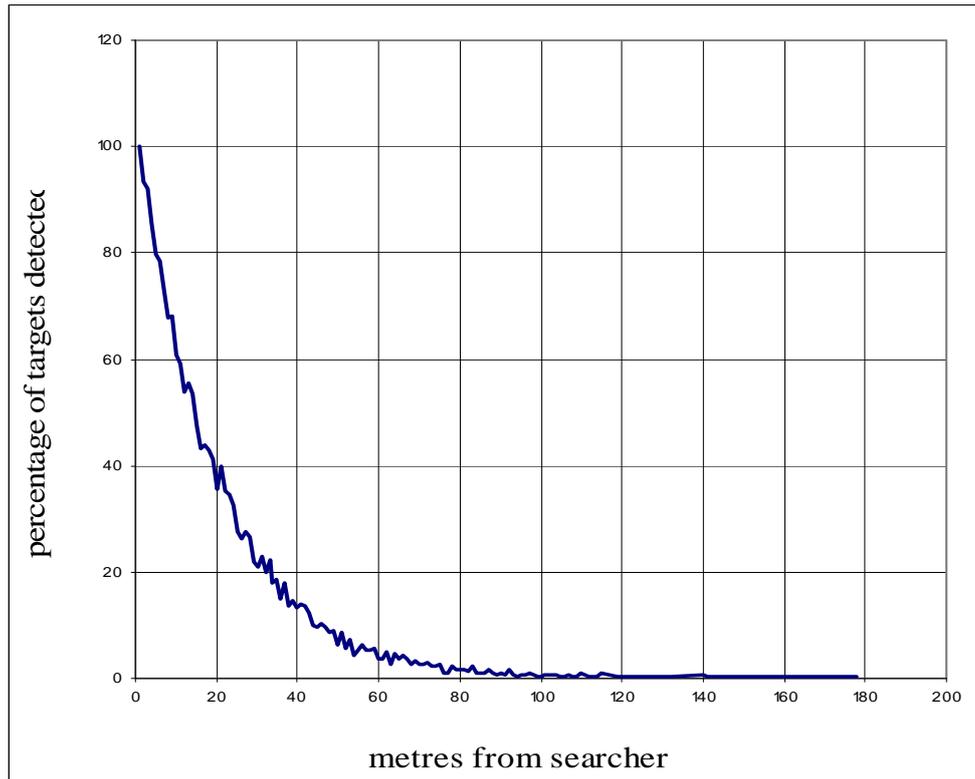


Fig.3: graph of percentage of targets detected against distance from the searcher for an obstacle density of 5%

Lateral range curves and effective sweep widths

All detections were made in a direction perpendicular to the searcher's line of travel, and therefore fig. 3 shows how the probability of detecting a target varies as the target's lateral distance from the searcher varies. Fig. 3 is therefore the lateral range curve for an obstacle density of 5%.

The effective sweep width (ESW) is defined as the area under the lateral range curve. The model used the trapezium (trapezoid) rule to calculate the area under this curve, and obtained a value of 20.4 metres.

Additional runs were carried out for obstacle densities of 2.5%, 7.5%, 10%, 12.5% and 15%; from now on these values, together with 5%, will be referred to as 'the standard set of experimental obstacle densities'. The resulting lateral range curves are shown in fig. 4. In each case the model calculated ESW as the area under the lateral range curve; the results are shown in the first two columns of table 1.

Fig. 4 shows that as the obstacle density increases the percentage of targets detected at a given distance decreases; for example, with an obstacle density of 2.5%, more than 60% of the targets at a distance of 20 metres are detected, whereas with an obstacle density of 10%

only 15% of the targets at that distance are detected. This is as you would expect; an increase in the number of obstacles makes it less likely that the searcher will detect targets at a given distance.

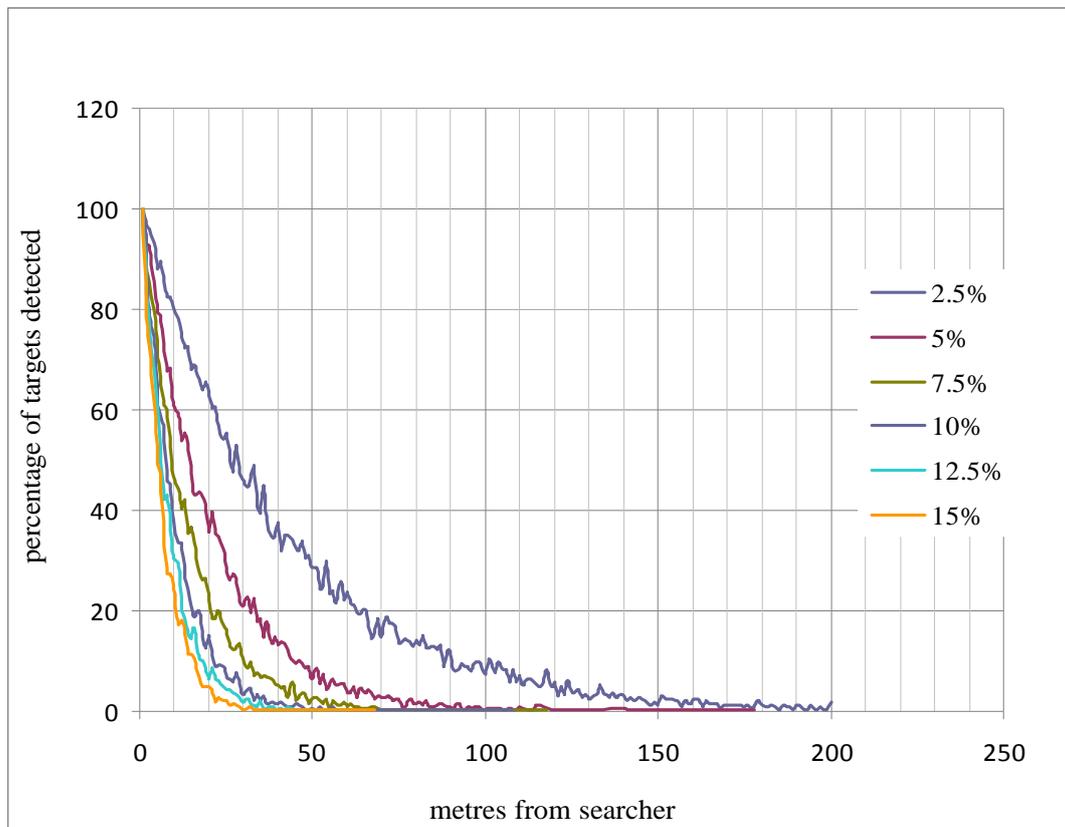


Fig. 4: lateral range curves for the standard set of experimental obstacle densities

Detections, non-detections and the crossover curve

The model counts the number of detections at each lateral distance, and since it knows how many targets there are at each of those distances, it can therefore calculate the number of non-detections at each lateral distance. From these it can calculate the cumulative number of detections and non-detections; when these are plotted as two separate curves on a graph, the distance from the searcher at which the lines cross is ESW.

Fig. 5 shows the detection and non-detection curves for an obstacle density of 5%. Note that in practice the curves extend to a distance of 500 metres from the searcher; extreme values have been excluded to improve the clarity, but their omission does not affect the crossover point. The crossover calculation of ESW was built into the model; the values for the standard set of experimental obstacle densities are shown in column 3 of table 1. Notice how similar the two values of ESW in columns 2 and 3 of table 1 are for each obstacle density.




Fig. 5: crossover curves (detections and non-detections) for 5% obstacle density

ESW and obstacle density

Table 1 shows that as the obstacle density increases then the ESW decreases. This is not unreasonable; an increase in obstacle density means that there are likely to be more obstacles on each row, and therefore an increased likelihood of an obstacle coming between the searcher and the target on that row. In the real world we would say that the vegetation is becoming denser, or the trees closer together, or the ground more undulating, and it therefore gets more difficult for the searcher to detect targets.

x		y		
percentage obstacle density	ESW as area under LRC (metres)	ESW as crossover distance (metres)	mean of two ESW values	xy
2.5	40.3	40.3	40.3	101
5	20.4	20.1	20.3	101
7.5	13.8	13.6	13.7	103
10	10.4	10.2	10.3	103
12.5	8.6	8.5	8.6	107
15	7.1	6.9	7.0	105

Table 1: effective sweep widths in metres calculated as area under lateral range curve and crossover distance for the standard set of experimental obstacle densities

Fig. 6 shows the mean of the two ESW values in table 1 (labelled y in table 1) plotted as a smoothed curve against obstacle density (labelled x in table 1). This curve is a hyperbola; hyperbolae have the form

$$xy = \text{constant}$$

where, in our case, x is the obstacle density and y is the sweep width. The values of the product xy are shown in the final column of table 1; the average of these values is 103, and therefore over the range of values in table 1, the equation of the hyperbola can be written as

$$xy = 103$$

This graph is also shown on fig. 6. The two curves in fig. 6 are almost identical; the reason for this relationship is not clear at the time of writing.

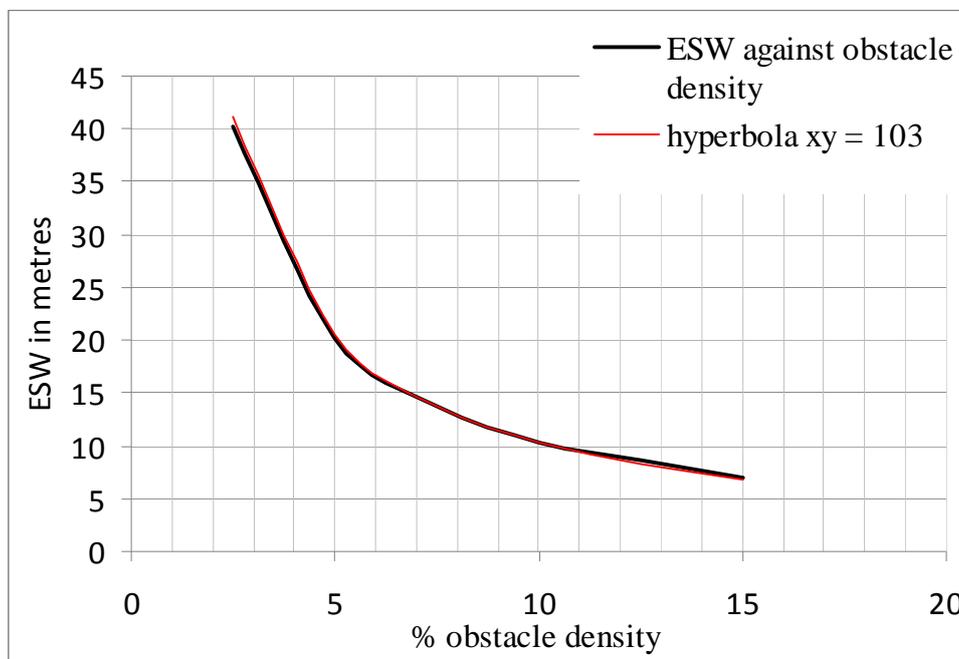


Fig. 6: graph of ESW against obstacle density, and the hyperbola $xy = 103$

Table 1 and fig. 6 show that when the obstacle density doubles then ESW is halved. For example when the obstacle density was 2.5%, ESW was found to be 40 metres; doubling the obstacle density gives 5%, which has an ESW of 20 metres, and doubling it again gives 10%, which has an ESW of 10 metres. The hyperbola supports this. If this relationship holds, it would be possible to determine intermediate values of ESW (for, say, an obstacle density of 11%) and to extend the curve in fig. 6 at either end to give the ESW for an obstacle density of say 20%.

Critical distance and critical separation

For each obstacle density, the model calculates and remembers the distance at which each target is detected, and can therefore calculate the mean of all of these distances. This is an interesting statistic, and is worth considering in some detail, but first we will remind ourselves about critical separation.

Critical separation² (CS) is defined as the distance between two searchers when an object that represents the item (person or clue) for which they are searching, placed midway between them, is on the limit of visibility of both searchers. In the field, the usual procedure is for the searchers to place an appropriate object on the ground and walk away from it in opposite directions, turning every now and then to look back at it to determine the first point at which they lose sight of it; they then move back towards the object until it reappears. They then circle round the object, increasing or decreasing their distance from it to keep it on the limit of visibility. We say that CS for these two searchers, for the object that they have used and in the terrain in which they have carried out this procedure, is the average distance between them as they circled the object. We say that circling the object and moving in and out to keep it just in sight gives a subjective averaging of the terrain, and allows for any obstacles that might come between the searcher and the object.

The concept of critical distance has been discussed elsewhere.³ It is the distance from one searcher to the object in the field procedure described above, and is therefore half CS. Critical distance is therefore a subjective measure of the average distance at which a searcher can detect a target in the type of terrain in which they are searching; in other words, it is the mean detection distance for a target in an environment with a particular obstacle density.

But this is the statistic that our model calculates for each obstacle density; therefore if we double this number we will have a value for CS.

There has been some conjecture over the relationship between ESW and CS; this has been expressed as

$$ESW = n \times CS \quad \text{for some value of } n \quad (1)$$

The model provides us with an opportunity to explore this relationship. Table 2 is a copy of the first three columns of table 1, with three additional columns labelled a, b and c. The additional columns show the following:

- a. the mean distance to the targets detected for that obstacle density, as calculated by the model
- b. CS, calculated as twice the distance in (a)
- c. n calculated as ESW / CS ; the value for ESW used here is the average of the two values given in table 2, and is equal to the value in column y in table 1

The values for n appear to increase as the obstacle density increases; fig. 7 bears this out, although the reason for this, or for the dip in the middle of the graph, is not clear at the time of writing. The mean of the six values for n in table 2 is 0.518.

percentage obstacle density	ESW = area under LRC (metres)	ESW = crossover distance (metres)	a	b	c
			mean distance to detected targets (metres)	CS (metres)	n
2.5	40.3	40.3	40.2	80.4	0.501
5	20.4	20.1	19.7	39.4	0.514
7.5	13.8	13.6	13.5	27	0.507
10	10.4	10.2	9.8	19.6	0.526
12.5	8.6	8.5	8.1	16.2	0.528
15	7.1	6.9	6.6	13.2	0.530

Table 2: $n = \text{ESW} / \text{CS}$ for the standard set of experimental obstacle densities



Fig. 7: graph of n against percentage obstacle density for the standard set of experimental obstacle densities

Giving n a value of 0.518 takes us into the realms of spurious accuracy. CS is a field measurement and does not pretend to be accurate to more than $\pm 5\%$ at best; in other words, if an experienced field team reports that they have determined CS to be 20 metres, we would take that as meaning that over the entire sector then the value would most likely lie somewhere between 19 and 21 metres. But that is the nature of CS; it is designed to give a quick and simple procedure for a field team to determine something about the average distance at which they could detect a target. We are not into tape measures or rangefinders.

It is perhaps more sensible to say that $n = 0.52$ over the likely range of obstacle densities, giving

$$ESW = 0.52 \times CS \quad (2)$$

Coverage, percentage detections and the coverage / POD curve

Coverage C is defined as the ratio of the search effort to the area being searched, or

$$C = \frac{ESW \times \text{track length}}{\text{area being searched}} \quad (3)$$

Track length is equal to the number of rows in the search area, since this is the distance that the searcher travels; the size of the search area is number of rows x number of columns. Varying the number of columns will give different values of C . For each of these values of C , the model calculates the number of targets that are detected, and since there is one target per row, the percentage number of detections (% POD) can be calculated as:

$$\% \text{ POD} = \frac{\text{number of detections made}}{\text{number of rows}} \times 100 \% \quad (4)$$

Therefore, for a given obstacle density, it is possible to obtain a set of values of C and POD, and hence plot a C / POD curve. The value of ESW used in the calculation of coverage was the original value found for a search area with 250,000 rows x 500 columns.

Fig. 8 shows the resulting coverage / POD curve for an obstacle density of 5 %, but in fact the values are independent of obstacle density. The random search curve⁴ is shown for comparison.

This is interesting in that the two curves in fig. 8 are similar. One of them is an established relationship from search theory, based on a sensor with a definite detection characteristic that moves in a random fashion through the search area; the other comes out of the model described in this paper.

But on reflection there is a problem here. The problem is that the model has considered only one searcher, and generally speaking searchers operate in field teams rather than on their own. What we should be looking at therefore is how the number of detections in the region between two adjacent searchers varies as their spacing varies; varying the searcher spacing will give different values of coverage, in just the same way that varying the number of columns in the search area gave different values of coverage; therefore we should be able to produce a coverage / POD curve for two searchers, which, by extension, will be the curve for a field team. This required a change to the model.

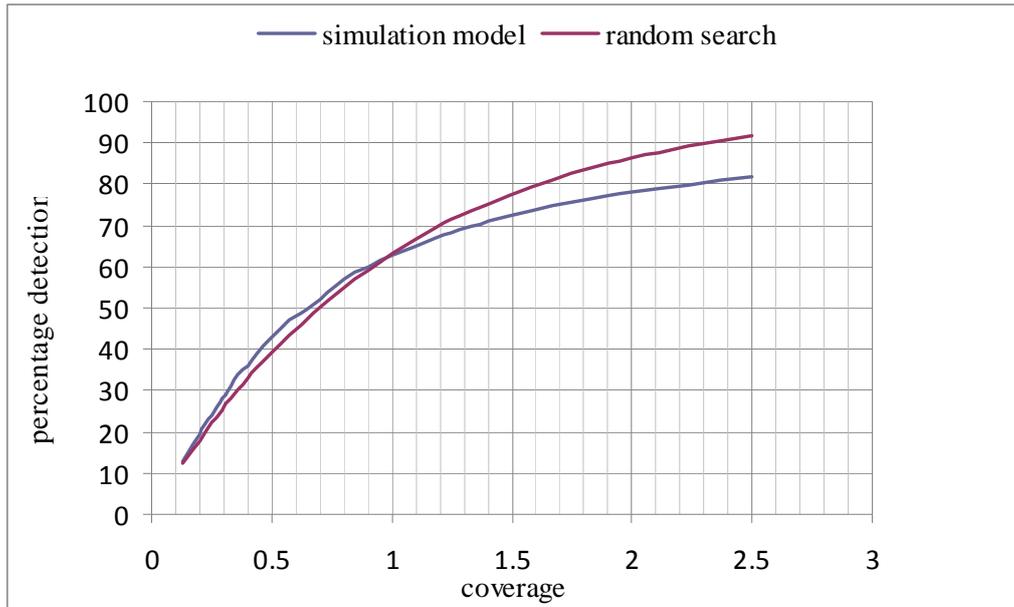


Fig. 8: graph of coverage against percentage detection from the model, with the random search curve for comparison

The two searcher model

A second searcher, **S2**, was added to the model (fig. 9). The second searcher initially travelled up column 101, keeping in line with searcher **S1**, and looking to their left along each row. In the model, searcher **S2** only came into action when searcher **S1** failed to detect the target on a particular row, thus avoiding the possibility of ‘double detections’ on rows with no obstacles, when both searchers were able to detect the target.

On row 2 of the search area shown in fig. 9, when **S1** failed to detect a target, **S2** scanned the row starting with column 100, then column 99 and detected the target. On row 3, the target is hidden by obstacles when viewed by either searcher, and therefore was undetected. When the searchers have finished, the percentage POD could be calculated as in equation (4); coverage was calculated using equation (3), using the value of ESW found for one searcher in a search area with 250,000 rows and 500 columns.

Searcher **S2** then moved one column to their left and the entire process was repeated, including the random placing of targets and obstacles in what was now a slightly smaller search area. This whole routine was repeated until the searchers were five columns apart. This gave a set of results showing how the number of targets detected, and hence POD, varied as the spacing between the searchers decreased.

In a previous analysis,⁵ a third searcher became involved when the searchers were close together. In the current model, there is no need to involve additional searchers because it is not possible to have a situation where a third searcher outside the original two would detect a target that they had missed. Furthermore, searchers **S1** and **S2** would themselves constitute obstacles as far as any additional searchers are concerned.




Fig. 9: the model with two searchers, **S1** and **S2**

The graph of coverage against POD obtained from the two-searcher model is shown in fig. 10, with the random search curve again shown for comparison. As one would expect, now that there are two searchers involved the percentage of detections has increased compared with the curve for a single searcher.

Some members of the land SAR community may not be comfortable with the concept of coverage. Therefore, the POD graph for two searchers shown in fig. 10 has been re-drawn using searcher spacing instead of coverage. Searcher spacing is generally a much more familiar concept to land SAR practitioners, but beware! The advantage of using coverage is that it allows for the fact that the searchers may not have searched the entire area that they were allocated; searcher spacing does not take into account of the size of the search area, or how much of it has been searched, but is concerned only with how far apart the members of the field team are. The POD / searcher spacing curve shown in fig. 11 assumes therefore that the search area is the area swept by the field team and extends from just beyond the searcher at one end of the team to just beyond the searcher at the other end. “Just beyond” means half CS if the searchers are at CS or further apart; if they are closer than CS then it extends for less than half CS beyond the end searchers.⁶ This approach is perfectly valid, but should be used with care.

In fig. 11, the spacing between the searchers is expressed in terms of critical separation; thus when the searchers are spaced at CS their spacing is 1, when they are at half CS it is 0.5 and so on.

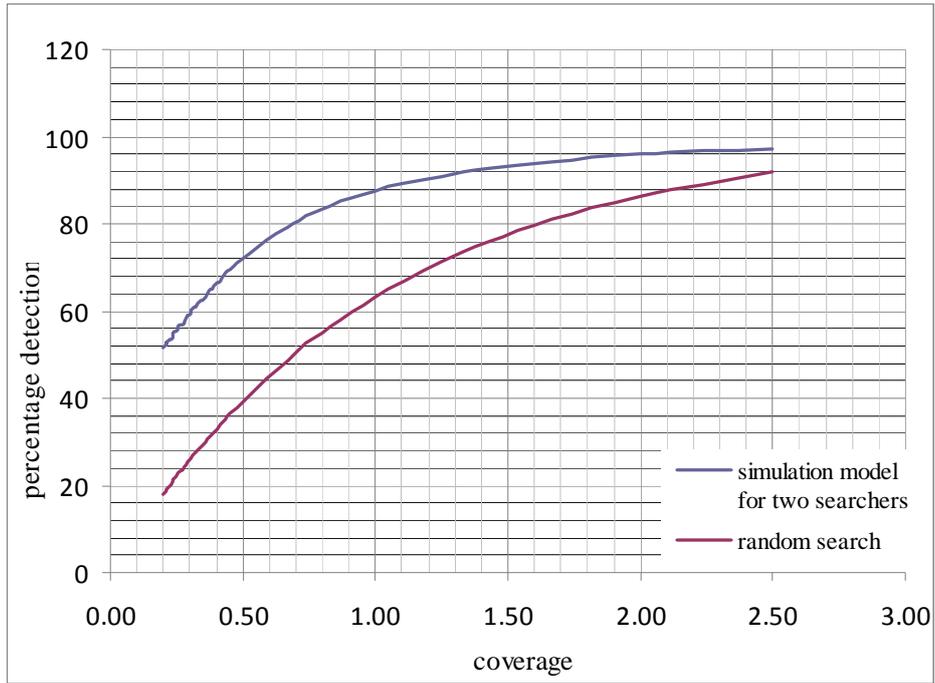


Fig. 10: graph of coverage against percentage detection from the two-searcher model, with the random search curve for comparison

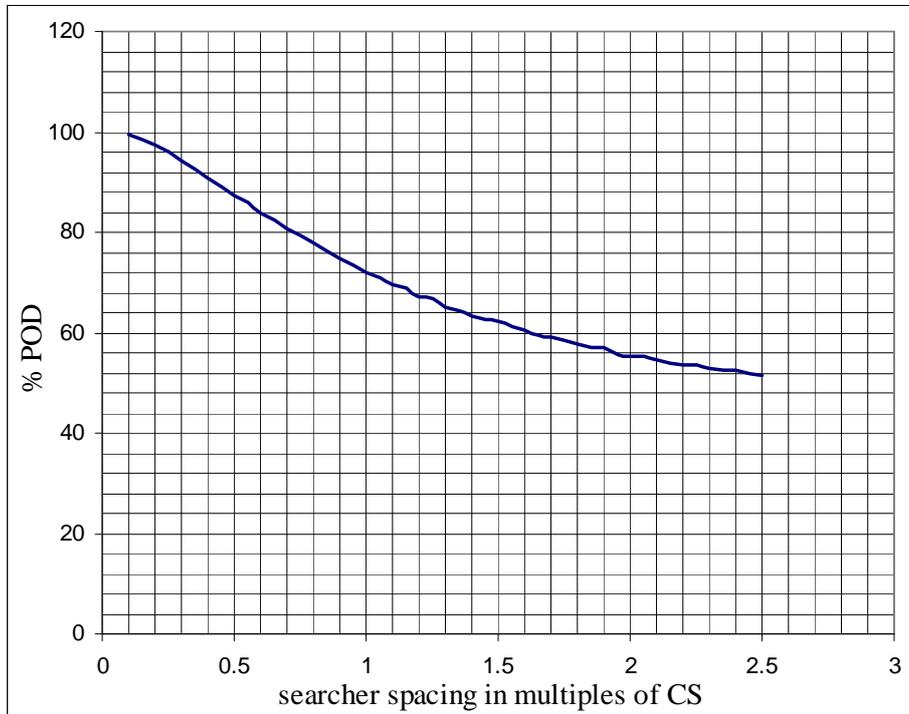


Fig. 11: graph of % POD against searcher spacing measured in multiples of critical separation, with fixed spacing and no purposeful wandering

The curve in fig 11 shows that the expected POD for searchers spaced at CS is considerably more than the 50% figure that has been quoted for a long time; in fact, the value from the graph is 72%. The reason for the difference in POD values is discussed in some detail later in the paper, but essentially it comes down to the fact that previous results were based on assumptions that oversimplified the problem.

Long range detections

The lateral range curves in fig. 4 clearly belong to the same family; as the lateral range increases from zero, they all start with a steep decline in the percentage of targets detected; the rate of decline diminishes as the lateral range reaches around twice the value of ESW and thereafter leads into a tail which seems to go on for some distance. Looking at early results from the model, it was clear that detections were indeed being made at much greater distances than might have been expected.

It is worth noting that in the real world it is not uncommon for a searcher to detect an object at a distance much greater than would be expected for the terrain in which they are searching.

An investigation into these so-called long range detections (LRDs) became one of the objectives of the study, and so each time the model was run it was made to report the furthest distance at which a target was detected both in grid units (“metres”) and in terms of critical separation. The ratio of the ‘distance to the furthest detected target’ to ‘critical separation’ for a particular obstacle density is referred to as the ‘LRD ratio’, thus:

$$\text{LRD ratio} = \frac{\text{distance to the furthest detected target}}{\text{critical separation}}$$

A series of ten runs through the model at each of the obstacle densities in the standard set of experimental obstacle densities took place; each time the model was given 50,000 rows and 500 columns (which is equivalent to a total of half a million detection opportunities for each obstacle density, or three million detection opportunities in total). A summary of the results is shown in table 3; this shows the minimum, maximum and mean value observed for the LRD ratio over the half a million detection opportunities for each obstacle density.

obstacle density %	minimum LRD ratio	maximum LRD ratio	mean LRD ratio
2.5	3.5	5.2	4.5
5	2.9	5.3	4.0
7.5	2.8	5.0	3.9
10	2.7	5.6	3.7
12.5	2.8	5.4	3.7
15	2.9	4.2	3.4

Table 3: long range detections for the standard set of experimental obstacle densities

As had been anticipated, the mean values of the LRD ratio decreased as the obstacle density increased; a simple explanation for this would seem to be that at low obstacle densities there are more ‘gaps in the vegetation’ for the searcher to see through than there are for higher obstacle densities. For the minimum and maximum values of the LRD ratio, however, the situation was erratic, which probably reflects the random nature of the distribution of targets and obstacles. Notice that in table 3, all except the largest obstacle density reported maximum values of the LRD ratio of at least 5 x CS. The largest observed value was the detection of a target at a distance of 5.6 x CS for an obstacle density of 10%; the model was giving a value of around 20 grid squares (“metres”) for CS for an obstacle density of 10%, and so this is equivalent to a searcher detecting a target at a distance of just over 110 metres in the type of terrain for which CS is 20 metres.

The smallest observed value of the LRD ratio was 2.7 x CS, also, as it happens, for an obstacle density of 10%. Using the same figures as above, this equates to a searcher detecting a target at a distance of 54 metres in terrain where CS is about 20 metres. This is still a target detection at a greater distance than one would anticipate.

The overall mean value of the LRD ratio over all three million detection opportunities was 3.9; in other words, the model is suggesting that **we should not be too surprised when a searcher detects an object at a distance of 3 x CS or more.**

Purposeful wandering

Experienced searchers often use a technique called purposeful wandering. Instead of walking in a straight line and keeping at a constant spacing from the searchers on either side of them, they move around within the strip of ground that they are responsible for searching in the manner described below. This enables them to see into locations that would otherwise be hidden from their view. The way in which they perform purposeful wandering is as follows:

1. the searcher takes up a position in the centre of the strip of ground that is theirs to search
2. they walk down the centre of the strip until their view of the outer edge of the strip is blocked by an obstacle
3. they then walk across to the obstacle (this is the purposeful wandering part) and take a look behind it
4. if there is nothing there, they return to their last position in the centre of the search strip
5. they continue in a straight line down the centre of the strip

In terms of the model (see fig. 2) purposeful wandering works like this: consider the searcher **S** travelling up column 0; when they reach row 2 their view along the row is blocked by the obstacle in column 1; they will therefore move across (they wander from their straight line path) to take a look behind the obstacle, and then move back to their original position. In fig. 2, if they did that then all that they would see would be the obstacle in column 3. Imagine, however, the same thing happening in row 998: the searcher **S** wanders across to take a look behind the obstacle in column 2, which gives them a clear view of the target in column 100.

This is relatively easy to incorporate in the model: all we need do is insert an algorithm that says that if the nearest thing to the searcher **S** in column 0 is an obstacle, then remove it. Therefore in fig. 2, the obstacles nearest to the searcher in rows 2, 3, 998 and 1000 would be removed; the targets in rows 3, 998 and 1000, which would have been undetected in the original model, are now detected.

It was decided that a useful measure of the effect of purposeful wandering would be the change in ESW.

Another version of the single-searcher model was produced that included the purposeful wandering algorithm. In this, the searcher has two passes through the search area; the first of these is exactly the same as in the original model, in other words no purposeful wandering. This is to establish the number of detections and hence ESW for the configuration of targets and obstacles. For the second pass, the purposeful wandering algorithm is applied as soon as the searcher reaches a row where the nearest thing to them is an obstacle. ESW was found for both passes by means of both the crossover analysis and the area under the lateral range curve, and averaged. The results are shown in table 4.

Table 4 shows the sweep widths with and without purposeful wandering (PW) for different obstacle densities; the final column shows the ratio of the two values of ESW. Based on these, it would be reasonable to assume that **the effect of purposeful wandering is to double the sweep width and hence double the coverage.**

obstacle density	ESW with no PW (a)	ESW with PW (b)	b / a
2.5	40.3	80.7	2.00
5	20.3	40.7	2.00
7.5	13.6	27.3	2.01
10	10.2	20.4	2.00
12.5	8.3	16.7	2.01
15	6.8	13.6	2.00

Table 4: sweep widths for various obstacle densities with and without purposeful wandering (PW)

Another consequence of this is that **where searchers are using purposeful wandering, equation (2) becomes**

$$\text{ESW} = 1.04 \times \text{CS} \quad (5)$$

or, more realistically

$$\text{ESW} = \text{CS} \quad (6)$$

As an illustration of the effect of purposeful wandering on POD, suppose that a field team is searching with a fixed spacing of CS. Their coverage is 0.5, and fig. 10 tells us that their POD will be 72%. If they use purposeful wandering then their coverage is 1.0, and their POD becomes 87%.

Comparisons with previous work

This is the third in a series of papers that has described the evolution of critical separation and its associated concepts since 1989. Some of the results in the paper differ from results in the previous two papers; this is an attempt to explain why.

The first paper² came out of a very simple analysis that looked at one searcher who had what turned out to be a linear lateral range curve (LRC), although the author was not aware of LRCs at the time. In this, CS was identified for the first time in writing, and it was stated that the percentage POD for searchers at CS is 50%. Well, in the region between two searchers with linear LRCs who are at CS and with no other searchers around it might well be 50%, but unfortunately the assumption about linear LRCs appears to be not quite right. In addition, this analysis considered only two searchers, and somehow overlooked the fact that when their spacing is less than half CS then a third searcher would be able to see into the region between them, which is where all the action is from the point of view of calculating POD; this third searcher could therefore make a contribution to the POD. Call it an oversight. But nevertheless, CS has stood the test of time as a field parameter, even if people running line search field trials did wonder why the participants usually found more than 50% of the objects they put out. Maybe this paper explains why.

Next, in 2008, came the paper The Critical Distance Method³. Critical distance, of course, amounts to nothing more than half CS. This analysis also involved a linear LRC (so there's a problem straight away) but resolved the issue of the third searcher. That paper was also an attempt to bridge the divide between land SAR and search theory; it included the concepts of LRCs, ESW and showed its results in the form of a coverage / POD curve. According to this, searchers at CS (coverage = 0.5) would expect to have a POD of 50% (as in the original paper), but as their spacing decreased (and coverage increased) it gave POD values in excess of the values from the original paper on account of the third searcher.

And now this study. This approached the problem in a completely different way; the search area was seen as a blank page onto which targets and obstacles could be dropped in random locations. The purpose of the model was to try to discover anything that might enhance our understanding of searching. Its results are different from the first two papers because it did not start off with any pre-conceived ideas about LRCs or the distance at which targets can be detected ... but in some respect the results are not all that much different. What it does do is show how CS is somehow central to all that happens, and that only helps to make it a more valuable parameter.

Limitations: what the model doesn't do

There are many things that the model does not do; here are some of them.

The first thing that the model doesn't do is allow its searchers to miss targets because they are inattentive or weary, even towards the end of a search area that is 250 km long. This is not as it is in real life, but fatigue and inattention were not designed into the model since exploring them was not part of the original intention.

The model has been built on the assumption that the searcher is as likely to see a target at a distance of 500 metres as they are at a distance of 1 metre, and that the only reason why they will not see it is because things will get in the way. In some ways this is slightly surreal (especially if we think of the subtended angle approach to some parts of search theory), but since the main objective of the model was to start with no preconceptions and see what happened, then why not.

Fundamental to the model is the notion of obstacle density; this refers to how many obstacles there are on average in each part of the search area. The model makes use of it but at no time does it tell us how we can recognise various amounts of it in the real world. This is not a problem: it uses obstacle density purely as a mechanism to produce results in terms that we can understand, such as CS, ESW and POD.

An alternative meaning of the term 'obstacle density' might mean how dense or opaque each obstacle is. While this might seem to be a trivial distinction, maybe it isn't. Consider for example the two people in the woods who were mentioned at the beginning, one of whom is wearing a green jacket, the other a high visibility orange coverall. A tree with fairly sparse foliage might act as an obstacle with regard to the person in the green jacket (a searcher would not detect them on the other side of it) whereas for the same searcher in the same location that tree would be unlikely to prevent them from detecting the person in the orange coverall; even though they might not be able to see all of that person, they would see enough of them to make a detection. At the moment the characteristics of the target and the environment are all taken together in the model, as they are in CS when we talk about "... for a particular target in a particular type of terrain in a particular level of visibility ..."; maybe it would be more correct to think of an obstacle in terms of its relationship with a target in terms of colour and environment. Perhaps a better way to think of an obstacle would be as a filter rather than a brick wall; that might enable us to build in rules about how much of the target they allow through, and whether or not the searcher would detect that at the distance under consideration. That was well beyond the scope of the current model.

Conclusions

1. The original objective of this project was to investigate detections in an environment containing things that would stop the searcher detecting a target if they got in the way. Prior knowledge of lateral range curves and detection distances was not allowed. The idea was that if there were any rules then perhaps the model would discover them. This seems to have been achieved.
2. The relationship between sweep width and critical separation has been the subject of some speculation; equation 2 seems to resolve this issue and could prove to be of

value to search planners. It provides a value for ESW, albeit slightly rough and ready, but which is nevertheless available in real time, i.e. when the search is taking place. It is also based on the conditions in the field at the time, and not on a figure obtained from elsewhere and somehow adjusted to fit local conditions.

3. It is felt that the coverage / POD curve (fig. 10) has sufficient justification for its validity, and should therefore be used to give POD values when coverage is known.
4. In situations where coverage is not known, fig. 11 should be used; this gives POD for searchers with a fixed spacing and no purposeful wandering; users should be aware of potential problems pointed out in the text. This curve was derived from the graph in fig. 10 and therefore comes with the same level of justification.
5. The graph in fig. 11 represents perhaps the final version of the spacing / POD graph based on critical separation; the evolution of this graph is described in the text, and would seem to suggest that previous problems and oversights have been resolved.
6. The analysis of long range detections is considered to be an interesting outcome of the study. The author is aware of these occurring, in well conducted field trials as well as on other occasions, but is not aware of any investigation into this aspect of searching. If nothing else, it will make the land SAR community aware that these can happen, provides a measure of it and offers an explanation for it.
7. The quantification of the effect of purposeful wandering is considered to be a most valuable outcome. In the past there has almost seemed to be an element of the fudge factor in operation when trying to explain the effect of using PW. In addition, if any field teams are not sure how purposeful wandering works then the explanation in the text and the algorithm should give them a better understanding.

References

- 1 Koester, R., Cooper, D.C., Frost, J.R. and Robe, R.Q., (2004), *Sweep Width Estimation for Ground Search and Rescue*, Potomac Management Group Inc, paragraph 3.3.4.1, page 33
- 2 Perkins, D., (1989), *Probability of Detection (POD) Research and Other Concepts for Search Management*, ERI International
- 3 Perkins, D., (2008), *The Critical Distance Method*, The Centre for Search Research, can be downloaded from URL: http://www.searchresearch.org.uk/www/published_papers/
- 4 Koopman, B.O., (1980), *Search and Screening*, Pergamon Press, New York
- 5 Perkins, D. and Lovelock, D., (2008), *Lateral Range Curves, Search Probabilities and Grid Searching*, The Centre for Search Research, page 7; this can be downloaded from the same URL as 3
- 6 as 3, page 8, observation 4

Acknowledgements

The author is grateful to David Lovelock, Department of Mathematics, University of Arizona, Tucson, and Pete Roberts, The Centre for Search Research, Ashington, Northumberland, for providing encouragement during the course of this work.