

Combining Random and Data-Driven Coverage Planning for Underwater Mine Detection

Jason R. Stack and Cheryl M. Smith

Naval Surface Warfare Center
Coastal System Station
6703 West Highway 98
Panama City, FL 32407, USA

Abstract—Unmanned autonomous vehicles are proving themselves to be effective means for conducting underwater mine hunting missions. The resulting efficiency, reduced search time, and covert search possibilities will facilitate larger mission areas requiring many agents searching for significant lengths of time (*e.g.*, many weeks or months). In search areas of this scale, complete area coverage may not always be feasible. Therefore, this research investigates a path planning scheme for incomplete coverage. This scheme divides the search area into cells and surveys each cell using a conventional line-sweep pattern with a row spacing that is larger than the sensor footprint. The rows of the line-sweep pattern are randomly spaced near the boundaries of each cell to decrease the probability of missing a line of evenly spaced mines. The spacing of the rows near the center of each cell are specifically determined from estimated possible mine locations. Bounds placed on the row spacing limits the amount of uncovered area to maintain an acceptable probability of detection. This method results in a probability of missing a mine that is less than the percent of unsearched area.

Keywords—Coverage path planning, mine countermeasures, unmanned underwater vehicles.

I. INTRODUCTION

In the near future, naval mine countermeasure (MCM) operations will be initiated by multiple autonomous agents surveying an area for the presence of underwater mines. This approach will facilitate significant improvements in search time and efficiency compared to the more traditional methods of a large manned vessel towing a sensor platform over the search area. In addition to requiring significant manpower, the mission planning phase of traditional MCM operations often consumes as much time as the survey itself. An autonomous MCM operation could be specified simply by programming a set of agents with a search boundary. The agents would then plan the details of the mission *en route* or in real time while searching. Additional advantages of autonomous MCM include reduced cost, fewer sailors in harm's way, and the option of performing covert missions.

When conducting underwater MCM within the surf zone, unmanned bottom crawlers are typically employed. For all waters outside of the surf zone, volume swimmers

or unmanned underwater vehicles (UUVs) are used. For these agents to survey an area effectively, a path or coverage scheme must be specified. There is an abundance of path planning algorithms in the literature. Many of these are concerned with navigation from point to point or through a set of waypoints while cleverly and efficiently avoiding obstacles. Within these techniques, there is a smaller subset known as coverage path planning algorithms. These algorithms are applicable to the problem of moving a sensor footprint over a given search area. Some popular applications receiving attention in the literature include autonomous lawn mowing, floor cleaning, and mine hunting.

II. COVERAGE PATH PLANNING

The goal of a coverage planning algorithm is to direct an autonomous agent to efficiently cover an area. Throughout the literature, the notion of what comprises an efficient search takes many different meanings. Example optimization criteria include minimizing the time of search, distance traveled, number of turns, or the agent's cost / complexity (*e.g.*, computational power or sensor capabilities).

A survey of coverage path planning algorithms is presented in [1]. This survey categorizes the existing coverage algorithms as either heuristic or complete. The heuristics-based algorithms rely on a simple set of rules or behaviors to guide the agent(s) through the search area. For example, many of the commercially available autonomous vacuum cleaners rely on a random search scheme. The vacuum cleaner simply travels in a straight line until an obstacle or boundary is encountered at which point it changes direction at a random angle, which still avoids the obstacle and remains within bounds. The main advantage of this method is its extreme simplicity, which is ideal for many consumer applications where the coverage time is not as important as product cost.

When search time is a factor, the random-based methods can be problematic. The study in [2] investigates autonomous bottom crawlers searching for mines in the surf zone using a random search scheme. This study concludes that the ability to cover a sufficient portion of the search area in a reasonable or allotted time is heavily dependant on the search parameters (*e.g.*, leg length, initial

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 01 SEP 2003		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Combining Random and Data-Driven Coverage Planning for Underwater Mine Detection				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Surface Warfare Center Coastal System Station 6703 West Highway 98 Panama City, FL 32407, USA				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002146. Oceans 2003 MTS/IEEE Conference, held in San Diego, California on September 22-26, 2003. U.S. Government or Federal Purpose Rights License., The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

agent positions, and search boundary). If any of these parameters are improperly specified or if conditions change during the search, the resulting coverage is poor.

Planning schemes classified as complete coverage algorithms typically involve partitioning the search area and then devising efficient techniques to cover each and every partition. In most algorithms, this is achieved by dividing the search area into “cells” where some criteria is used to define the boundaries of the cells. Each cell is typically covered using a simple line-sweep pattern (also known as a lawnmower pattern, seed sowing path, or boustrophedon path) illustrated in Fig. 1. The advantages of the line-sweep pattern include its simplicity and the absence of overlapping tracks. With this search pattern, the problem of coverage planning is then reduced to how to define each cell and to find an efficient path between adjacent cells.




Figure 1. Line-sweep pattern.

Many algorithms in the literature place an emphasis on finding a complete coverage while accounting for obstacles within the search bounds. In [3], a technique is proposed that attempts to limit the number of cells used in describing the search area so as not to create redundancies or inefficiencies at the cell boundaries. This method’s completeness of coverage is provable, and it can handle arbitrarily shaped (*i.e.*, non-polygonal) obstacles.

The authors in [4] address coverage planning in the floor-cleaning problem. In their research, path planning is accomplished while addressing the non-holonomic constraints of the agent (*e.g.*, preventing collisions with the boundary while making a U-turn of nonzero radius). The same authors introduce a method to handle uncovered areas at the conclusion of the search in [5]. This is useful in the event that areas were originally missed while avoiding temporary obstacles (*e.g.*, other search agents). A coverage algorithm relying only upon touch / contact sensor information is presented in [6]. This algorithm also decomposes the search area into cells and employs a line-sweep pattern for coverage. Completeness of coverage for this algorithm is provable and this technique is extended to multiple cooperating agents in [7]. When a line sweep pattern is employed, the authors in [8] claim that minimizing the number of turns is the most important factor in an efficient search. They propose a cell

decomposition and coverage path algorithm that tries to optimize this criterion.

Other algorithms that address the same problems as coverage planning algorithms are found in the field of emergent behavior. A popular example is the research in [9] in which multiple robots search and retrieve geologic samples. These agents wander randomly, and once an object is found, they collect it and follow a beacon to a collection station. This scheme takes the additional step of instilling cooperation between the agents by allowing them to drop markers on their way to the station. Thus if a trail of markers is encountered during the wander stage, it is followed (and partially picked up) to a potential source of desired samples. The trail is again reinforced on the way back to the station if samples are found. While these emergent behaviors are not explicitly categorized as coverage path planning algorithms, they are applicable to many of the same problems.

III. COVERAGE PLANNING FOR UNDERWATER MCM

A typical coverage path plan for a naval underwater MCM mission is composed of one or more line-sweep patterns designed to cover 100% of the search area. Generally, the coverage path is planned well before any sensors are placed in the water. An example mission, coverage plan, and results are described in [10] for an MCM operation performed by an autonomous UUV. In this mission, an area of 3.1 km² was surveyed in approximately 15 hours (divided over 3 days) by one UUV. The coverage plan consists of 3 partly overlapping cells surveyed with line-sweep patterns, which completely cover the search area. While this example hints at some of the promising possibilities of autonomous MCM, the full scale of the naval MCM problem is much greater.

The emergence of multiple autonomous agents for naval MCM will contribute to improved search efficiency and reduced search time. This will in turn lead to more ambitious MCM missions. Consider the two different size MCM missions illustrated in Table I. This table illustrates two different scales of the MCM search problem. The sizes illustrated in this example are arbitrary yet realistic and indicative of the mission areas that need to be considered. These coverage rates are based on a perfect sensor with a 40 m swath traveling at a constant 4 knots and only wasting 15% of each 12 hour search day for turning and travel time. The rationale behind this example is that a “small” area is one that a few agents can search in one or two days (here 5 agents can complete the search in one day). This problem is straightforward since the coverage planning could consist of one or more sets of line-sweep patterns designed for complete coverage of the area (as in the example of [10]). Then the agent or agents could cover the entire area in a single expedition.

The “large” search area is significantly more complicated. With a search problem of this scale, complete coverage could require many agents searching for many weeks or months. For the example in Table I, 10 agents will take 46 days to search this area. Additionally as the

search area becomes sufficiently large, the travel time from the recharging station to the starting point of that day's search begins to dominate each agent's time. Therefore, each agent may spend the majority of its 12-hour day traveling, which will significantly reduce its effective search time. Hence, 457 robot-days to cover the large area could be a gross underestimate of the required time depending on the search area geometry, recharging logistics, etc. Therefore, the survey of the large area needs to be meticulously planned since inefficiencies could add days or weeks to a mission of this scale.

TABLE I.
Arbitrarily sized example search areas illustrating the scale of the MCM search problem

Small Area	Large Area
2 nm x 2nm	10 nm x 40 nm
14 km ²	1372 km ²
4.7 robot-days	457 robot-days

*Based on 3 km² coverage in a 12 hour day at 100% coverage

When naval operations demand one of these large search areas is to be cleared of mines, significant resources (agents) and time are required. In this instance, it is possible that complete coverage of the area is not feasible. This could be due to lack of search agents; time limitations; or to an unforeseen, premature termination of the search partway through. If the problem is a deficiency of resources, there are essentially two options: reduce the size of the search area or survey the entire area with less than 100% coverage while still remaining within the mission's acceptable level of risk. A variation on this later option is to divide the search area into a region that must receive complete coverage and a remaining region where incomplete coverage is acceptable. In either instance, there is a need for an efficient coverage plan designed to survey an area at a specified percent coverage that is less than 100%. The specific percent of coverage produced by the coverage planning scheme should be adjustable to fall within the acceptable level of risk for each individual MCM mission.

IV. RANDOM + DATA DRIVEN COVERAGE PLANNING

One solution to this problem is to decompose the search area into cells and simply not search some of the cells. In this instance, the probability of a missed mine is equal to the percent of unsearched area. This statement and the remaining developments in this research assume a perfect detector (*i.e.* if the sensor footprint covers a mine, it will be detected). The disadvantage of this coverage plan is that entire regions of the search area receive no coverage. Therefore, there could be entire minefields in these unsearched cells and there would be no chance for their detection. A better option is to design a coverage plan that ensures the uncovered regions are relatively small (*e.g.*, on

the order of the mine spacing) and distributed over the entire search area. This scheme is beneficial because mines are typically placed in lines rather than as isolated entities. Therefore, keeping the unsearched regions small and scattered will decrease the likelihood of missing an entire mine line.

One way to implement this principle is to cover every cell with a line-sweep pattern where the spacing between each row is greater than the sensor swath. This produces a coverage pattern with gaps between each row in the line-sweep pattern. The advantage of this scheme is that there are no large, unsearched regions; there are only unsearched gaps between each row. The disadvantage of this scheme is that it still produces a probability of missed mines equal to the unsearched area. Therefore, a mine line containing a significant number of mines could go undetected. This scenario is illustrated in Fig. 2 where the mines and the rows of the sweep pattern have similar spacing. The mine line can then fall in the gaps between each row of the sweep. For the illustration in Fig. 2, only 37% of the cell is covered by the detector; however, 100% of this unsearched area is still vulnerable to mines (*i.e.*, a mine line could fall anywhere in the unsearched area and pass completely through the cell without being detected).




Figure 2. A cell surveyed using a line-sweep pattern with incomplete coverage. A mine line can fall anywhere in the unsearched area and not be detected.

Typically, when underwater mines are initially deployed, they are placed in lines or fields with regular spacing. This allows the navy that installed the minefield to navigate safely around or through the field or to collect the mines safely at a later date. The actual spacing between underwater mines depends on the lethal range of the mines used; however, there are several categorizations of mine placements commonly used in the public literature. A "mine cluster" generally refers to one or more rows of mines with spacing less than 300 m and row lengths less than 1000 m. A "mine line" generally refers to one or

more rows of mines with spacing up to 300 m and row lengths of 1.5 km to 11 km. While these distances vary throughout the literature and depend on the specific mine type used, they provide a general idea of the scale of the problem. In this paper, the term “mine line” refers to any linear placement of evenly spaced mines.

If each potential mine line is assumed to contain evenly spaced mines, then randomly varying the spacing between each row in the line-sweep will decrease the probability of missing an entire mine line. This is because the random spacing of the rows will effectively reduce the possible locations that evenly spaced mines could occur. This principle is illustrated in Fig. 3. In this figure, approximately 2/3 of a cell is searched from left to right using randomly spaced rows. The overlapping “*” symbols in the middle of the graph represent the x-axis coordinate of every possible location in which an evenly spaced mine line could exist and remain undetected. This assumes the mine line begins on the left side of the cell (in the gap between 40 m and 90 m) and continues to the right. For this illustration, the minimum allowable mine spacing is 40 m (the sensor footprint width) and the maximum allowable mine spacing is 100 m. From this figure, it is seen that there are regions in the right half of the search area where no mines can exist under these assumptions. The implication is that the probability of a missed mine is less than the percent of unsearched area.




Figure 3. A cell surveyed using a line-sweep pattern with randomly spaced rows. Only 2/3 of the cell is surveyed, and all possible locations for an undetected, evenly spaced mine line are indicated by overlapping “*” symbols in the middle of the figure.

The illustration in Fig. 3 only indicates the x-axis coordinates of possible mine locations. An actual mine line could occur at any arbitrary angle relative to the rows; however, as long as it traverses multiple rows in the line-sweep, the x-axis coordinates of the mines must still fall in these same gaps to go undetected. In the situation where the mine line is exactly parallel to the rows, it could pass

completely through the cell without detection. This possibility supports the rotation of the sweep direction by 90° in the adjacent cells, which is illustrated in Fig. 4. With this orientation, mine lines that fall parallel to the rows in the center cell are more likely to be detected by the search in the top or bottom cell. Therefore, the search in the center cell is concerned with minimizing the possible undetected mine line locations that enter the cell from the left or right.

Fig. 5 shows all possible locations for mine lines that enter the cell from the left or from the right. In this example, 48% of the cell is surveyed by the detector; however, undetected mines can occupy only 69% of the unsearched area. This is a significant improvement over the 100% vulnerability of the unsearched area for the evenly spaced line-sweep in Fig. 2. This illustrates how a line-sweep pattern with randomly spaced rows can result in a probability of missed mines that is less than the percent of unsearched area.




Figure 4. Rotating the line-sweep orientation by 90° in adjacent cells helps detect mine lines parallel to the rows in the center cell.

To extend this concept and further reduce the probability of missing a mine, a non-random, data-driven placement is employed for the rows in the center of the cell. Consider a cell where an agent has surveyed the first few (randomly spaced) rows in the line-sweep (e.g., Fig. 3). The agent then estimates the remaining possible undetected mine locations and then spaces some number of rows near the center of the cell to specifically eliminate “clusters” of possible mine locations. Using this approach, a further reduction in the probability of a missed mine is achieved.

For example, the path illustrated in Fig. 6 is the same as in Fig. 5 except for the two rows in the middle of the cell. In Fig. 5, all rows are randomly spaced. However, in Fig. 6, two rows in the middle of the search area are specifically spaced by the search agent to eliminate groups of possible mine locations. The row at 570 m in Fig. 5 is moved to 540 m in Fig. 6, and the row at 690 m is moved to 640 m. This reduces the amount of unsearched area that undetected mines could occupy from 69% in Fig. 5 to 51% in Fig. 6.

From this discussion, it is seen that this research computes the percent of space on the x-axis that undetected mines can occupy rather than explicitly estimating the probability of missed mines in the cell. This is because computing the vulnerable space along the x-axis direction is significantly less complicated, requires fewer assumptions and is easier to perform in real time. This facilitates real-time computation onboard a search agent without consuming processing resources. This method is also slightly conservative because it does not account for the travel in the x-axis direction that connects each row.




Figure 5. All possible undetected mine locations from mine lines originating from the left or from the right. These mines occupy 69% of the unsearched area.




Figure 6. The same pattern as Fig. 5 except two rows (here at 540 m and 640 m) have been repositioned by the search agent. This reduces the vulnerable amount of unsearched area from 69% in Fig. 5 to 51% here.

In practice, this coverage scheme is implemented by each agent during each cell search. Upon reaching a new cell, the search agent surveys the first several rows of the line-sweep pattern with random row spacing. Then the agent generates randomly spaced row locations for the remainder of the cell and estimates the x-axis coordinates of possible mine line locations. The rows in the middle of the cell are then repositioned to minimize the possible undetected mine locations. Therefore, the placement of the rows on the right and left sides of the cell is random while the placement of the rows in the middle is a function of the estimated undetected mine locations. Bounds are placed on the minimum and maximum row spacing to ensure the desired percent coverage of the entire search area.

For the search agents to compute the non-random or data-driven row spacing in the middle of the cell, the agents need to estimate the x-axis coordinates of the possible mine locations. The estimation calculation should be computationally tractable and inexpensive. While there are many ways to perform this estimation, this research employs a numerical method for its computational simplicity. The *a priori* assumptions required include the minimum and maximum allowable mine spacing and that the undetected mine line will originate from either the left or right side of the cell and traverse at least one row in the line-sweep pattern. It also assumes a perfect probability of detection for any mines within the sensor footprint. The width and shape of the sensor footprint can be modified to account for specific sensor types.

Given a fully or partially surveyed cell containing N gaps or uncovered lanes between the rows, the possible undetected mine locations are estimated as follows. Define the x-coordinates of the n^{th} uncovered lane in the cell by a_n and b_n (e.g., the x-coordinate of the left side of the 2^{nd} uncovered lane is a_2 while the x-coordinate of the right side is b_2). A template mine line of M mines and allowable spacing is then constructed where L_m is the x-coordinate of the m^{th} mine in the template mine line. If mine m falls in lane n , then the inequality in (1) is valid.

$$(L_m - a_n)(b_n - L_m) > 0 \quad (1)$$

This expression can be expanded as

$$L_m a_n + L_m b_n - L_m^2 - a_n b_n > 0, \quad (2)$$

and then rewritten in the matrix form of (3).

The first k consecutive rows in the result of (3) that contain at least one positive value indicate that the first k mines in the template mine line L will go undetected. This calculation is then repeated for template mine lines that span the range of allowable mine spacing and mine line starting positions. The advantage of this method is that possible mine line positions are verified by calculations involving only matrix multiplication and can thus be efficiently performed in real time.

$$\begin{aligned}
& \begin{bmatrix} L_1 & L_1 \\ L_2 & L_2 \\ \vdots & \vdots \\ L_M & L_M \end{bmatrix} \begin{bmatrix} a_1 & a_2 & \cdots & a_N \\ b_1 & b_2 & \cdots & b_N \end{bmatrix} \\
& - \begin{bmatrix} L_1 & 0 & \cdots & 0 \\ 0 & L_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & L_M \end{bmatrix} \begin{bmatrix} L_1 & L_1 & \cdots & L_1 \\ L_2 & L_2 & \cdots & L_2 \\ \vdots & \vdots & \ddots & \vdots \\ L_M & L_M & \cdots & L_M \end{bmatrix} \\
& - \begin{bmatrix} a_1 & a_2 & \cdots & a_N \\ a_1 & a_2 & \cdots & a_N \\ \vdots & \vdots & \ddots & \vdots \\ a_1 & a_2 & \cdots & a_N \end{bmatrix} \begin{bmatrix} b_1 & 0 & \cdots & 0 \\ 0 & b_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & b_N \end{bmatrix} > 0
\end{aligned} \tag{3}$$

V. CONCLUSIONS

This research is concerned with coverage path planning for MCM operations large enough to require many agents searching for many weeks or months. In search problems of this scale, it is likely that complete coverage is not always possible. Therefore, this research investigates methods of path planning for incomplete coverage. A method is proposed in which the search area is divided into cells and each cell is surveyed with a line-sweep pattern. The rows in the beginning and end of the line-sweep in each cell are randomly spaced while the rows in the middle are specifically placed to minimize the possible locations for undetected mines.

This scheme produces a probability of missed mines that is less than the percent of unsearched area given a perfect detector. This method also keeps the unsearched regions relatively small and spread out, which decreases the likelihood of large mine lines or clusters going

undetected. An additional motivating force in the use of randomly spaced rows is to achieve unpredictability in the MCM missions. If MCM strategies designed for incomplete coverage employ patterns or probability distributions in the coverage design, these assumptions will be exploited in the deployment of the minefields. An efficient method for estimating the x-axis coordinates of undetected mine locations is also presented. This method is computationally simple, efficient, and suitable for real-time implementation.

REFERENCES

- [1] H. Choset, "Coverage for robotics—a survey of recent results," *Annals Mathematics & Artificial Intelligence*, vol. 31, pp. 113-126, 2001.
- [2] L.R. Howell, W.C. Littlejohn, C. Guastella, and N. Rodriguez-Casanova, "Defining surf zone crawler search strategies for minefield reconnaissance," *IEEE / MTS OCEANS*, vol. 1, pp. 85-96, 2001.
- [3] H. Choset, "Coverage of known spaces: The boustrophedon cellular decomposition," *Autonomous Robots*, vol. 9, pp. 247-253, 2000.
- [4] C. Hofner and G. Schmidt, "Path planning and guidance techniques for an autonomous mobile cleaning robot," *IEEE/RSJ Proc. Intelligent Robots and Systems*, vol. 1, pp. 610-617, 1994.
- [5] G. Schmidt and C. Hofner, "An advanced planning and navigation approach for autonomous cleaning robot operations," *IEEE/RSJ Proc. Intelligent Robots and Systems*, vol. 2, pp. 1230-1235, 1998.
- [6] Z.J. Butler, A.A. Rizzi, and R.L. Hollis, "Contact sensor-based coverage of rectilinear environments," *IEEE Symp. Intelligent Control/Intelligent Systems and Semiotics*, pp. 266-271, 1999.
- [7] Z.J. Butler, A.A. Rizzi, and R.L. Hollis, "Cooperative coverage of rectilinear environments," *IEEE Proc. Robotics & Automation*, vol. 3, pp. 2722-2727, 2000.
- [8] W.H. Huang, "Optimal line-sweep-based decompositions for coverage algorithms," *IEEE Proc. Robotics & Automation*, vol. 1, pp. 27-32, 2001.
- [9] L. Steels, "Cooperation between distributed agents through self-organization," *IEEE Proc. Intelligent Robots & Systems*, pp. 8-14 suppl., 1990.
- [10] R. Stokey, *et al.*, "Very shallow water mine countermeasures using REMUS AUV: A practical approach yielding accurate results," *IEEE / MTS OCEANS*, vol. 1, pp. 149-156, 2001.