An efficient method of combining detection and identification of seafloor objects using Gavia AUV

Hordur Johannsson
Department of Computer Science
University Of Iceland
Reykjavik, Iceland
Email: hordurj@hi.is

Torfi Thorhallsson
Hafmynd
Reykjavik, Iceland
Email: torfi@gavia.is

Hjalmtyr Hafsteinsson
Department of Computer Science
University of Iceland
Reykjavik, Iceland
Email: hh@hi.is

***Abstract*— This paper describes a method of dynamically searching for objects on the seabed using side-scan sonar images. Traditionally this has been a two step process where first the area is mapped and then later objects are chosen for inspection in a second pass. In the paper a model is described to explore different search tactics. Using this model we show that it can be beneficial to use an on-line detection algorithm and inspect objects as soon as they are detected. This is because the vehicle will position the objects more accurately and we will not need to travel again through the whole area. The benefit will depend on the number of objects that are in the search area and the quality of the online detection algorithm. An CAD/CAC algorithm was developed as a part of the project and simplified inspection patterns are also explored.**

I. Introduction

An application of considerable interest is the detection and identification of mine like objects (MLOs) on the seabed. For reliable identification of such small objects it is necessary to perform a high-resolution acoustic or photographic survey using a high-frequency side-scan sonar or a video camera. The fact that both these sensors have a very short range severely limits the rate of area coverage when applied in an exhaustive search.

Over the last few years algorithms for computer aided detection and classification (CAD/CAC) of mine like objects have been developed by several research groups that enable automatic detection of target candidates from side-sonar images. While some of these algorithms are computationally expensive, some have been shown to work in real-time onboard an AUV [1].

As shown in trials [1], an on-board CAD/CAC algorithm can automatically select for transmission to the operator small amounts of target data during a SCM mission over an available low-bandwidth channel. Using this information the operator can initiate a RI phase while the SCM survey is still underway by directing another AUV to inspect the target location.

Observing the fact that the navigation errors of an inertial navigation system accumulate slowly over time, the relocation error can be made extremely small by having the same vehicle revisit the target soon after it initially detects it. The experiment was carried out using a Gavia AUV equipped with dual frequency sonar, digital camera and an INS/DVL system for accurate positioning.

In this paper we develop this idea and propose a new method to search for objects on the seafloor combining the SCM and the RI phase using a single vehicle. In section II we describe a model that we used to compare different search methods. In section III we describe the search method and the search patterns developed. Finally in section IV we give a brief description of a CAD/CAC algorithm that was developed to be used for testing the dynamic search method.

II. Search methods

In this section a model to estimate the cost of searching an area will be described. Using the model we will give two different methods for executing the mission. Then those two methods will be compared for different values of the parameters.

The problem is to conduct a search of an area and then gather detailed information for objects found. This information will then be used to identify the objects. The model consists of a set of parameters, random variables and derived values.

Parameters and Variables: The primary input is an area that needs to be searched. This area is scanned using a side-scan sonar that is set at some specific range. The range is how far it sees to the left and right, so the whole swath that the side-scan sonar covers is two times that. It is also common to have some overlap. That overlap is controlled by changing the spacing of the legs. This is done to increase the probability that nothing is missed. For instance, the area directly under the vehicle gives shorter shadows and often makes objects harder to detect. One common method is to alternate the spacing as 50% and 150% of range. For example if we have the range set to 20 meters we would start with two legs that have 10 meters spacing and then 30 meters to the next, and then 10 meters again.

A	- Size of area to search in m^2
R	- Range of sensor in m
D	- Distance traveled/scanned

$$D = \frac{A}{2R} \tag{1}$$

The sonar data is split into a set of images. Each image contains a certain number of pings. The length of a single image could be set at two times the swath as is the case with the Marine Sonic Sonar. False alarm rate is the expected

1-4244-0115-1/06/$20.00 ©2006 IEEE

number of detections that do not turn out to be objects that interest us.

L - Length of image
p_{fa} - false alarms per L meters (FAPI)

In each image a certain number of objects will be classified as objects of interest. In the mine-hunting case it will be MLOs. The CAD/CAC algorithm will have a false alarm rating and a probability of detection. An area will have some number of true objects.

$p_d p_c$ - Probability of detection.
O - Number of objects

The total number of detections made by the algorithm is the sum of the number of objects correctly classified and the number of incorrect classifications. The expected number of correct classification depends on the number of objects in the area being searched and the algorithm's $p_d p_c$. The expected number of incorrect classifications depends on the size of the area and the false alarm rate of the algorithm.

O_t - Total number of objects found (true positive)

$$O_t = O \cdot p_d p_c \tag{2}$$

O_f - Total number of false alarms (false positive)

$$O_f = \frac{D \cdot p_{fa}}{L} \tag{3}$$

O_d - Number of detections

$$O_d = O_t + O_f \tag{4}$$

The total length of the mission will depend on the tactic used. It will at least contain the distance to scan the area and then in addition the length of executing the RI missions for objects detected.

L_{RI} - Length of a single RI pattern
M - Total mission length

The parameters of the model are the area, sensor range, length of images, false alarms per image, probability of detections and the length of RI patterns used. The number of objects in the area can be considered a random variable, but will depend on the application. Sometimes we are only searching for a single object, so either the object is in the area or it is not. In other cases we are searching an area for multiple objects. The total mission length is a random variable that will depend on the tactic used. The number of objects detected and the number of false alarms is also a random variable. Later when we compare the tactics we will be looking at the expected length of the missions.

A. Tactic 1

Search tactic 1 is the method most commonly used today when using side-scan sonar for mine hunting in the littoral environment. It consists of three steps:

1) SCM Run
2) Analyze data and detect objects
3) RI Run

Fig. 1. Tactic 1 SCM

For the SCM mission we have two legs. The total length of the mission is around 1500 meters. Three objects were detected.

Fig. 2. Tactic 1 RI

After the SCM mission a RI mission is planned to identify the objects detected earlier. Each RI pattern consists of 3 times 6 legs under 3 different angle. This is to cover 30 meters at 5 m range. The total mission length is estimated 6100 meters.

The SCM run is used to gather sonar images of the area so it will be possible to detect and classify possible targets. This data is usually processed by human operators. The length of the SCM mission is D. After the data has been processed, a set of RI missions is generated to collect more detailed information for the targets. This detailed information can then be used to identify the mines. Then a diver team is sent to neutralize the mines.

The length of RI mission is at least the shortest path through the detected targets and for each target execution of the RI pattern. We can consider this an instance of the Euclidean traveling salesman problem (TSP). It has been shown [2] that the shortest tour through N locations uniformly distributed over a rectangular area of A units is at least:

$$L_{opt}(N, A) = K\sqrt{NA} \tag{5}$$

where K is defined to be

$$0.765 \leq K \leq 0.765 + \frac{4}{N}$$

If we assume that all the detections will be uniformly distributed over a rectangular area, then equation (5) gives us a lower bound on the cost of traveling through all the detected objects. The expected number of detections is O_d. We execute one RI pattern for each so the cost of all the RI patterns is $O_d L_{RI}$.

Adding the length of SCM and RI run we get the lower bound on the total mission length for this tactic.

$$M > D + 0.765\sqrt{O_d A} + O_d L_{RI} \tag{6}$$

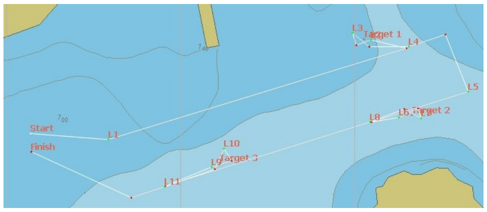

Here the simplified inspection pattern is executed as soon as the objects are detected. The estimated length for this mission is 2200 meters.

Fig. 3. Tactic 2

B. Tactic 2

Search tactic 2 is the method we are proposing in this thesis. It consists of the following step:

1) SCM/RI Run

While executing the SCM mission the side-scan sonar images are processed in real-time. When an object is detected the vehicle can decide to inspect it immediately. If it is possible to detect the object again later, because of image overlap, the vehicle can defer the inspection until later.

The total mission length is again the length of scanning the whole area in addition to the inspection runs that are carried out. This gives an upper bound on the total length that is

$$M < D + O_d(L_{RI} + D_O) \tag{7}$$

Here D_O is the distance the vehicle needs to travel to reach the object and then return back to the location where the object was detected. The maximum distance to an object is the length of one image or $2 \cdot R$. The maximum cost is then two times the maximum distance and the cost of two 180 degree turns. By controlling how much data is processed at a time, i.e. control the size of each image, we can lower the maximum distance.

1) Improvements: There are several improvements that we can do to shorten the mission length. Let us look at some of them.

If more than one object is detected in a single image, then it is possible to plan a RI pattern that would inspect all the objects that were detected in the image. In our analysis we calculate the cost for the RI patterns as if one object was inspected at a time, and we returned to the point it was located at when the detection was made.

When we execute missions that have overlap we can make a choice on whether to inspect an object at the time it is initially detected or do it later when we pass it on the next leg. One parameter that affects the length for each inspection is the image size and that affects the backtrack distance D_O when an object is detected.

Another thing that is possible is to perform off-line analysis of the data when the vehicle returns. The results from that could then be used to inspect objects that were missed by the online detection algorithm. We could also skip the inspection of objects that are close to the start of the mission, that would mean that in a secondary mission we would not have to go far from the start point for inspection. We will not analyze the effect of this last idea here.

C. Comparison of tactics

When comparing the two tactics we want to look at to how the different sizes of RI patterns and the different number of objects detected affect those two methods. Let L_1 and L_2 denote the length of the RI patterns for tactic 1 and 2 respectively. Furthermore let O_1 and O_2 be the number of objects detected by tactic 1 and 2 respectively. The distance D and the area A remain the same. Then equation (6) becomes

$$M > D + 0.765\sqrt{O_1 A} + O_1 L_1 \tag{8}$$

and equation (7) becomes

$$M < D + O_2(L_2 + D_O) \tag{9}$$

Let us compare where those two bounds are equal. By (8) and (9) we get

$$D + 0.765\sqrt{O_1 A} + O_1 L_1 = D + O_2(L_2 + D_O) \tag{10}$$

This simplifies to

$$0.765\sqrt{O_1 A} + O_1 L_1 = O_2(L_2 + D_O) \tag{11}$$

From this last formula we can see that important factors are the false alarm rates of the algorithms and the length of the inspection patterns. Further more we see that if the density of the objects in the area grows then the parameter D_O will start to matter, so we want to have $L_2 + D_O < L_1$.

Let us explore the case where we have $L = L_1 = L_2 \cdot K$ where K is some constant. Then equation (11) can be written as:

$$(K \cdot O_1 - O_2)L + 0.765\sqrt{O_1 A} - O_2 \cdot D_O = 0 \tag{12}$$

A negative left hand side means that tactic 2 is longer than tactic 1. From the first part we see that the scale factor between the inspection patterns gives us some flexibility in the false alarm rate of the detection algorithm. The second part tells us that that we will get more benefit from tactic 2 where the density of objects in the area is less.

Let us consider an example where the area is 500×250 meters, range is 40 and there are 5 objects in the area. Both methods have perfect detection capability, that is $p_d p_c = 1$. In chapter III we will look at sizes for different RI patterns. From that we let $L_1 = 1000$ and $L_2 + D_O = 300$. In figure 4 we compare the mission lengths of the two tactics for different false alarm rates. From the graph we can see that the mission length for tactic 2 is less if the on-line detection algorithm has equal or two times the FAPI of the off-line algorithm. But if the on-line detection has FAPI four times higher then the off-line algorithm, it will be worse to use the on-line detection when the FAPI for the off-line algorithm is more than 0.6.

Using a detection algorithm, that fuses results of multiple algorithms has been able to achieve probability of detection 100% and 0.22 FAPI [3]. That means that a vehicle with that algorithm could compete with an perfect off-line algorithm, i.e. 0.0 FAPI. If it is possible to improve the location accuracy in tactic 1 such that an RI pattern of length 600 would be

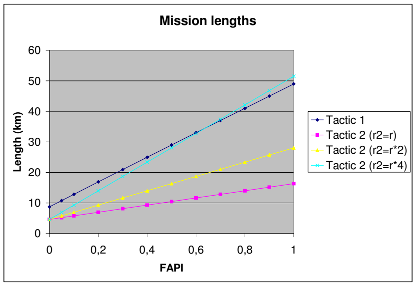

Comparisons of tactic 1 and 2 with $p_c p_d = 1$, $L_1 = 1000$ and $L_2 + D_O = 300$. The lines show how the mission length grows with the false alarm rate. Two of the lines show how the mission length for tactic 2 grows when the false alarm rate is 2 and 4 times higher than for tactic 1.

Fig. 4. Comparing tactics

sufficient, the it would mean that tactic 1 and 2 would be of equal length.

From the results above we see that which tactic is better will depend on the situation. The major factors are the false alarm rate and the size of the inspection pattern. The size of the inspection pattern will depend on the location accuracy, the more accurate the location the smaller pattern we can allow.

III. Planning

This section discusses the implementation of the second search tactic described in section II. As previously described the search tactic works using on-line processing of side-scan sonar data and then inspecting objects of interest that are detected. The implementation builds on existing behaviors in Gavia, such as mission execution, device management, line tracking, and bottom following.

A. Dynamic search methods

There has been some research into various search methods for AUVs. A technique for adaptive bathymetric contour tracking was worked on by Bennet et al. [4]. Olafsson has looked at search methods for the Gavia [5]. Some of the methods he covered were an adaptive search using fuzzy inference system, gradient search using scalar measurements and fixed patterns when searching for moving targets like a school of fish.

B. Object Search

The object search is a pilot crew-member that can be activated in the mission plan. It is responsible for activating the object detection operator. After that it waits for reported detections from the object detector.

When a detection is received from the object detector a decision on whether to inspect the detected object is made. This is done by comparing the cost of inspecting the object now or possibly later if the object will be passed again. It is common that we pass objects more than once because a typical search pattern is a lawnmower pattern, where we go back and forth over an area.

If it is decided that the detection should be inspected a new sub-mission is created. The inspection mission can include any of the patterns that are described in section III-D. This mission is then sent to the mission manager for execution. When the mission manager receives the mission it will halt the current mission while it executes the inspection mission. After the inspection mission is completed it will go back to the location it was at when the inspection mission was started. It will then resume the main mission.

In the mission plan it is possible to configure what type of inspection mission is executed and what sensors should be active when inspecting. It is also possible to specify some parameters for the object detector. This can be maximum number of objects notified per image, height of object and length of object. In the future it will be possible to have different detectors depending on the application. Different detectors could also need different parameters but those will not be specified here. When an object is detected the pattern will be centered over the estimated location of the object and then sent to the mission manager for execution.

An example of a pattern is the cross hatch pattern for digital cameras and a lawnmower patterns at different angles for the side-scan sonar. Also a simple pattern that would use both a camera and side-scan sonar.

C. Path generation

The Gavia can be considered as a uni-directional vehicle with bounded curvature. A path generator was implemented that created a path of minimal length and with bounded curvature using cubic spirals. The path generator works on the fact that given an initial configuration it is possible to get to another configuration using two cubic spirals. A configuration is a location and a heading. Liang et al. gave an algorithm that finds a minimal length path, of bounded curvature, between two configurations [6]. They built on work by Kanayama and Hartman on smooth paths using cubic spirals [7].

This was used here to explore various patterns. One thing that was noticed was that when executing small patterns the distance traveled when the vehicle turns is a considerable part of the length. Using the path generator an optimal path for a pattern of 3 times 2 legs was found. This was done by varying in what order the legs where tracked. Each leg was of length 30m, for a vehicle with turn radius 15m an intuitive way of executing the pattern was 708m but the optimal was 577m. Comparing those distance to the 180m ($6 \times 30m$) for the 6 legs shows that the cost for all the turns is considerable. This pattern was dropped later for a simpler pattern described in the next section.

D. Patterns

From the search model in section II we saw that the inspection pattern was a key factor in the size of the overall

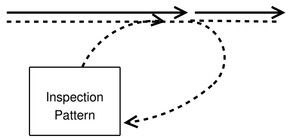

Fig. 5. An object is detected

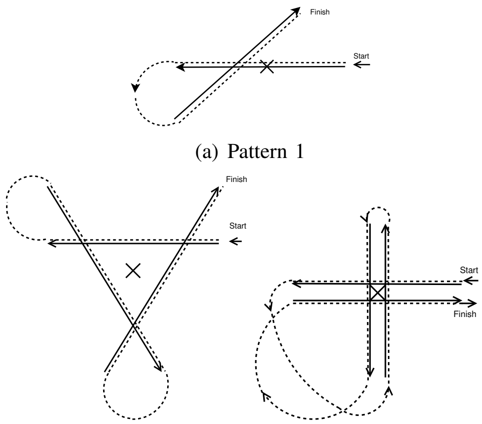

(a) Pattern 1

(b) Pattern 2

(c) Pattern 3

Fig. 6. Example of different inspection patterns

In figure 6(a) the vehicle first goes directly over the target to get digital images and then passes it at an offset on they way back to get high resolution side-scan sonar image. Figure 6(b) shows a simplified RI pattern that gives high resolution side-scan sonar images at 3 different angles. The pattern in 6(c) is cross-hatch pattern used to get digital images of the object.

mission. So at first we started looking into how the inspection pattern could be improved.

RI missions are often legs at different angles laid out over the estimated location of the target. This is because there is always a certain amount of uncertainty in the location of the objects and images of the objects from different angles help with identification.

First we looked into optimizing existing patterns. This way we could get some improvement. One of the potential benefits of doing the inspection soon after detection is that then the relative positional accuracy is very good, allowing us to use simpler patterns. There were three basic patterns that we looked into, a simple 2 line pattern, 3 times 1 RI, and finally a 2 times 2 crosshatch. These patterns can be seen in figures 6(a), 6(b), and 6(c) respectively.

We chose to implement the simple pattern because that would give the biggest gain. The simple pattern consists of a single line directly over the target, giving a photograph of the object. Then another leg close to the object, this leg should give a high resolution side-scan sonar image of the object. The main drawback of this pattern is that we have to navigate very accurately to be able to go directly over the object in one pass

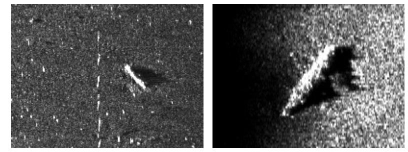

(a) Detected at 20m range (b) Inspected at 5m range

Fig. 7. The object that was inspected using pattern 1

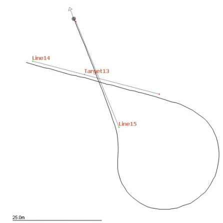

Fig. 8. Real execution of pattern 1

and we only get a side-scan sonar of the object from one angle.

Another issue that needs to be considered is when to act when a contact is made, for instance if the object lies closer to the next leg then it could be better to wait and do the close inspection when the object is closer to the AUVs main path.

E. Trials

The object search implementation was tested by having it create a mission plan for a simulated execution. The mission plan could then be used to test how the mission looked like.

Several sea trials where conducted to test how the vehicle would execute the patterns. We wanted to see how long they would be and how closely the vehicle would be able to track the selected path.

In the trials we initially scanned an area in the Reykjavik oil harbor and ran the files through a detection algorithm. A brief description of the algorithm is given in section IV. From this data we picked two targets that we decided to use as tests for pattern 1 and pattern 2. The real path of the Gavia can be seen in figure 8 and 9. The length of the first pattern was 152 meters and 263 meters for the second pattern. Figure 7 gives an example of one of the objects that was inspected, the figure shows image from the 20m range that is used for detection. Then the object is shown again at 5m range from the inspection mission.

If we look at the vehicle path for pattern 1 we see that there is around 1m tracking error. This is most likely caused by current in the harbor area. It is important that we tightly follow the track specified, especially when flying directly over the object for photographing. At $1.5m$ altitude the image covers about $1.5 \times 1.5m^2$ area. So a one meter offset will cause us to miss the object. This is something that needs to be improved for the dynamic search method to work.

IV. Detection algorithm

For the experiments a CAD/CAC algorithm was developed. The detection algorithm works in few phases. Filtering,

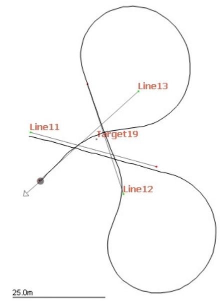

Fig. 9. Real execution of pattern 2

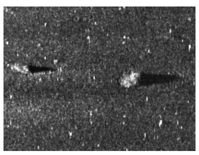

(a) Sonar Image

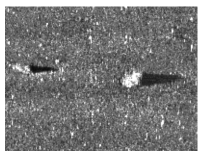

(b) Normalized

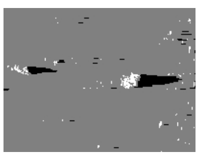

(c) Segmented

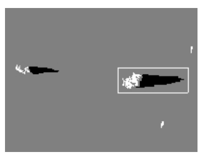

(d) Detected

Fig. 10. The object detection process

normalization, segmentation, detection, feature extraction and finally classification.

The normalization used was a Serpentine Forward-Backward Filter developed by Dobeck [8]. The segmentation was then carried out using adaptive thresholding, and separating the image into background, highlight and shadow.

Detection was implemented by grouping highlights and shadows into blocks and blocks the matched a given size where taken as detection. Then using that information several features where detected and classified by a simple rule based classifier. An example of the algorithm execution can be seen in figure 10.

V. Conclusion

We have shown that real-time detection can save time when searching for objects on the seafloor using an AUV with side-scan sonar. Real-time detection will not help in all cases, it depends on how many objects we are searching for, how hard they are to find, and the quality of the CAD/CAC algorithm.

We looked at what kind of inspection pattern could be used when we have improved position accuracy. This allowed us to use a much simpler patterns and that is one of the biggest gain of doing the on-line detection.

When doing real-time detection these results open up a host of new possibilities. We could for instance use objects that are detected as navigation aids. In the SCM missions we are executing a lawnmower pattern and often with some overlap. This leaves the possibility of us seeing the same object more than once. Position measurements of the object could be used to estimate our position for improved AUV navigation [9].

In the two-phase tactic we could use the real-time detection in the second phase. This would allow us to update our position estimate of the object so we could execute simpler patterns.

For further improvement of the search method presented it is possible to look into more detection methods. The detections methods, like for mine classification, are designed to find specific type of objects. There are no general classifiers the classify detected objects into many categories.

Another improvement is how we inspect the objects. We noted earlier that it could be possible to plan an inspection mission for all the detected objects in a image. Also tuning how much data we process at a time is an important issue. In newer versions of the sonar software used it will be possible to receive each ping. This allows for more flexibility of controlling the processing size.

The search methods previously described are not limited to detection with side-scan sonar system. Usage with a multi-beam detector is also something that could be a viable option and sub-bottom profilers have been used for buried targets.

Acknowledgment

The authors would like to thank the staff at Hafmynd for giving access to the necessary data and use of the Gavia AUV.

References

[1] C. M. Ciany, W. Zurawski, D. G.J., and D. Weilert, "Real-time performance of fusion algorithms for computer aided detection and classification of bottom mines in the littoral environment," in *OCEANS 2003*, vol. 2, Dec. 22–26, 2003, pp. 1119–1125.

[2] M. G. Norman and P. Moscato, "The euclidean traveling salesman problem and a space-filling curve," unpublished.

[3] T. Aridgides, M. Fernandez, and G. J. Dobeck, "Volterra fusion of processing strings for automated sea mine classification in shallow water," R. S. Harmon, J. T. Broach, J. H. Holloway, and Jr., Eds., vol. 5794, no. 1. SPIE, 2005, pp. 358–369. [Online]. Available: http://link.aip.org/link/?PSI/5794/358/1

[4] J. Bennett, A.A. Leonard, "A behavior-based approach to adaptive feature detection andfollowing with autonomous underwater vehicles," *IEEE Journal of Oceanic Engineering*, vol. 25, no. 2, pp. 213–226, Apr. 2000.

[5] A. Olafsson, "Leit," Master's thesis, University Of Iceland, Reykjavik, 2000.

[6] T.-C. Liang and J.-S. Liu, "A path planning method using cubic spiral with curvature constraint," Institute of Information Science, Academia Sinica, Taiwan, TR-IIS 02-006, 2002.

[7] B. Kanayama, Y. Hartman, "Smooth local path planning for autonomous vehicles," in *Robotics and Automation, 1989. Proceedings., 1989 IEEE International Conference on*, vol. 3, May 14 – 19 1989, pp. 1265–1270.

[8] G. J. Dobeck, "Image normalization using the serpentine forward-backward filter," in *Detection and Remediation Technologies for Mines and Minelike Targets X*, vol. 5794, 2005.

[9] Y. L. D. Tena Ruiz, I. Petillot, "Improved auv navigation using side-scan sonar," in *OCEANS 2003*, vol. 3, 22–26, 2003, pp. 1261–1268.