Rendezvous Point Technique for Multivehicle Mine Countermeasure Operations in Communication-Constrained Environments

AUTHORS
Veronika Yordanova
Hugh Griffiths
University College London

ABSTRACT

Advantages of using a multivehicle network over a single autonomous underwater vehicle platform include extended coverage area, potential cost and time efficiency, and more robust performance. A common issue that slows advancement in the field is the limited available communication between the platforms. The approach we propose is based on assigning a sequence of rendezvous points (RPs) where the vehicles can meet and exchange information. The work we present in this paper applies principally to mine countermeasure and suggests that, despite the disadvantage of time to allow for the vehicles to reach the RPs, there are techniques that can minimize the losses and provide advantages such as easier coordination and access points for operator monitoring and system modifications. The results we present in this paper give an estimate of the reduction in loss if such an approach is employed. We make a comparison between the RP and a benchmark case by analyzing numerical simulations.

Keywords: rendezvous, multivehicle, mine countermeasures

Introduction

There has been increased interest in the last decade in using autonomous underwater vehicles (AUVs) in a network configuration. However, the inability to maintain robust communication continuously is a major constraint for the development of such technology. Radio frequency and optical signals do not propagate well underwater, and therefore, acoustic transmissions are the most common choice. However, there are severe limitations in range and bandwidth of the acoustic communication, and the network operation must take this into account (Akyildiz et al., 2005).

The approach we propose is based on assigning dynamically a sequence of rendezvous points (RPs) throughout the mission, a location and time where all agents in the network agree to meet and the vehicles can exchange information. This way a complete lack of connection can be assumed outside of the RP perimeter, thus providing a means for the system to operate under severe channel conditions. The work we present applies mainly to mine countermeasures (MCMs). This paper presents the idea of applying RP to MCM by proposing adaptive RP scheduling. The overall goal of this new approach is to enable an adaptive reallocation of system resources to maximize search area while making explicit the rule of revisiting all contacts on the way.

The remainder of the paper is organized as follows: Section 2 gives details on different types of MCMs; Section 3 focuses on relevant work published recently on underwater networks and autonomy; Section 4 gives an overview of the RP idea; Section 5 explains the methodologies adopted for evaluating the RP method and analyzing the suitable conditions for application; Section 6 shows simulation results; Section 7 presents the analysis and limitations of the approach. The last section concludes the work and suggests future directions.

Evolution of MCM Systems

The Korean and Gulf wars are examples where effective mine warfare was applied. Warships were damaged and amphibious assaults aborted due to the inability of the navies to counter this asymmetric threat. Currently, the sheer number of existing naval mines is another reason to treat the problem as a challenging and diverse task: it is estimated that a million mines, of more than 300 types, are stored by 60 navies worldwide (this excludes U.S. weapons); mine production exists in more than 30 countries, and export is done by more than 20. These figures do not account for improvised explosive devices, which are considered affordable and relatively easy to make (Truver, 2012). The MCM problem arises from the difficulty of distinguishing between the real mines and the false alarms (FAs) due to mine-like seafloor objects (Sariel, Balch, & Erdogan, 2008), as well as from inefficient means of clearing them.

Due to the large diversity of mine types, their actuation mechanisms and means of deployment, it is hard to identify a single best method to deal with the problem of minefields. Currently, the conventional MCM approaches are sweeping and hunting. Minesweeping is used for removing mines by causing their detonation or capturing them. The design of the minesweeping vessel should be stealthy such that it does not trigger the mine itself, but instead the explosion occurs at a safe distance where the towed body with the triggering mechanism is. Minesweepers can also capture the chain or cable of moored mines, which was the predominant mine type until WWI, but not that common presently. The disadvantage of using such a sweeping technique is that there can be no assurance that the area is clear of mines. The other commonly employed technique, mine hunting, involves prior detection and classification before any neutralization action is taken. This brings the advantage of providing a probabilistic evaluation of the threat level of the area. The sensor used to detect mines is sonar, and the acquired imagery is processed by human operators to classify any contacts that could be actual mines. Once a decision is made, the object can be neutralized.

However, some countermeasures are becoming obsolete with the advancement of mine technology, and new solutions are being sought.

In minesweeping, the acoustic and magnetic vessel signature that would actuate a mine is mimicked by the minesweeper vessel in an attempt to trigger it prematurely. This is becoming less effective as modern mines rely on multiple signatures, which are not always possible to simulate all together (Truver, 2012). On the other hand, even if we disregard for a moment the complicated multiple triggering mechanism, the sweeping method does not result in any certainty that an area is mine free (Cornish, 2003). There is the possibility that a mine did not activate even though it detected a suitable target. The actuation mechanism sometimes involves randomized control that selects a target from a sequence of detections in order to avoid multiple mines being triggered by the same contact or hit only the first vessel from a convoy.

Mine hunting is considered more reliable than sweeping as at the end of the mission a level of confidence can be reached that can be input to a decision-making process on whether to drive a ship or convoy through an area (Cao & Bell, 1999). The process includes a detection, classification, and identification stage performed using imagery from sonar towed by a ship. Once a certain target is located, a neutralization unit, usually a remotely operated vehicle (ROV), is sent to dispose of it. However, there are some issues with traditional mine hunting techniques. Mine hunting ships require design with a minimal vessel signature so that it does not trigger the mines in its vicinity (Schwarz, 2014). The bigger issue that remains is that no matter how stealthy the ship is, there still needs to be people on board to control the mission. An alternative is to use a remotely operated vessel that is controlled from a safe base on the shore. While this is technology that is advancing (Benjamin & Curcio, 2004), there is still the issue of the sonar not being able to explore the contacts from close proximity. The method also relies on a single sensor that makes repeated scans over the area, which might require long mission times. The interest in using AUVs for mine hunting has been increasing in the last decade, due to cumulative work in multiple relevant fields to allow collaboration between vehicles and lowering the price of commercial hardware. Recent advances that have contributed to the area include work in communication, navigation, localization, mapping, vehicle design, underwater swarms, autonomy (surface and underwater), networking, international experiments with defense and scientific applications, etc. (Kalwa et al., 2015; Dugelay et al., 2015).

Another reason to adopt autonomous vehicles for MCM applications is that there are still some unconventional methods in use, such as sending divers or mammals to perform the neutralization and search phases. Although not that common, these techniques do exist. Mine neutralization methods conducted by human divers remain the most reliable. Mammals have more endurance and could be trained for MCM purposes, and there are several existing programs that have trained dolphins and sea lions for such a mission. However, the issue of misunderstanding between the animal and the handler exists during a mission. Overall, for both humans and mammals, the major disadvantage is risking their lives by sending them into a minefield (Brown et al., 2012).

Using AUVs for MCM gives the advantage of keeping all personnel at a safe distance by allowing for autonomous operation. When a group of networked vehicles is available, this has the potential to reduce time, cost, and efforts compared to single platforms and current conventional methods. Such a configuration could also introduce distributed and more efficient area coverage.

Related Work

Efforts to improve underwater sensor network (UWSN) performance are commonly aimed at modem development and network protocol design (Partan et al., 2007; Kong et al., 2005; Cui et al., 2006). When moving platforms are the focus of the network configuration, solutions are driven by ideas adopted by the robotics community. Examples include adapting coordination techniques, such as auction mechanisms (Sariel et al., 2008; DeMarco et al., 2011). Often such approaches do not take into account the limitations imposed by the communication in the underwater channel. Some methods for reducing reliance on communication using prediction models are also available (Sotzing & Lane, 2010); however, they also rely on anticipated environmental conditions.

The work in this paper was developed with the focus of adapting optimization techniques to the appropriate application constraints. Relevant ideas have been used for a group of networked surface and underwater vehicles to adapt their formation to the water basin borders (Kemna et al., 2015). The possibility of a group of vehicles reconfiguring their positions has also been recognized when the nodes need to adapt to unexpected conditions or to seek optimal placement (Braca et al., 2014; Yilmaz et al., 2008). However, such solutions aimed at improving autonomy often discount or neglect the issue of communication.

The idea of synchronous rendezvous has been recognized and adopted for ad hoc networks of mobile autonomous agents (Cortes et al., 2006); however, it is not a typical approach for UWSN. Although it provides a means to avoid the communication restriction by allowing all nodes to meet and plan further actions, one major drawback is that part of the resources in the system are sacrificed to allow the nodes to travel to the appointed place. To reduce the lost time, these points can be preplanned in a static sequence to guarantee optimality. This, however, introduces rigidness and lack of adaptability to external events. Therefore, we propose a rendezvous approach that allows dynamic online point allocation based on the information gathered by the vehicles and their future goals.

RP Approach

To improve the resource usage in the system, we consider a specific application and scenario to measure the loss and evaluate the significance of parameters for optimizing the RP scheduling. This paper looks into applying the RP approach to MCM and the typical operation specifics are presented. Furthermore, the loss mechanism is explained with regards to the selected scenario.

MCM Phases and Operation for AUVs

A typical mine hunting operation has five phases:

- Search: An area is scanned for mine-like objects (MLO). To secure complete coverage, often the platform is moved in a lawnmower pattern.
- Detection: Contact data are received from sensors, location is recorded, and a message of the contact and its location is created.
- Classification: A decision is made as to whether an object is a mine-like or non-mine-like. This is done by using autonomous target recognition software or a human operator. However, currently this decision is not trusted to be made autonomously.
- Identification: This determines the type of the mine so further neutralization strategies can be employed. Often, there is a long delay between Identification and Neutralization phases, due to the system lacking the ability of autonomous Classification and Identification, which means the data first have to be recovered and processed off-board at the end of the search mission and before the neutralization phase (Brown et al., 2012).
- Neutralization: A mine is considered neutralized once its location is defined so it can be avoided. In case the platform cannot evade the mine, other measures are adopted. Those can include destroying the mine, disabling its detonation ability, or disabling its ability to detect.

The work in this paper is concerned with the Search and Identify phases, without taking into account the sensor specifics or the autonomous target recognition restrictions. Instead, it focuses on optimizing the collaboration between multiple platforms. The Neutralize phase is excluded as often a ROV, rather than AUV, is used to properly guide a disabling mechanism.

A typical multivehicle MCM mission configuration includes two types of sensor packages: Search-Classify-Map (SCM) and Reacquire-Identify (RI) (Freitag et al., 2005). The SCM relies on a coarse side-scan sonar allowing faster speed during the searching phase. The RI phase makes use of a high-resolution sensor, such as multibeam sonar, that collects images for final identification and decision making. Often these two tasks are performed by separate vehicles. The SCM or search vehicle follows a lawnmower pattern, while the RI vehicle relocates the target and further examines it.

RP Loss in MCM Scenario

The limitation of using separate vehicles for distinct tasks is not necessarily due to hardware restrictions. Modern vehicles can have multiple sensors mounted on the same platform. However, coordinated position and task reconfigurability of multiple networked vehicles is still a big issue underwater. Using RPs throughout the mission enables communication between the AUVs and thus allows for dynamic task reallocation. This can aid higher resource utilization in the network. On the other hand, there is a trade-off with the time spent for the platforms to travel to RP. Evaluating and minimizing this time loss is vital for the approach to be applicable.

In order to schedule an RP, all vehicles in the system have to agree on the most convenient time and location to meet. This is done multiple times throughout the mission, and the intervals between RPs are adapted based on the number of contacts found. We divide the RP scheduling into two stages: (1) select the time and (2) select the coordinates of the RP. To minimize the vehicle traveling time toward the RP, we want to minimize the total number of RPs and push the next point as far in time as possible. On the other hand, we want to give the system nodes regular chances to adjust future strategy and provide updates to the operator. It is useful to define a suitable time interval that satisfies the opposing demands by showing when the resource loss becomes prohibitive for the mission.

The resource loss per vehicle is defined as the time spent by each platform to travel to the RP instead of doing mission-related task. That is, travel time from the point when the vehicle stops its search function and goes to the RP. Equation 1 gives this relation by calculating what fraction of the time window between RPs is spent for reaching the RP:

$$\text{loss} = \begin{cases} \dfrac{x}{v \times t_{\text{RP}}} \times \dfrac{1}{2}\left(1 - \dfrac{1}{n}\right), n > 1 \\ \dfrac{x}{v \times t_{\text{RP}}} \times \dfrac{1}{2n}, n = 1 \end{cases} \tag{1}$$

where x is the width of the search area, n is the number of vehicles (the cases and transformations relevant to n are explained further), v is speed, and t_{RP} is the time until the next RP. An important note is the reasoning behind the choice to evaluate the loss per single vehicle and per single RP interval. First, a single platform loss gives flexibility for reconfigurability throughout the mission. Second, the single RP interval evaluation, as opposed to total mission time, comes from the fact that all RP intervals vary, as they are a function of the number and location of contacts found during the search phase. Therefore, the resulting loss for adopting the RP approach will be additive, but nevertheless each time interval will be unique.

In order to apply Equation 1, some assumptions are made: the area is searched sequentially; the vehicles are homogeneous and have the same speed. Since the resource loss calculation is very dependent on the geometry of the search area, Figure 1 gives a graphical representation of the scenario considered in this paper for evaluating the RP approach. After defining the favorable parameters for minimizing the loss, some conclusions can be drawn on when this approach might not be applicable due to prohibitive losses.

The mission scenario considered for the remainder of the simulations in this paper is demining a strip near the shore. Typically, the width would be much smaller compared to the length of the area. Another assumption is that the vehicles will not be able to cover the whole area before their batteries are exhausted. The overall purpose of the mission is to explore as much of the area as possible, while having the constraint of revisiting the detected MLOs.

In Figure 1, x and y are the dimensions of the searched area in the simulated scenario. The red circle on the top left side pinpoints the starting position of three available vehicles. They are all equipped with search (side-scan sonar) and RI (multibeam) sensors and can perform both SCM and RI tasks. The first RP is predefined and the speed of the vehicles is known, so this can give a good idea of the location of the next meeting point as well as the area that will be covered by each platform. Since the vehicles are homogeneous, all three search areas will be equal ($area_{veh1}$ = $area_{veh2}$ = $area_{veh3}$).

The right-hand side of Figure 1 gives an example of the continuation of the mission after the first RP. At the RP, all vehicles have shared the location of the contacts they have encountered. The path and the required time to revisit them have been calculated. A decision has been made that one vehicle will be reallocated with an RI task and follow all known contacts (the route is drawn in red in Figure 1) and two platforms will continue in search mode, but now with changed area patterns.

Given the scenario from Figure 1, the loss calculation from Equation 1 can be further clarified. Normally, the AUVs would search in a lawnmower pattern, as shown in the first vehicle's search box (*'veh 1'*, left-hand side of Figure 1). At the time of the RP, this will bring the AUV either at the far or close corner of its search

FIGURE 1

Rendezvous point approach: Left—starting point to first RP. Right—example scenario between two RPs. (Color version of figures are available online at: http://www.ingentaconnect.com/content/mts/mtsj/2016/00000050/00000002.)

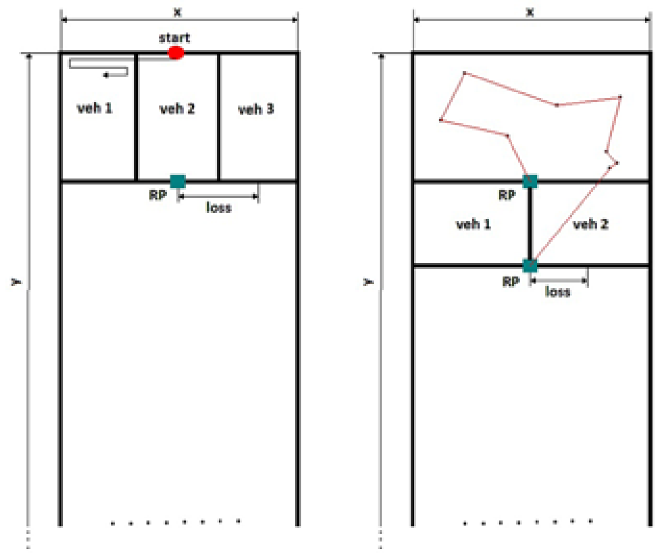

box, which will define the distance and the subsequent resource loss to the RP. However, the simulation is not optimized to always position the nodes at the near corner. To account for this, the function in Equation 1 is adjusted to be proportional to the number of vehicles performing the search phase. The special case when $n = 1$ is required as otherwise it would result in loss = 0. The current penalty for this case resembles the loss of $n = 2$. With this correction, the calculation always assumes the distance between the middle point of the platform's search area (on the y axis, same as where the RP is) and the RP point (this distance is noted on both sides of Figure 1). The distance penalty is proportionate and increases when increasing the number of search platforms. The loss calculation also depends on the speed of the platforms—the faster they are, the less time is wasted to travel to the RP. And lastly, the time to the next RP defines what fraction of the total time will be used for task-related purposes and what part will be traveling to the meeting point.

It is obvious that the longer the time between the rendezvous, the smaller fraction of the time will be lost in traveling there. On the other hand, there is a limitation on this time depending on how often the operator would need an update from the network. To evaluate when the loss becomes prohibitive or to define the minimum time for RP, we have parameterized Equation 1, and the result is shown on Figure 2. The graph presents the loss for a group of two and three vehicles, respectively, for different speed values, v, and width of the searched area, x, while the RP time interval and the length of the search area are kept constant ($RP = 1$ h, $y = 5000$ m). Then, in Figure 3, we have selected favorable, but realistic, conditions—3 vehicles, speed, 2 m/s and width of 2,500 m. The plot shows what percentage of the total time is lost if we vary the time of the RP. This can be used as a rule of thumb guidance of the minimum time limit of the next RP.

In the scenario we have selected in Figure 3, RP between 1 and 2 h gives a loss between 5% and 12% of the total time. If this is considered unacceptable by the operator, the time window can be moved further in time. The calculations in the next sections adopt the parameters used in Figure 3 ($n = 3$, $x = 2{,}500$ m, $v = 2$ m/s) and define an RP interval of 1 h assuming there is no other parameter to base the decision on. The scenario from Figure 1 is used throughout all simulations for the remainder of the paper.

RP Scheduling and Analysis

Defining the lower boundary for scheduling RP is useful when there is no other information available. However, the main advantage of the RP approach comes from the ability to reallocate vehicle tasks based on the information they have gathered during the search.

For example, if we assume the scenario from Figure 1, on the right, but without utilizing the RP technique, at the start of the mission two vehicles will be performing search tasks and they will send the locations of the detected MLOs to the third vehicle that is tasked with RI. However, if there is no prior information about what number of targets to expect, the mission can be completed with few detected MLOs or with a very large number. In the former case, the RI vehicle will be idle most of the time; in the latter, the mission will be incomplete with targets observed only with low-resolution sensor. This scenario was assumed as a benchmark case and adopted in all

FIGURE 2

Loss calculation for two vehicles (lines) and three vehicles (dots): different speed values and width of search area (x), fixed RP time (1 h), and fixed length of the y dimension of searched space (5,000 m).

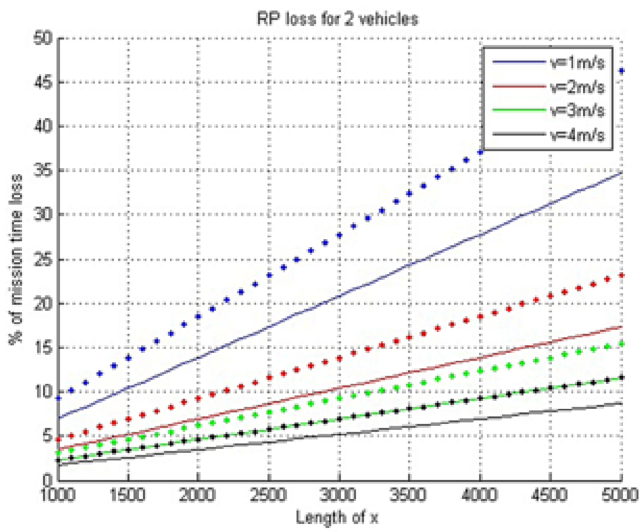

simulations throughout the remainder of this paper for comparison with the RP approach.

A higher amount of the system resource could be utilized if the vehicles are retasked adaptively according to the number of detected contacts. During an RP, a decision is made whether there is a need to retask a search vehicle into an RI vehicle and vice versa based on how many detections were made during the previous RP interval. Therefore, the number of detected targets or the time it takes for the RI vehicle(s) to revisit them is the main parameter that decides how to utilize the available vehicles as well as how to schedule the RP:

$$t_{RP} = f\left(t_{RI} + t_{reaquire}\right), [\min, \max] \quad (2)$$

where t_{RI} is a calculation of the time required for a single vehicle to travel in an optimal path and visit all detected MLOs to identify them. Additional time for reacquiring the target locations once the AUV is at the reported coordinates and collecting high-resolution sensory information was added by the $t_{reaquire}$ term. For more realistic calculation, this parameter needs to be adjusted depending on the mission conditions and the platform sensors. The [min, max] interval is derived by using Equation 1 and Figure 3.

The second stage of scheduling the next RP is the position estimation of the vehicles at the time when they will be advancing toward it. This calculation is based on predicting the vehicles' positions at the next t_{RP},, or the area coverage of the search AUVs. The $Area_{search}$ parameter in Equation 3 gives the area that will be covered by the search vehicles in the time interval between RPs. This calculation is performed during the meeting at RP, after the following RP time is defined by Equation 2. The position of the search vehicles coincides with the y-axis from Figure 1, where the RP is positioned.

$$Area_{search} = v \times sw \times t_{RP} \times loss \times n \quad (3)$$

where v is speed of a vehicle, sw is swath width of the vehicle's sensor, t_{RP} is time until next RP, loss is the resource waste for traveling to RP using Equation 1, and n is number of AUVs. However, since all platforms are assumed to be of the same type, at time

FIGURE 3

Loss of time resource (y axis) vs. total time between RP (x axis)—calculation based on Equation 1 (parameters used: n = 3; x = 2,500 m; v = 2m/s).

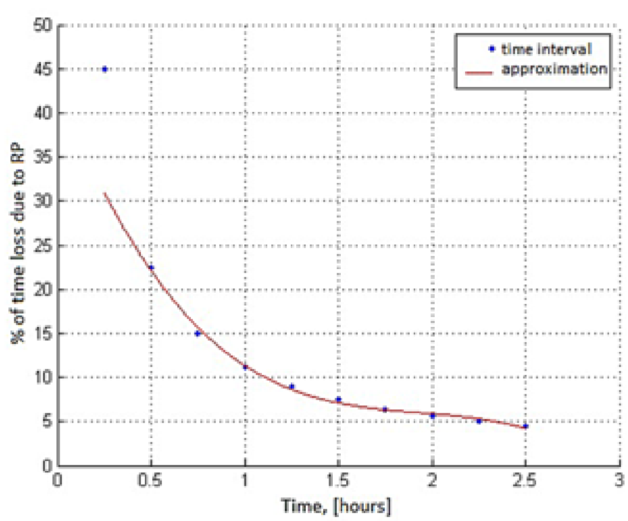

t_{RP}, all nodes will be aligned at the same y coordinate in the search space. Therefore, the location of the RP is the middle of the width of the search area at this y coordinate, as seen in Figure 1. Equation 3 accounts only for the positions of the search vehicles. The RI vehicles follow the shortest path between the contacts they will be revisiting and reach back to the current RP. The advancement over the y-axis made by the search vehicles has not been considered for the RI vehicles.

These relations were used to explore the gains of network reconfigurability by applying the RP approach. The steps used in our simulation are summarized in the pseudo code in Algorithm 1, as well as the rule-based decision making for how many platforms to employ search and RI tasks.

At line 3, the first RP is predefined. It is assumed that no prior information is available at the start of the mission; hence, the RP time is solely driven by the loss calculation or given the minimum value from Equation 2. The simulation runs until the mission_time or the sum of the RP time windows exceeds a predefined threshold. This threshold was selected as a percentage of typical battery capacity of an AUV —70% of 10 h (line 7). Depending on the selected parameters, the average resource loss and search area are calculated (lines 8 and 9). Once all platforms reach the first RP, they will share information about the detected targets. Communication at RP is assumed available. The MLOs are simulated by generating a random number of targets, limited in number, and with locations constrained within the area that have been searched in the time window (line 10). The shortest path to revisit them is then calculated (line 11).

ALGORITHM 1

Dynamic allocation of vehicle tasks and RP.

```
1: procedure SCHEDULE RPs
2:     min_int ← min loss limit from Graph 4
3:     t_RP = min_int ← assign time for first RP
4:     n ← number search vehicles
5:     targets = 0
6:     mission_time = t_RP
7:     while mission_time < threshold do
8:         Calculate loss using Equation 1
9:         Calculate search area using Equation 3
10:        targets = +new_targ ← generate MLOs for RI (applied limit t_RI < min_int*3)
11:        t_RI ← calculate RI time (Nearest Neighbor)
12:        if (t_RI > min_int * 3) then
13:            break
14:        else
15:            if t_RI ∈ [min_int; min_int * 2) then
16:                t_RP = t_RI
17:                n = n − 1
18:                targets = 0
19:            else
20:                if t_RI ∈ [min_int * 2; min_int * 3) then
21:                    t_RP = t_RI/2
22:                    n = n − 2
23:                    targets = 0
24:                else
25:                    t_RP = min_int
26:                    targets = new_targ
27:                end if
28:            end if
29:            mission_time = +t_RP
30:        end if
31:    end while
```

The decision making on when to schedule the next RP and how many vehicles to send is given between lines 12 and 30. Simple rule-based logic is used. If the RI time falls in the interval between *min_int* and *2*min_int* (line15), the time for the next RP is selected as the time it will take an RI vehicle to check all targets, the search vehicles are reduced by one and the targets in the searched area are considered identified. In case the time for the RI task exceeds the *2*min_int* limit – line 20 (due to too many targets), then two vehicles are tasked to perform the RI task and the time for the next RP is half of the RI time. The number of search vehicles is reduced by two, and also all targets are considered identified. In the case when the RI time is below the *min_int* threshold, it is better to leave the targets unidentified until the next RP cycle. The advantage is that all vehicles will perform search rather than one vehicle tasked with RI and then being idle for portion of the time. The disadvantage is that, on the next cycle, when there are enough targets to identify, the path to travel will be longer. In the next section, this trade-off is explored further. Lines 12 and 13 ensure that the simulation will be terminated in the case that more than two RI vehicles are required to identify the contacts from the previous RP interval.

Results

The function from Algorithm 1 was used in a MATLAB simulation to evaluate the performance of the RP approach when variable numbers of MLOs are detected. This assumes the case where no prior intelligence is available about the expected number of contacts, and thus, the mission operators would be unable to manage the resources in the system offline.

The aim of this work is to show how applying the RP method gives the opportunity for the system to adapt to an unexpected and varying need for retasking the vehicles. On the other hand, a benchmark scenario was designed where the vehicles' functionality is predefined, as would be the case if no autonomy was implemented. Our expectations were that the results would show such architecture makes the system rigid and inefficient. To represent this base case, at the start of our simulated missions, from the total of three vehicles, to match the RP case, two AUVs were designated to perform search and one was tasked with identifying a detected contact. The reason to put more vehicles in search was that the objective of the mission is to maximize the overall search area. The advantage here is that there is no loss introduced from regularly traveling to a meeting point as when the RP method is applied. The benchmark case assumes that the vehicles have the ability to broadcast target locations continuously throughout the duration of the mission with no collision or message loss.

Figure 4 shows three graphs, each representing simulation results with different numbers of detected targets. At the top is a scenario with low number of targets (0–10 generated targets per RP window), the middle graph shows average number of targets (0–20 per RP), and at the bottom is a large number of simulated contacts (0–30 per RP). The choice of target number intervals is somewhat arbitrary, as this parameter is not based on literature or experiments. Since this variable is hard to determine, we have adopted these intervals based on the load that would be generated within the assumed system resources. Thus, each graph shows results for 100 repetitions of Algorithm 1 where the randomization of number and location of MLOs accounts for variability in losses and gains in the system.

The graphs in Figure 4 show the total area searched by the vehicles (depicted on the y axis) throughout the available mission time (shown in minutes on the x axis). The scattered black dots are full mission simulations with the RP approach applied, while the red dots are simulations with designated vehicles tasks or the benchmark case. It is essential to clarify that the mission time per simulation is defined by the RP approach, as described in Algorithm 1 on line 7, and then this limit is applied to the benchmark case, where the same input of target distribution and number is used to calculate the overall area searched by this architecture. This creates a pair of RP and benchmark case simulations that give a comparable output as all inputs are the same. In Figure 4, there are 100 pairs in each graph, which allow an evaluation of the loss or gain of applying the different methods, even though multiple variables are present in the simulations, such as number and location of the targets, as well as mission duration.

FIGURE 4

Comparison between the RP and benchmark approaches applied to MCM mission by plotting the search area gained by each method. The three graphs show how the results change when the number of targets increases: top graph gives a simulation with low number of targets (0 to 10), middle one doubles the targets (0 to 20), and bottom graph shows a cluttered environment (0 to 30 contacts per RP window).

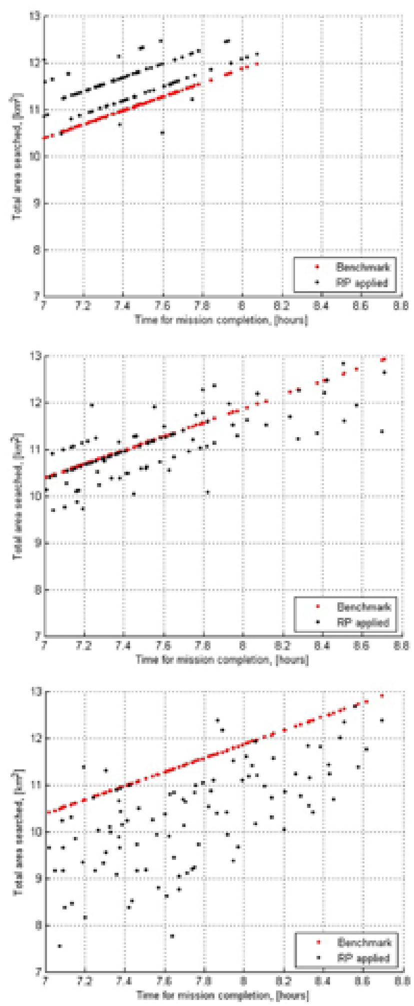

It can be seen on the graphs that in all cases the red dots follow a linear relationship—the more mission time available, the larger area is searched. On the other hand, the black dots are scattered below and above this red line, showing lower resource utilization if they are below the red line and higher if they are above.

As expected, in the example of fewer targets (top graph in Figure 4), most of the simulations show better performance for the RP approach, which allows for function reconfigurability, compared to the case of predefined tasks for the AUVs. The reason is that the RI vehicle from the benchmark mission is being idle for the majority of the time, as there are not enough targets to fill its schedule. On the other hand, the third vehicle in the RP simulations is performing a search function during most of the RP windows as there were not enough targets found to justify the RI task. It can be observed that with increasing the number of targets, the advantage of the RP is lost due to multiple vehicles tasked with RI throughout the mission. However, in the case of a cluttered environment with many contacts present, the vehicles from the benchmark case are not able to revisit all detections as only one vehicle is tasked with RI, and its resources are not enough. The comparison between the two approaches in their ability to revisit and identify the contacts has been analyzed further in this paper, and simulation results are available in Figure 7.

Figure 5 compares the distribution of the search resource in the RP and benchmark case. The y axis gives the average number of search vehicles utilized per simulation, and the x axis provides the simulation number. Since the base case is constantly using two search vehicles, regardless of the mission circumstances, there are 100 equally spaced black dots at $y = 2$ for every simulation number on all three graphs. When RP is applied, the search platforms change their number at every RP window depending on the available MLOs, as sometimes the RI vehicles can account for 0, 1, or 2 of the total number of vehicles (Algorithm 1, lines 15–28). Each data point from the blue lines in Figure 5 corresponds to the average number of search vehicles utilized throughout each simulation.

FIGURE 5

Comparison between the RP and benchmark approaches by giving the average number of search vehicles used per mission. The graphs show the decreasing availability of search resource (top to bottom) in the RP case by increasing the number of detections in the simulations.

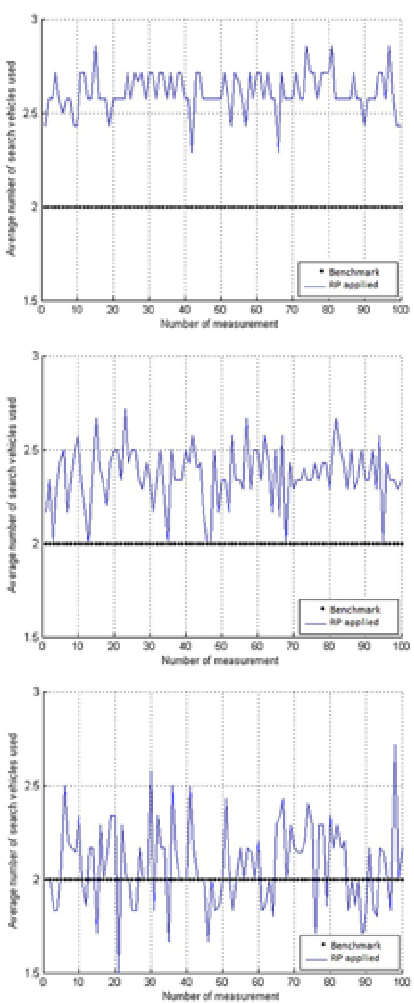

This result can be correlated with the graphs in Figure 4, explaining the higher search area covered in the case when less vehicles are tasked with RI (top graphs in both figures where less targets are found) and thus the average search platform number is much higher. The bottom graph in Figure 5, where the RP network of vehicles had to adapt to a large number of targets, shows that in a proportion of the simulations less than two vehicles on average were available for the whole mission. This can explain the majority of black scattered dots moving toward the bottom right corner in the third graph of Figure 4. This is the undesirable situation when more time spent at the MCM mission yields less area searched.

On the other hand, even in the unfavorable event of encountering many targets, the RP guarantees that all detections are revisited. The approach was built with this core objective in mind and it defined a condition driving the whole decision process in Algorithm 1. For the benchmark case, some contacts were not revisited as the RI vehicle resource would not be enough. This is another advantage for the base scenario, together with the unlimited communication mentioned earlier, which makes its results to look more favorable. However, the objective to provide high-resolution images of all MLOs is violated in the benchmark case and thus making the RP approach show more pessimistic results than it would in a fair comparison. The RI success rate for both methods is further analyzed in the next section.

Analysis

In order to give a more suitable representation of the outcomes showing the efficiency gains and losses of the RP method, Figure 6 displays the same data as used to present results in Figure 4, but in a different format. To achieve a better visualization to distinguish between the different approaches, the benchmark case data points were rotated and translated to coincide with the $x = 0$ axis and plotted as a red line. Then, the difference between the area searched by using RP and benchmark was calculated. This emphasized the discrepancy in overall area between the RP and base case pairs. To capture the dynamics of this variation, these pairs were further sorted in a descending order, which resulted in the blue plotted line on all graphs in Figure 6. Such result representation is easier to evaluate by clearly differentiating how likely it is for the RP approach to be beneficial. Above the red line or the zero of the y axis, the RP approach gains additional search area even though the vehicles waste time resource for multiple meetings. If the blue line is below the red one, then the simulation instance results in a loss of search area.

In the top graph of Figure 6, showing results from a search with low

FIGURE 6

Normalized and sorted simulation results comparing the gain of search area achieved by applying RP and base case. The top graph gives a simulation with low number of targets (0 to 10), middle one doubles the targets (0 to 20), and the bottom graph shows a cluttered environment (0 to 30 contacts per RP window).

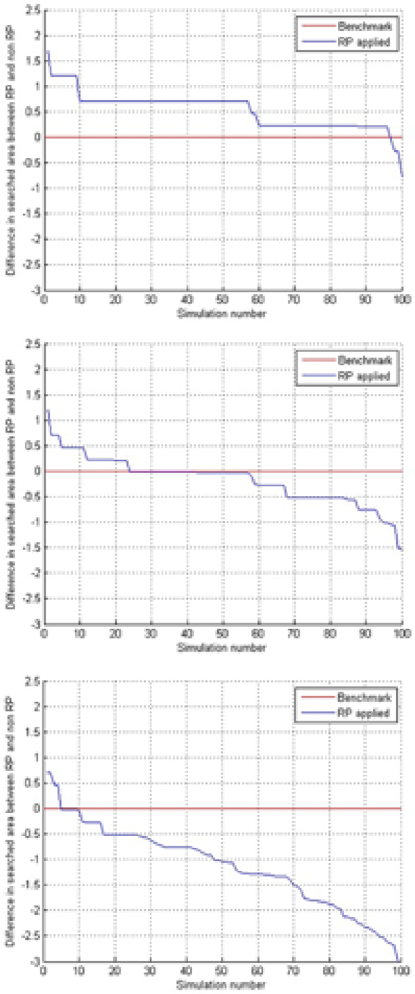

FIGURE 7

Normalized and sorted simulation results comparing the time required to reacquire and identify (RI) all detected contacts for applying RP versus the benchmark case. The top graph gives a simulation with a low number of targets (0 to 10), the middle one doubles the targets (0 to 20), and the bottom graph shows a cluttered environment (0 to 30 contacts per RP window).

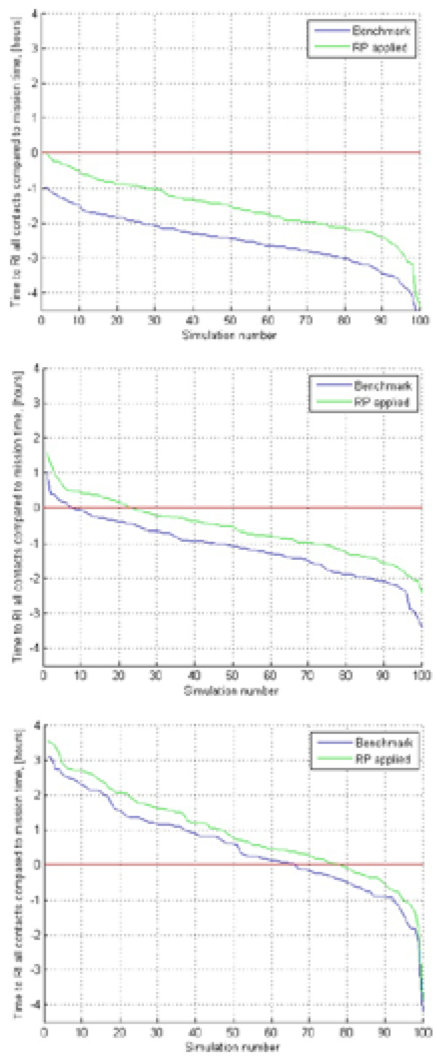

number of contacts, the RP gives a significant gain over the benchmark case in the majority of the simulations, even though the loss for vehicles meeting is accounted for. It is obvious that this is the most advantageous situation to apply RP.

The middle graph gives the resulting area search difference when doubling the number of targets from the previous case. Now it is visible that the majority of the blue line is under the red line; i.e., most simulations indicate a loss to the system when applying the RP. However, a large part of the simulation points fall very near the red line. In about 40% of the simulations, the results are within 0.2 km^2 of total search area. This accounts for about 3 to 4 min of a single vehicles operation time, according to Equation 1. This shows that, half of the time, the RP approach is almost indistinguishably as good as the benchmark case. However, it adds the operational advantages discussed in Section 4 for system monitoring and periodic data gathering.

The final graph presents a state when there is an extensive number of contacts simulated during the mission and shows the expected undesirable results from the RP approach—a lot of resources would be spent on relocating these contacts, and this will add to the time loss brought from the regular meetings. In reality, most of the contacts would be FAs. Environments that are characterized by many FAs in the vehicle sensor, such as rocky sea bottom, are unfavorable for the RP approach. It can be seen on the graph that only very few data points have the blue line above the red, and in the majority of the cases this approach brings loss. In addition, the loss propagates to significant difference in search area, in the order of 1.5 km^2 and above, for the 30% of the cases.

So far, only the search resource in the system has been analyzed. Figure 7 gives an insight how the benchmark and RP approach spend their RI resources. The x axis on all graphs is the simulation number and the y axis measures the time it would require to spend in RI phase per mission. In order to make the comparison clearer, the results of the 100 repeated simulations were sorted in a descending order. The red line is the mission time (different for each iteration of the algorithm). The blue line follows the benchmark results, and the green line follows the corresponding RP results. It is clear that the benchmark case always uses less resources compared to when RP is applied (the blue line is always below the green line). This comes from the fact that the RI vehicles in the RP simulations have to go back to each contact location only after its route is defined at the meeting point, which adds a time overhead. On the other hand, the benchmark case assumed perfect communication, so the two search vehicles can instantly send their contact locations and the RI vehicle can select the shortest path without delay.

In all the cases when the benchmark is above the red line or the limit of the mission time, some contacts remain unidentified, as the RI resource is not sufficient. This is evident in the bottom graph where the majority of the simulations result in contacts requiring extra RI time. On the other hand, the RP approach adaptively reallocates part of its resource from search into RI, resulting in more than one RI vehicle per mission on average, as can be seen in Figure 5. Therefore, all contacts are identified by the end of the mission. This flexibility gives an advantage to the RP approach that can surpass the overall loss of time to meet and the reduced total search area.

A disadvantage of the RP in relation to the RI resource is that the contacts are always revisited after a meeting point. This results in a delayed reaction, and the current simulation does not allow for the contacts from the last RP interval to be revisited (as the algorithm terminates once the mission time reaches the selected threshold—line 7 in Algorithm 1). This is an oversight of the approach, but it does not violate the conclusions made in this paper, as both the benchmark and the applied RP simulations disregard these contacts. In a real MCM or experimental setting, where revisiting the contacts is crucial, this can be easily amended by setting the last RP interval to force each search vehicle to perform the RI task for its contacts before it goes back to the mission end point. This, however, would not contribute to the current evaluation and was not included in the simulations.

Conclusion and Future Direction

This paper has presented the idea of adaptive scheduling of RPs for MCM application with AUVs. The benefit of using RP is that the vehicles can be utilized at a constant rate independent of the number of targets detected throughout the mission. In contrast, when adopting a typical configuration where the functions of search and ID vehicles are separated from the start of the mission, the overall time to achieve the same result could be significantly increased if there are large numbers of MLOs or the ID vehicle could be underutilized if their number is low. The conventional approach thus becomes less efficient compared to the presented RP method.

This work can be improved by relaxing or refining some of the assumptions which forced pessimistic calculations, such as adjusting the position of the search vehicle when calculating the loss and the time required for reacquiring a target. Another important change would be allowing different kinds of vehicles with varying speed and sensors that will add diversity to the simulation and make it more realistic. Finally, relaxing the assumption of sequential search allows for the use of a probabilistic model of selecting which area to be given priority for exploration and which one needs a repeated coverage. This also calls for more sophisticated decision making, such as using Markov processes.

Acknowledgment

This project is funded by Atlas Elektronik UK and EPSRC Doctoral Training Centre no: EP/G037264/1. The authors would like to thank A. Gibbert, A. Charlish, and R. Brind for being available and providing early feedback to this work.

Authors

Veronika Yordanova and Hugh Griffiths
Department of Electronic and Electrical Engineering
University College London
Torrington Place,
London WC1E 7JE, UK
Email: veronika.yordanova.11@ucl.ac.uk; h.griffiths@ucl.ac.uk

References

Akyildiz, I.F., Pompili, D., & Melodia, T. 2005. Underwater acoustic sensor networks: research challenges. Ad Hoc Netw. 3(3): 257-79. http://dx.doi.org/10.1016/j.adhoc.2005.01.004.

Benjamin, M.R., & Curcio, J.A. 2004. COLREGS-based navigation of autonomous marine vehicles. In: Proceedings of Autonomous Underwater Vehicles. pp. 32-9. Sebasco, ME: IEEE/OES.

Braca, P., Goldhahn, R., LePage, K.D., Marano, S., Matta, V., & Willett, P. 2014. Cognitive multistatic AUV networks. In: 17th International Conference on Information Fusion (FUSION), 1-7. Salamanca: IEEE.

Brown, T., Damiano, J., Jhala, S., Moore, R., Morgan, B., Nguyen, V., & Turk, J. 2012. Next generation mine countermeasures for the very shallow water zone in support of amphibious operations (No. NPS-SE-001). Naval Postgraduate School Monterey CA Department of Systems Engineering.

Cao, P., & Bell, M.J. 1999. New methods to evaluate the detection performance of a minehunter (No. DSTO-TN-0200). Defence Science and Technology Organisation Canberra (Australia).

Cornish, G.J. 2003. US Naval Mine Warfare Strategy: Analysis of the Way Ahead. Army War Coll Carlisle Barracks PA.

Cortes, J., Martínez, S., & Bullo, F. 2006. Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions. IEEE T Automat Contr. 51(8):1289-98. http://dx.doi.org/10.1109/TAC.2006.878713.

Cui, J.H., Kong, J., Gerla, M., & Zhou, S. 2006. The challenges of building mobile underwater wireless networks for aquatic applications. IEEE Netw. 20(3):12-8. http://dx.doi.org/10.1109/MNET.2006.1637927.

DeMarco, K., West, M.E., & Collins, T.R. 2011. An implementation of ROS on the Yellowfin autonomous underwater vehicle (AUV). In: OCEANS 2011, 1-7. Kona: MTS/IEEE.

Dugelay, S., Williams, D.P., Okopal, G., Connors, W.A., Midtgaard, Ø., Sæbø, T., … Hesby, A. 2015. MANEX'14: Experimental description and preliminary results. Tech. Rep.: NATO STO Center for Marine Research and Experimentation, La Spezia, Italy: CMRE-FR-2014-013.

Freitag, L., Grund, M., Von Alt, C., Stokey, R., & Austin, T. 2005. A shallow water acoustic network for mine countermeasures operations with autonomous underwater vehicles. In: Undersea Defence Technology (UDT), Amsterdam: UDT. pp. 1-6.

Kalwa, J., Pascoal, A., Ridao, P., Birk, A., Glotzbach, T., Brignone, L., & Bibuli, M. 2015. EU project MORPH: Current status after 3 years of cooperation under and above water. IFAC-PapersOnLine. 48(2): 119-24. http://dx.doi.org/10.1016/j.ifacol.2015.06.019.

Kemna, S., Caron, D.A., & Sukhatme, G.S. 2015. Constraint-induced formation switching for adaptive environmental sampling. In: OCEANS 2015, 1-7. Geneva: MTS/IEEE. http://dx.doi.org/10.1109/oceans-genova.2015.7271361.

Kong, J., Cui, J.H., Wu, D., & Gerla, M. 2005. Building underwater ad-hoc networks and sensor networks for large scale real-time aquatic applications. IEEE Military Communications Conference, MILCOM 2005. 3:1535-41. http://dx.doi.org/10.1109/MILCOM.2005.1605894.

Partan, J., Kurose, J., & Levine, B.N. 2007. A survey of practical issues in underwater networks. ACM SIGMOBILE Mobile Computing and Communications Review. 11(4): 23-33. http://dx.doi.org/10.1145/1347364.1347372.

Sariel, S., Balch, T., & Erdogan, N. 2008. Naval mine countermeasure missions. IEEE Robot Autom Mag. 15(1):45-52. http://dx.doi.org/10.1109/M-RA.2007.914920.

Schwarz, M. 2014. Future mine countermeasures: No easy solutions. Naval War Coll Rev. 67(3):123.

Sotzing, C.C., & Lane, D.M. 2010. Improving the coordination efficiency of limited communication multi-autonomus underwater vehicle operations using a multiagent architecture. J Field Robot. 27(4):412-29. http://dx.doi.org/10.1002/rob.20340.

Truver, S.C. 2012. Taking mines seriously. Naval War Coll Rev. 65(2):1-37.

Yilmaz, N.K., Evangelinos, C., Lermusiaux, P.F., & Patrikalakis, N.M. 2008. Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. IEEE J Oceanic Eng. 33(4):522-37. http://dx.doi.org/10.1109/JOE.2008.2002105.