

FRACTIONAL BANDWIDTH REACQUISITION ALGORITHMS FOR VSW-MCM

Benjamin Cook, Daniel Marthaler, Chad Topaz, and Andrea Bertozzi
Department of Mathematics
Center for Nonlinear and Complex Systems
Duke University
Durham, NC 27708

Mathieu Kemp
Nekton Research
Durham, NC 27713

Abstract Autonomous underwater vehicles are gradually being recognized as key assets in future combat systems.¹ Central to this attitude is the realization that teams of vehicles acting in concerted fashion can accomplish tasks that are either too costly or simply outside the range of capabilities of single vehicles. The VSW-MCM target reacquisition problem is the primary driver of underwater multi-agent research.² Because of the VSW's inherent high vehicle attrition rate and unreliable communication, it is felt that vehicle coordination must be done off-site. In this paper, we suggest an alternative to this which permits on-site coordination despite loss of vehicles and communication. Because of its generality, this approach might also be valuable in land-based and aerial applications.

1. INTRODUCTION

The underwater mine countermeasure (MCM) problem refers to the cleanup of an underwater landing lane in advance of a sea-based military operation.² MCM is currently accomplished by diver-dolphin teams. Because it is extremely hazardous to personnel (low temperature, heavy gear, proximity to explosives, and obscurity), one of the objectives of the US Navy is to replace the diver-dolphin teams by unmanned underwater vehicles (UUV).³

The UUV-based MCM problem is normally partitioned into three phases and three regions: detection / reacquisition / neutralization, and shallow water

/ very shallow water / surf-zone. Penetration into the shallow water region is well understood. Current Navy efforts are focused on the very-shallow water (VSW) region.⁴

In this article, we focus on the VSW-reacquisition problem. Simply stated, the problem is to visit as quickly as possible a large number of near-shore targets of known location. The obstacle to achieving this are the harsh conditions inherent to the VSW: rock, kelp, surge current, unstable ground, low visibility unreliable communication, etc.² The rate of UUV attrition is expected to be high, at least 20% per 8 hours.³ Current military specifications require coverage of 100 square nautical miles in under 3 days, with a launch point a few kilometers from the field of operation.³ Given existing small-UUV technology (typical coverage rates are of order 0.02 square nautical mile/hour/vehicle and operational duration of 8 hours⁴), this operation is expected to require the coordination of 75 vehicles.

One of the scenarios under consideration uses centralized control from a remote platform.³ Vehicles are launched from the remote platform with a list of targets to visit. After visiting those targets, they return and update the central node's target registry. While the vehicles are being recharged, the control node computes new target lists, which are then downloaded to the vehicles. The cycle repeats until the mine field has been fully visited. The central node adjusts for vehicle attrition by keeping track of losses and assigning the missed targets to other vehicles. The only penalty for attrition is the smaller rate of target reacquisition. This approach is free from the unreliability of VSW communication, because all coordination is handled remotely.

In this article, we suggest a complementary approach to this view, which we call "fractional bandwidth" algorithms, that achieves on-site vehicle coordination while being robust against loss of vehicle and communication. The value of this approach is particularly high when the attrition rate is large, or when centralized control is not available. The article is organized as follows: Section 2 introduces this approach and Section 3 examines its performance. The results are summarized in Section 4.

2. FRACTIONAL BANDWIDTH APPROACH

2.1 VSW-MCM model

We start with some definitions and a description of the VSW-MCM model. The target reacquisition problem can be stated as follows: Following a target identification sweep by a separate vehicle, T target locations are identified for reacquisition by a team of N UUVs. The objective is to visit

these targets, identify their type, communicate these results, and to do so in time-optimal manner.⁵

The target field geometry is a square of side L . The vehicles are modeled as point objects, moving at constant speed while approaching a target, and stationary while classifying it. The mean time between failure of individual vehicles is modeled with an exponential random process. The unreliability of inter-vehicle communication is modeled with a finite communication range, i.e. vehicles within the communication range hear the transmission otherwise they do not. The vehicles are initially launched at random locations within the target field.

2.2 Philosophy of Approach

We start with two algorithmic observations that motivate our approach.

2.2.1 On-site centralized coordination. We consider an algorithm where a team of vehicles are under the direct control of an on-site central node (in contrast to the approach discussed in the Introduction where the central node is off-site). The central node knows where all targets are, what targets remain to be visited, where all vehicles are, and what their health is. It continually determines the optimal deployment of the remaining vehicles, and communicates those requests. This algorithm is robust against loss of vehicles, as long as one vehicle remains, and can be made time-optimal; its weakness however is that it breaks down when communication is lost, either because vehicles fall outside range of the central node, because of environmental effects, or because of a central node failure.

2.2.2 Parallel execution. The other extreme is an algorithm where each vehicle is assigned a pre-mission plan (i.e. a list of targets), then sent on its way with no further need for coordination (again, this is different from the approach discussed in the Introduction). This algorithm is by construction robust against communication loss, however not so against vehicle loss.

These two extremes show that methods for on-site coordination cannot rely entirely on communication, yet cannot be free from it.

The fractional bandwidth approach that we propose is based on this observation. It is based on two principles:

- (1) each vehicle must individually be capable of accomplishing the entire mission.
- (2) communication is a performance accelerator.

Next, we present three examples that illustrate these principles.

3. FRACTIONAL BANDWIDTH PERFORMANCE

These principles can be used to augment the performance of the remote coordination VSW-MCM methodology discussed in the Introduction, where each vehicle is assigned a set of non-overlapping targets within a sector. We propose three variants, where instead of being independent while on-site, the vehicles put their targets in common, and coordinate their motion to reacquire them. As we show, this increases the robustness of the mission against vehicle loss, yet is robust against loss of communication.

3.1 Greedy algorithm

The structure of the greedy algorithm is the following:

1. each vehicle keeps a map of all targets, i.e. their position and a status flag. The status flag takes one of three values: "available", "selected", and "done". All targets are initially "available".

2. each vehicle looks at its map and selects the nearest "available" target. If no targets are "available", the vehicle selects the nearest "selected" target. If none are left, the vehicle ends its mission.

3. after target selection, each vehicle broadcasts its choice. Vehicles within communication range receive this data and update their map as follows: if the current value of the received target's status is "available", it is upgraded to "selected"; if it was "selected", it remains "selected"; if it is "done", it remains "done". This step is crucial to coordination, as two vehicles cannot knowingly select the same target (unless all targets have already been selected).

4. each vehicle travels to its target; once within the search radius, it executes the classification.

5. when done with its target, each vehicle updates the status of the target to "done", then broadcasts the identity and updated status of the target. Each vehicle within communication range updates its maps as follows: if the current value of the received target's status is "available" or "selected", it is upgraded to "done"; if it was "done", it remains "done".

6. return to step 2 until all targets are "done".

The effectiveness of this algorithm is surprisingly high, especially considering the low rate of communication (~15bits/target). Although inter-vehicle communication is parsimonious, it provides effective coordination. It provides local coordination, since two vehicles within range do not go after the same target. It also provides long-range coordination, since the locality of the target selection algorithm favors growth by domains (movies can be viewed at www.math.duke.edu/~daniel/search.html).

The algorithm is robust against loss of vehicles (Figure 1). The probability of mission failure is less than 1% up to an attrition rate of

.67/mission. The algorithm maintains a high level of performance even as vehicles are lost (time only increases by 75% at the highest attrition rate), since the targets a dead vehicle would have acquired are eventually selected by other nearby vehicles.

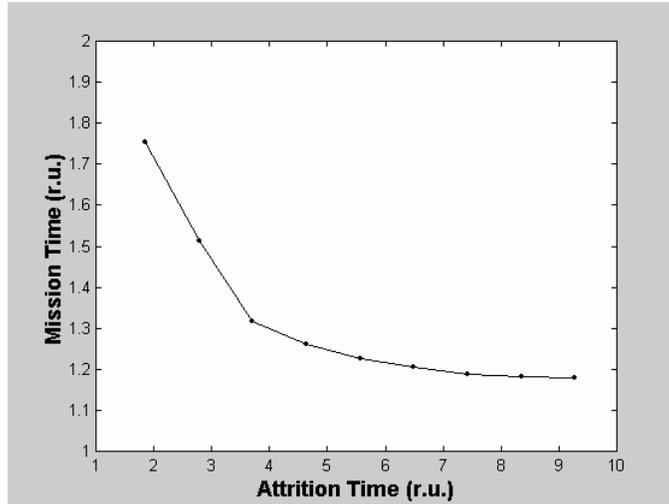


Figure 1. Dependence of the mission time on vehicle attrition. Mission time and attrition time are measured in units of the time taken without attrition. The mission ends when one of the vehicles has acquired all the targets. The simulation performed on a 100x100 target field, with 50 targets, 10 vehicles, unit speed, and infinite communication range.

The performance of the algorithm is relatively sensitive to the communication range (Figure 2). Full coordination is achieved when the communication range is comparable to the size of the target field. Mission time increases exponentially down to 10% of the field size, at which point the vehicles are effectively ex-communicado. We find that the rate of communication loss depends strongly on inter-target spacing, but that it is only weakly dependent on inter-vehicle spacing. The sensitivity to the communication range can be reduced substantially by a minor alteration discussed in Section 3.2.

Mission time increases as the square root of the number of targets (Figure 3), and inversely with the number of vehicles (Figure 4). This dependence is expected: Without attrition and with perfect communication, the target field is scavenged uniformly. Because individual vehicles select the target nearest to them, the transit time between targets goes as $1/\sqrt{T}$. Since each of the N vehicles acquires on average T/N of the T targets, mission time scales as \sqrt{T}/N .

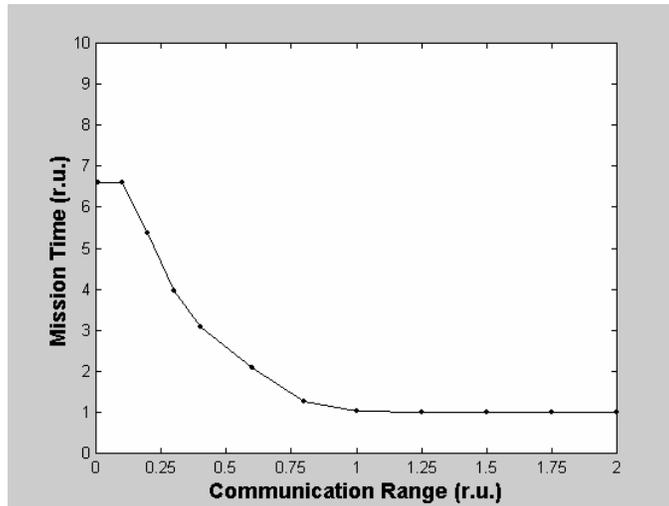


Figure 2. Dependence of mission time on communication range. Mission time is measured in units of the time taken when the communication range is infinite. Communication range is expressed in units of field size L . Simulation done on a 100×100 target field, with 50 targets, 10 vehicles, unit speed, and zero attrition.

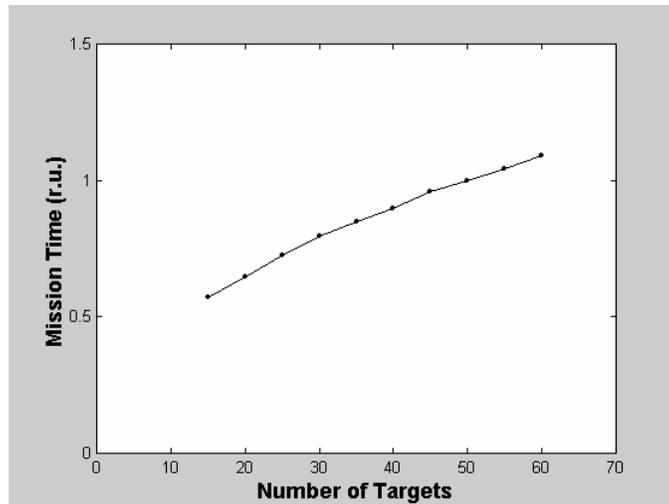


Figure 3. Dependence of mission time on number of targets. Regression onto a power law dependence gives an exponent of ~ 0.5 . Mission time is measured in units of the time taken when there are 50 targets. Simulation done on a 100×100 target field, with 10 vehicles, unit speed, zero attrition, and infinite communication range.

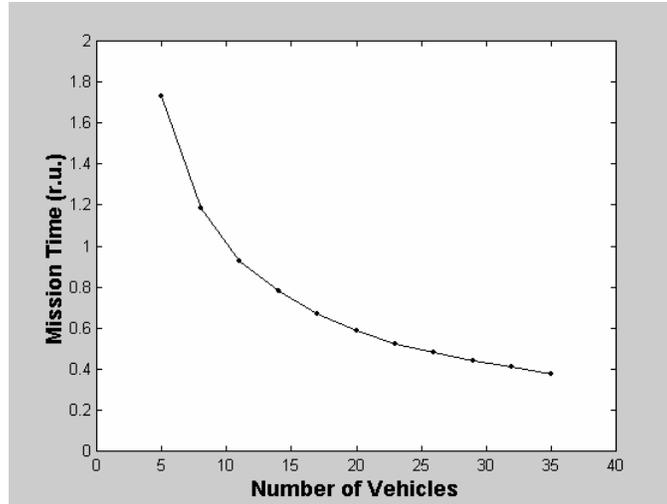


Figure 4. Dependence of mission time on number of vehicles. Regression onto a power law dependence gives an exponent of -0.8, indicating less than optimal vehicle coordination. Mission time is expressed in units of the time taken when there are 10 vehicles. Simulation done on a 100x100 target field, with 50 targets, unit speed, zero attrition, and infinite communication range.

3.2 Fast Greedy

Loss of long-range communication affects the performance of the greedy algorithm because distant vehicles are not coordinated. A solution to this problem is to increase the information content of each broadcast. Let us illustrate this on a simple example: say that vehicles A and C are within the range of B but not within range of each other. Instead of only transmitting the identity of the target under consideration (i.e. the greedy algorithm), A now transmits its entire target map to B. When B transmits, C receives A's data. This way, two vehicles that never meet can still coordinate their motion. This is the concept behind the fast greedy algorithm.

Fast greedy is implemented by altering Step 5 of the greedy algorithm:

5. when done with its target, each vehicle updates the status of the target to "done", then broadcasts the identity and updated status of *all targets in its map*. Each vehicle within communication range update its maps as follows: "available" targets take the received value; "selected" targets upgrade to "done" if the received value is "done"; "done" targets remain "done".

Fast greedy is much more robust to loss of communication than the greedy algorithm, achieving nearly complete coordination down to 20% of the

field size (Figure 5). As expected, the gain in robustness is offset by a substantial bandwidth penalty (200bits/target compared to 15).

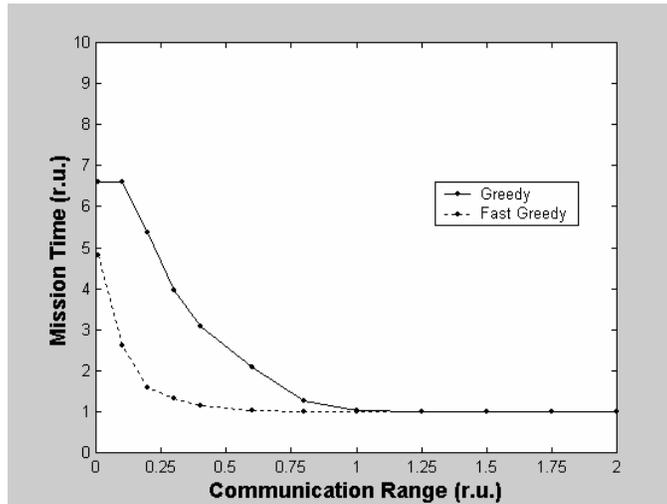


Figure 5. Comparison of mission time for fast greedy and greedy algorithms. Mission time is expressed in units of time taken by the greedy algorithm with infinite communication range. Communication range is expressed in units of field size L . Simulation done on a 100×100 target field, with 50 targets, 10 vehicles, unit speed, and zero attrition.

3.3 Divide and Conquer

Divide and Conquer separates the target field into sectors. Communication only happens when a vehicle selects a sector, otherwise each sector is searched quietly:

1. the target field is split into P sectors and a nexus. All targets belong to a sector. The nexus consists of a single target. The nexus belongs to all other sectors.
2. each vehicle keeps a map of all sectors; this map includes the identity of all target within the sector, their position and a sector status flag. The status flag takes one of three values: "available", "selected", and "done". All sectors are initially "available".
3. one of the vehicles is assigned to the nexus.
4. each vehicle, except that which occupies the nexus, looks at its map, and selects the nearest "available" sector. If no sectors are "available", the vehicle selects the nearest "selected" sector. If none are left, the vehicle ends its mission.

5. each vehicle broadcasts its choice to the nexus. The nexus updates its map as follows: if the current value of the received sector's status is "available", it is upgraded to "selected"; if it was "selected", it remains "selected"; if it is "done", it remains "done".

6. each vehicle travels to the "selected" sector, visits its targets, then returns to the nexus.

7. when done, each vehicle updates the status of the sector to "done", then broadcasts the identity and updated status to the nexus. The nexus receives the data and updates its map as follows: if the current value of the received sector's status is "available" or "selected", it is upgraded to "done"; if it was "done", it remains "done". It then broadcasts the status of all sectors to the vehicle, who then updates its list similarly. If the vehicle does not receive a reply, it assumes that the nexus vehicle is dead and takes its place.

8. return to step 4.

The nexus vehicle is not required if all vehicles are within communication range. Alternatively, an acoustic buoy could play the role of the nexus.

Simulations show that Divide and Conquer has the lowest communication cost of the three algorithms (1 bit/target vs 15 and 200). Mission time is also independent on communication range, since only short-range communication with the nexus is needed. Finally, the performance is robust against vehicle loss, since sectors left unfinished because of in-transit vehicle death are eventually selected by other vehicles.

4. CONCLUSION

We have presented three multi-agent algorithms that provide on-site coordination despite loss of vehicles and communication. These algorithms are example of the "fractional bandwidth" approach that we propose as a vehicle coordination strategy for harsh environments. We find that this strategy increases overall system effectiveness with regards to vehicle attrition and loss of communication.

Fractional bandwidth is valuable when centralized control is not available on-site and when the attrition rate is high, as is the case with the VSW-MCM problem. Our simulations indicate that by sharing their targets and using a minimum of communications, vehicles can substantially augment the effectiveness of their mission.

ACKNOWLEDGMENTS

We thank Mark Huber for his insights into the N-traveling salesmen problem. BC, DM, CT, and AB acknowledge the support of the Army Research Office (grant DAAD19-02-1-0055), the Office of Naval Research

(grant N000140310073), and the National Science Foundation (grants DMS-9983320 and DMS-0074049). MK acknowledges the support of DARPA MTO (contract DAAN02-98-C-4030).

REFERENCES

1. B. Fletcher. UUV Masterplan: A Vision for Navy UUV Development. *Proceedings Oceans 00 Conference*. 2000.
2. M.R. Blackburn, R.T. Laird, and H.R. Everett. Unmanned Ground Vehicles: Lessons Learned. *SPAWAR Technical Report 1869*. 2001.
3. NAVSEA Indian Head. Request for Proposal: Man Transportable Robotic System, *NAVSEA Indian Head Contract Office*. 2002.
4. See for example: R.D. Rhinehart, VSW/SZ MCM Reconnaissance Concept Assessment and Development, *ONR Ocean Atmosphere Space Annual Reports*. 2001; A.H. Rupert, Tactile Technologies for VSW/SZ MCM Missions. *ONR Ocean Atmosphere Space Annual Reports*. 1999; D. Darling and J. Jalbert. Morpheus AUV Development. *ONR Ocean Atmosphere Space Annual Reports*. 2001; D.R. DelBalzo et al. Environmental Effects on MCM Tactics Planning. *Proceedings Oceans 02 Conference*. 2002; T.C. Liu, H. Schmidt. AUV-Based Seabed Target Detection and Tracking. *Proceedings Oceans 02 Conference*. 2002.
5. A.J. Healey, Y. Kim. Control And Random Searching With Multiple Robots, *Proceedings IEEE CDC 00 Conference*. 2000; A. Healey. Application of Formation Control for Multiple Vehicle Robotic Minesweeping. *Proceedings IEEE CDC 01 Conference*. 2001.