Chemical Plume Source Localization with multiple Autonomous Underwater Vehicles

Bernardo Maciel, João Borges de Sousa and Fernando Lobo Pereira

Underwater Systems and Technologies Laboratory
Electrical and Computer Engineering Department
Faculty of Engineering of the University of Porto
Rua Dr. Roberto Frias, s/n 4200-465, Porto, Portugal

{bernardo.maciel, jtasso, flp}@fe.up.pt

***Abstract*—We developed a novel, multiple vehicle extension to a typical, single vehicle chemical plume source localization algorithm. The approach is implemented and tested in MATLAB®. We discuss simulations results to assess the potential of using the multiple vehicle approach. It occurs that there can be an advantage from using multiple vehicles.**

I. Introduction

In recent years, there has been some effort by the research community on finding efficient algorithms that solve the problem of Chemical Plume Tracing (CPT) [1]. CPT is a problem that can described as how to follow a plume, characterized by the concentration of a chemical substance, down to its source. There are many applications to the solutions to this problem, e.g., the prospecting for (and mapping of) hydrothermal vents, mine detection and effluent discharge monitoring.

In water, Autonomous Underwater Vehicles (AUVs) can be used to tackle this problem [2]. There exist reports of successful attempts [3]–[6] but using only one AUV. To the best of our knowledge, there have not been any attempts at using multiple AUVs to perform CPT tasks *underwater*. The main hindering factors pointed out are the poor capabilities of underwater communication [1, Sec. 6].

Here, we detail the work developed towards an expansion of an underwater CPT-related algorithm. One of the tasks that are performed in a CPT action is the localization of the source of the plume. We consider such task as outlined in [5] and present an expansion that makes use of multiple vehicles.

It is assumed that the reader is familiar with the work presented in [5]. Moreover, the notation used follows largely the one of [5].

II. Problem statement

We begin by stating the problem we set out to tackle. It is stated as

how to develop an algorithm for chemical plume source localization using multiple AUVs that performs better, according to metrics of accuracy and dispersion, than one that uses a single AUV.

Now, for this problem, an experimental scenario is considered.

A. Envisioned experimental setup

The setup considered for real experiments consists of having a team of AUVs operating at constant depth. These use acoustic modems to communicate with an Autonomous Surface Vehicle (ASV) which is stationary, working at the surface and located at the geometric center of the operating area (the area of interest, typically rectangular). Figure 1 depicts the envisioned experimental setup.

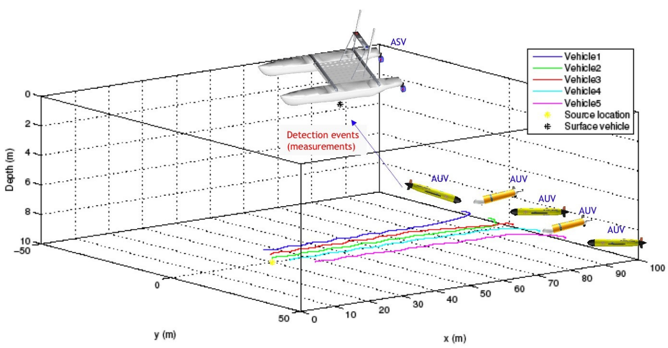

Fig. 1. Example of the envisioned experimental setup with 5 vehicles. AUVs are also depicted by their trajectories.

Moreover, AUVs are equipped with chemical concentration and flow measurement sensors. There exists also a localization system that allows an AUV to be aware of its own position, with some accuracy (e.g., a Long BaseLine (LBL) system).

B. Assumptions

With such setup in mind, we outline some assumptions before explaining the approach developed for chemical plume source localization.

From previous work: The assumptions taken in [5] are also used here and can be mentioned throughout the rest of the document. For example, it is assumed that the operating area is divided into rectangular cells [5, Sec. III] and that flow has to be approximated [5, Sec. IV].

978-1-4244-5222-4/10/$26.00 ©2010 IEEE

Formation control: We decouple the problem of mapping the source of the chemical plume from the vehicle motion planning part as in [5]. It is assumed that formation control of a team of AUVs is in place, i.e., we have an algorithm for it that is "independent" from the one described here for mapping.

Synchronization of time at mission start: Initial time values are synchronized at mission start and have negligible drift.

Communication: The communications between the AUVs and the ASV are considered to have negligible delays - i.e., lower than the time step of the algorithm - and minor packet losses. While the latter is simply assumed for simplification purposes, the former is explained as follows.

The communication delay here considered is the time difference between the observation of an event by an AUV until the acknowledge of such happening by the ASV. It is approximately [1] given by

$$\begin{aligned} \text{delay} &= \text{message assembly time} + \text{transmission time} + \\ &\quad + \text{message disassembly time} \\ &\approx \text{transmission time} \\ &= \text{propagation time} + \text{message sending time} \\ &= \frac{\text{AUV-ASV distance}}{c_{\text{sound in water}}} + \frac{\text{message size}}{\text{bit rate}} \end{aligned}$$

Each of the unknowns can be computed as described next.

- *AUV-ASV distance:* Considering a mission setup such as the one depicted in figure 1, this distance is higher when an AUV is at any corner of the square-shaped operating area. Assuming a maximum operating depth of $30m$ and an operating area with $100m$ sides, the AUV-ASV distance is upper bounded by

$$\sqrt{30^2 + \sqrt{(100/2)^2 + (100/2)^2}^2} \approx 77m$$

- $c_{\textit{sound in water}}$: The speed of sound in water is assumed to be $1500ms^{-1}$.
- *Message size:* Each message assembled by an AUV contains information about an event. Such information is
  - which event happened (detection/non-detection) - 1 bit -,
  - a flow measurement - 16 bit -,
  - vehicle location (NED coordinates) - $3*64$ bit - and
  - a timestamp - 32 bit.

  Message size is then equal to overheads $+ 1 + 3*64 + 32 + 16 \approx 256\text{bit} = 32\text{byte}$.
- *Bit rate:* This value is highly variable, but we assume it to be low so it can be representative of existing, available and relatively cheap acoustic modems. Thus we take the bit rate to be lower bounded by $512bps$.

[1]Although the setup implies near vertical propagation through water, its effects (if any exists) are neglected for simplicity.

Delay is therefore lower than or equal to $\frac{77}{1500} + \frac{256}{512} \approx 551ms$. Considering that our algorithm runs at $1Hz$, delay becomes negligible as its upper bound is lower than $1s$.

III. Source Probability Mapping with Multiple Vehicles

To obtain the location of the source of a plume, we follow the ideas in [5] and expand them to obtain a source probability map. This map makes the correspondence from all cells that divide the operating area to the probability that the plume source is located in the cells.

A. Approach

The simplest approach consists of considering that the ASV performs all map computations, i.e., the asynchronous source probability map updates are carried out at the ASV. This means that the source mapping algorithm is centralized. On the other hand, chemical concentration and flow measurements are distributed amongst the AUVs.

AUVs communicate events of detection or non-detection (of a chemical substance) to the ASV. Accordingly, there will be N events that have to be incorporated in the source probability map update process. There are N_D detections and $N_{\bar{D}}$ non-detections ($N = N_D + N_{\bar{D}}$), i.e., $D_{j_p}, p \in \{1, \ldots, N_D\}$ and $\bar{D}_{j_q}, q \in \{1, \ldots, N_{\bar{D}}\}$, where the j_p, j_q are the vehicles' locations.

It is assumed that all D_{j_p}, $\bar{D}_{j_q}$ and the sequence of previous events, $B(t_{k-1})$, are independent and conditionally independent given A_i [2]. We follow [5, Sec. V] to obtain, by Bayesian theory,

$$\begin{aligned} &\alpha_i(t_k) = \\ &= \Pr\left(A_i | B(t_{k-1}), D_{j_1}(t_k), \ldots, D_{j_{N_D}}(t_k), \bar{D}_{j_1}(t_k), \ldots, \bar{D}_{j_{N_{\bar{D}}}}(t_k)\right) \\ &= \frac{\Pr\left(B(t_{k-1}), D_{j_1}(t_k), \ldots, D_{j_{N_D}}(t_k), \bar{D}_{j_1}(t_k), \ldots, \bar{D}_{j_{N_{\bar{D}}}}(t_k) | A_i\right) \Pr(A_i)}{\Pr\left(B(t_{k-1}), D_{j_1}(t_k), \ldots, D_{j_{N_D}}(t_k), \bar{D}_{j_1}(t_k), \ldots, \bar{D}_{j_{N_{\bar{D}}}}(t_k)\right)} \\ &\qquad \text{(by Bayes's theorem)} \\ &= \frac{\Pr\left(B(t_{k-1}), D_{j_1}(t_k), \ldots, D_{j_{N_D}}(t_k), \bar{D}_{j_1}(t_k), \ldots, \bar{D}_{j_{N_{\bar{D}}}}(t_k) | A_i\right) \Pr(A_i)}{\Pr(B(t_{k-1})) \prod_{p=1}^{N_D} \Pr\left(D_{j_p}(t_k)\right) \prod_{q=1}^{N_{\bar{D}}} \Pr\left(\bar{D}_{j_q}(t_k)\right)} \\ &\qquad \text{(by independence of events)} \\ &= \frac{\Pr(B(t_{k-1})|A_i) \Pr(A_i)}{\Pr(B(t_{k-1}))} \frac{\prod_{p=1}^{N_D} \Pr\left(D_{j_p}(t_k)|A_i\right)}{\prod_{p=1}^{N_D} \Pr\left(D_{j_p}(t_k)\right)} \frac{\prod_{q=1}^{N_{\bar{D}}} \Pr\left(\bar{D}_{j_q}(t_k)|A_i\right)}{\prod_{q=1}^{N_{\bar{D}}} \bar{D}_{j_q} \Pr((t_k))} \\ &\qquad \text{(by conditional independence of events given } A_i\text{)} \\ &= \left(\frac{1}{\Pr(A_i)}\right)^{N_D} \alpha_i(t_{k-1}) \prod_{p=1}^{N_D} \frac{\Pr\left(D_{j_p}(t_k)|A_i\right) \Pr(A_i)}{\Pr\left(D_{j_p}(t_k)\right)} \times \\ &\qquad \times \prod_{q=1}^{N_{\bar{D}}} \frac{\Pr\left(\bar{D}_{j_q}(t_k)|A_i\right)}{\Pr\left(\bar{D}_{j_q}(t_k)\right)} \end{aligned}$$

[2]A_i is the event that there is a source in cell C_i.

$$=M^{N_D}\alpha_i(t_{k-1})\prod_{p=1}^{N_D}\Pr\left(A_i|D_{j_p}(t_k)\right)\prod_{q=1}^{N_{\bar{D}}}\frac{\Pr\left(\bar{D}_{j_q}(t_k)|A_i\right)}{\Pr\left(\bar{D}_{j_q}(t_k)\right)}$$

and, therefore, the following recursive, source probability map update rule,

$$\alpha_i(t_k)=M^N\alpha_i(t_{k-1})\prod_{p=1}^{N_D}\beta_{ij_p}(t_0,t_k)\prod_{q=1}^{N_{\bar{D}}}\frac{\gamma_{ij_q}(t_0,t_k)}{\sum_{i=1}^{M}\gamma_{ij_q}(t_0,t_k)} \quad (1)$$

where $\beta_{ij_p}(t_0,t_k)$ is the probability that there is a source in cell C_i, given that there is detectable chemical in cell C_{j_p} at time t_k, and $\gamma_{ij_q}(t_0,t_k)$ is probability of not detecting the chemical in cell C_{j_q} at time t_k due to the continuous release from a source in cell C_i. This map update rule was implemented in MATLAB® and validated through simulations.

Independence of simultaneous events: The assumption that the simultaneous events (that are incorporated in the map update rule) are independent is an important one. This is because such independence is not verified in practice [6, Section 2.1] [7]. However, in this work we did not try to circumvent this problem. We assume the referred independence for simplicity and leave it to be tackled in future work (see section VII).

Before discussing the results of the simulations ran, a description of the simulation setup is in order.

IV. Simulation setup

The simulation setup developed to test the above described algorithm consisted of 2 parts: simulating actual missions, where AUVs execute a plume tracing algorithm, and running our algorithm on the data gathered during the first part.

A. Gathering data

For the first part, we used an previously written application [8], entitled "Real-time simulation of odor release", that simulates not only a chemical plume but also its tracing by AUV(s). This application permits the gathering of data to use in our algorithm: vehicles' location, chemical concentration and flow speed. A picture showing the visual output of the application is shown in figure 2.

The technique for plume simulation used in the application is described in detail in [8]. Many simulation parameters can be tweaked, from water flow speed to plume growth rate. In our case, flow is kept approximately constant and we changed plume parameters so that there would be several degrees of plume persistence, meaning that all simulated vehicles could track the source with more or less difficulty.

The vehicles simulated in the application are programmed to use a CPT strategy that is inspired in the behaviour of moths. This biomimetic approach is one option for solving the CPT problem [1, Sec. 4]. Moreover, we consider every vehicle simulated to be independently running its strategy from the other ones, building on the "Formation control" assumption mentioned in Section II-B.

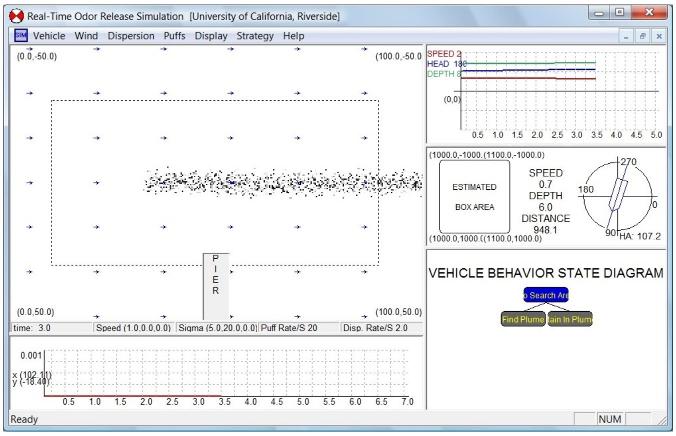

Fig. 2. "Real-time simulation of odor release" application. ©University of California, Riverside.

In this way, collision avoidance is not enforced at simulation time. However, the data gathered is analyzed afterwards to check if a collision occurred, in which case the corresponding simulation is discarded.

The application was run various times, in 3 batches. Every batch of simulations consisted of runs with a number of vehicles ranging from 1 to 5 (batches 1 and 3) or 3 to 5 (batch 2). The first and second batches were controlled using a visual stopping condition, i.e., when the vehicles "seem close to the source". From batch 1 to batch 2, the AUVs' strategy and a plume parameter were changed in order to allow more diverse data to be gathered. Batch 3 consisted of longer runs as the stopping condition was time-based. Moreover, the initial location of the vehicles for this batch formed a tighter group than for the other ones.

B. Algorithm testing

The data gathered was used to test the map update rule, detailed in Section III-A, in MATLAB®. The code written can be split in 3 parts: loading the data gathered, running the algorithm on the loaded data and analyzing the results.

Loading data: First of all, the data gathered was loaded as if it had been received at the ASV. Then, we plotted the data in different ways, in order to appreciate it visually. Moreover, we determine if any collision occurred by looking at the minimum distance between vehicles. If this value is lower than the length of the AUVs ($1.2m$), the simulation under analysis is discarded.

Running the algorithm: The data loaded in the previous step was, for every time instant, used to update the source probability map according to the map update rule described in section III-A.

The source mapping algorithm was also ran as if the vehicles had ran single. This means that, at the previous simulation part, only data from one vehicle at a time was loaded and then used in the algorithm. Such runs serve as bases for comparison with the multiple vehicles approach.

The maps obtained over time were stored for further analysis, both numerical and visual.

Analyzing results: The analysis of results comprises both a visual and a numerical part. Visually, comparisons are made between the final maps (maps at the simulation's last time instant) resulting from the multiple vehicles and single vehicle runs. Numerically, two metrics are computed and the results are also compared.

The two metrics computed are

- *Metric 1:* the distance from the map's maximum likelihood location (i.e., the map cell with highest source presence probability) to the actual source location and
- *Metric 2:* the area of the significant portion of the map, i.e., the number of cells with probability larger than a threshold ($1 * 10^{-6}$ in this case).

Metric 1 can be used to show how far is the algorithm's output from reality, thus being an accuracy estimate. On the other hand, metric 2 quantifies the dispersion of the source probability map values [3] . If the significant area is small, then the map has low uncertainty on the cell(s) that contains the plume source (being it accurate or not). Furthermore, analyzing both metrics 1 and 2 over time sheds light on the speed of convergence to the final result and its uncertainty.

V. Results

This section presents the results obtained and comments on them. For every simulation that was not discarded due to the occurrence of collisions [4] , we show

- plots of the data gathered,
- plots of the final maps obtained and
- graphs of the metrics computed.

We also include the runs with 1 simulated AUV, that show that the algorithm also works outside a multiple vehicles scenario.

A. All simulations

We first show, in table I, the values obtained for the metrics for all simulations that were not discarded.

We now discuss the results obtained for each simulation.

B. Batch 1

Simulation with 1 AUV: Figure 3 shows the data gathered for the 1 AUV simulation of batch 1.

The run with 1 AUV turned out to be fairly inaccurate, which is visible from figure 4.

Metric 1 shows poor accuracy in source localization, although source probability values are not too scattered. Metric 2 is, indeed, just below 2%, which is a considerably small portion of the map. However, convergence rates in the two

[3]The dispersion metric used in [5] is the 95% probability area. Here, it was found that it was easier to code the computation of the significant area than the computation of the 95% probability area. Therefore, the area of the significant portion of the map was chosen as a metric. Comparisons between this work and [5] are, nevertheless, still possible, if desired. In fact, both the different metrics reveal the spread of map values.

[4]All multiple vehicles simulations in batch 3 and one in batch 2 (with 4 AUVs) were discarded due to the occurrence of collisions.

TABLE I

Final values for metrics 1 ("Source dist.") and 2 ("Sig. area") for all simulations. Runs as if the AUVs had been single are labelled by vehicle number and their average by "Avg." ("Vehicle(s)" column).

Batch #	# vehicles	Vehicle(s)	Source dist. (m)	Sig. area
1	1	All	40	0.0184
1	2	All	13	0.1012
		1	40	0.0253
		2	13	0.0290
		Avg.	26	0.0272
1	3	All	7	0.0045
		1	40	0.0253
		2	17	0.0030
		3	13	0.0300
		Avg.	23	0.0194
1	4	All	30	0.0126
		1	40	0.0253
		2	22	0.0096
		3	5	0.0090
		4	12	0.0836
		Avg.	20	0.0319
1	5	All	30	0.0056
		1	40	0.0253
		2	7	0.0624
		3	18	0.0054
		4	20	0.0002
		5	14	0.0196
		Avg.	20	0.0226
2	3	All	11	0.0001
		1	24	0.0144
		2	6	0.0360
		3	37	0.0672
		Avg.	23	0.0392
2	5	All	11	0.0001
		1	18	0.0054
		2	26	0.0612
		3	4	0.0132
		4	1	0.0032
		5	37	0.0630
		Avg.	17	0.0292
3	1	All	5	0.0189

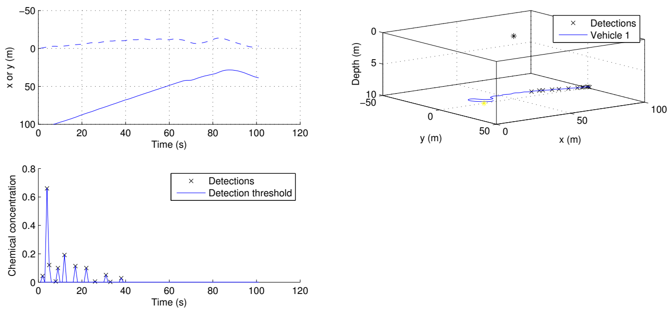

Fig. 3. Plots resulting from analysis of the data gathered with 1 (simulated) AUV. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left).

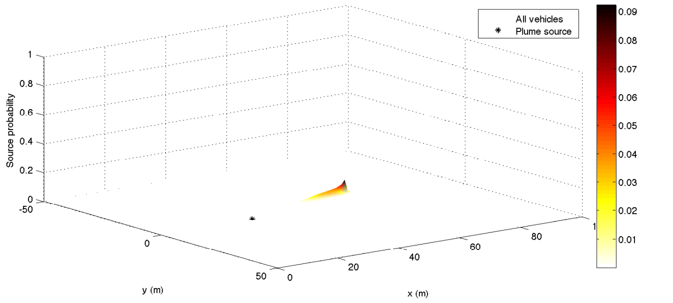

Fig. 4. Final source probability map for the 1 AUV simulation of batch 1.

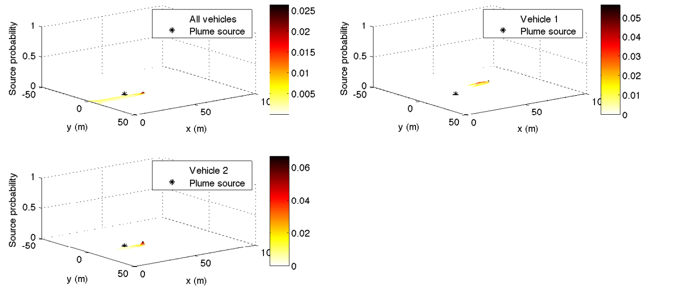

Fig. 7. Final source probability maps for the 2 AUV simulation of batch 1. The final map for the multiple vehicles run is shown (top left) as well as the final maps for the runs as if the vehicles had ran single (other plots).

metrics seem to differ (see figure 5), as metric 1 converges fairly slower than metric 2.

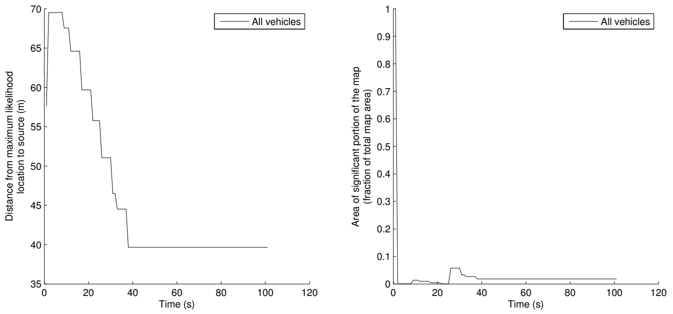

Fig. 5. Evolution of metrics 1 (left) and 2 (right) over time for the 1 AUV simulation of batch 1.

Simulation with 2 AUVs: Figure 6 shows the data gathered for the 2 AUVs simulation of batch 1.

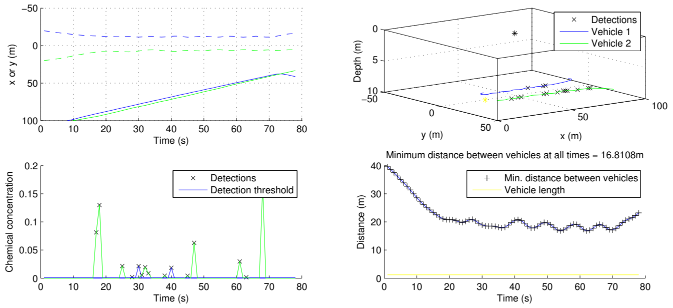

Fig. 6. Plots resulting from analysis of the data gathered with 2 (simulated) AUVs. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left); inter-vehicle distances (bottom right).

With 2 AUVs, the results seem better (from figure 7) when compared with the 1 AUV run. However, when comparing with the other single vehicle runs, that is not always the case.

Metric 1 shows that localization accuracy is fair with the multiple vehicles approach and better than the average of the other runs. It is, however, as accurate as the run of vehicle 2. On the other hand, metric 2 is shows a higher dispersion of values when compared to the other runs, average or not. It can be seen in figure 8 that at a time around $60s$, the map's significant area rises with the multiple vehicles approach whereas it remains low with the single vehicle approaches. This is perhaps caused by the lack of detections from approximately $50s$ to $60s$ (see figure 6), which the algorithm is not capable of making up for in terms of metric 2.

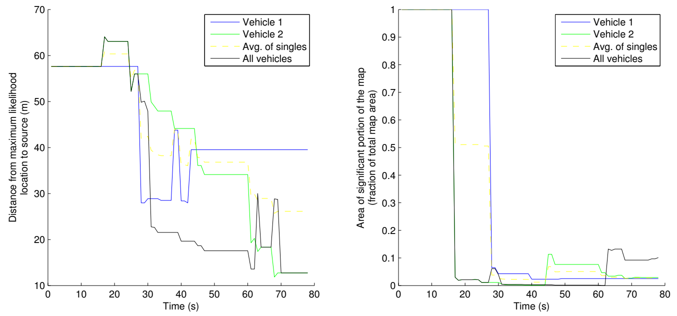

Fig. 8. Evolution of metrics 1 (left) and 2 (right) over time for the 2 AUVs simulation of batch 1. Runs as if the vehicles had ran single are shown, along with their average.

Therefore, in this case there is not a clear advantage, in source localization performance, from using 2 vehicles instead of 1. Nevertheless, note that (from figure 8) convergence appears to be faster, for both metrics, with the multiple vehicles approach.

Simulation with 3 AUVs: Figure 9 shows the data gathered for the 3 AUVs simulation of batch 1.

The simulation with 3 AUVs shows good results for the multiple vehicles run when comparing with the runs as if the vehicles had ran single. Figure 10 shows a final map with a probability peak that is close to the source and small in area. This contrasts with the other plots in figure 10, which show either diffused or inaccurate values.

The results for both metrics 1 and 2 reinforce that the multiple vehicles approach brings an advantage. From figure 11, we can see that localization with the 3 AUVs is not only accurate but also fast to converge to the final values. Table I confirms the former, showing a $7m$ inaccuracy from the

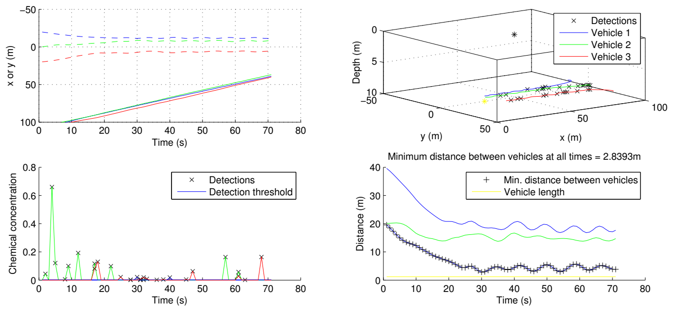

Fig. 9. Plots resulting from analysis of the data gathered with 3 (simulated) AUVs. This simulation is part of batch 1. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left); inter-vehicle distances (bottom right).

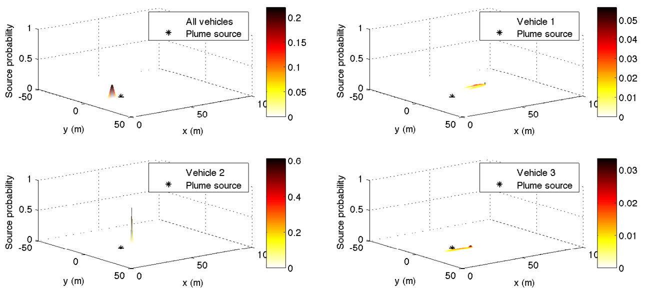

Fig. 10. Final source probability maps for the 3 AUVs simulation of batch 1. The final map for the multiple vehicles run is shown (top left) as well as the final maps for the runs as if the vehicles had ran single (other plots).

maximum likelihood location to the actual source location and a significant area below 1%. In general, both metrics are higher valued for the single vehicle runs; also, the average values are significantly larger than those of the multiple vehicles run.

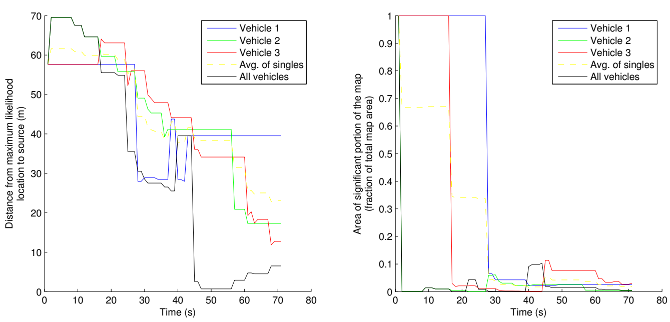

Fig. 11. Evolution of metrics 1 (left) and 2 (right) over time for the 3 AUVs simulation of batch 1. Runs as if the vehicles had ran single are shown, along with their average.

In this case, using 3 AUVs clearly brings an advantage to the mission. The fact is, localization performance is better from all points of view.

Simulation with 4 AUVs: Figure 12 shows the data gathered for the 4 AUVs simulation of batch 1.

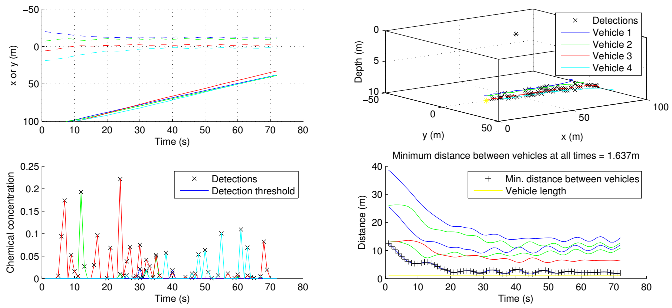

Fig. 12. Plots resulting from analysis of the data gathered with 4 (simulated) AUVs. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left); inter-vehicle distances (bottom right).

The 4 AUVs simulation of batch 1 resulted in a multiple vehicles performance worse than that of 3 AUVs simulation. From figure 13, the final source probability map shows a peak which is quite off the actual source location. On the other hand, at least half of the single vehicle runs end up in more accurate maps.

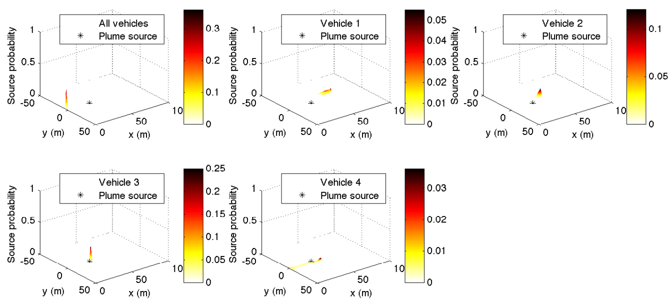

Fig. 13. Final source probability maps for the 4 AUVs simulation of batch 1. The final map for the multiple vehicles run is shown (top left) as well as the final maps for the runs as if the vehicles had ran single (other plots).

This is confirmed by the plots in figure 14 and the values in table I. In fact, the final values for metric 1 are worse with the multiple vehicles approach than with most of the single vehicle runs and their average. For metric 2, this is not the case anymore: although 2 single vehicle runs yield final values lower than 1%, the others, including the average, do not. Thus the multiple vehicles value, 1.26%, is acceptable.

Nevertheless, looking at figure 14, we can see that up to a time around 60s, the multiple vehicles run results in faster convergence and lower metrics values than the other runs. At that time, both metrics increase for the multiple vehicles run. While the algorithm ends up reducing the value of metric 2 before the end of the simulation, it does not improve localization accuracy. The cause for this error is not clear from figure 12.

In the 4 AUVs simulation, using multiple vehicles is not beneficial as in the 3 AUVs simulation. However, high convergence rates do suggest that slight improvements to the algorithm may lead to a superior localization performance.

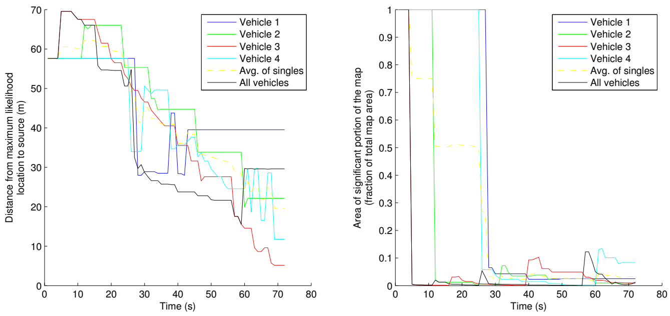

Fig. 14. Evolution of metrics 1 (left) and 2 (right) over time for the 4 AUVs simulation of batch 1. Runs as if the vehicles had ran single are shown, along with their average.

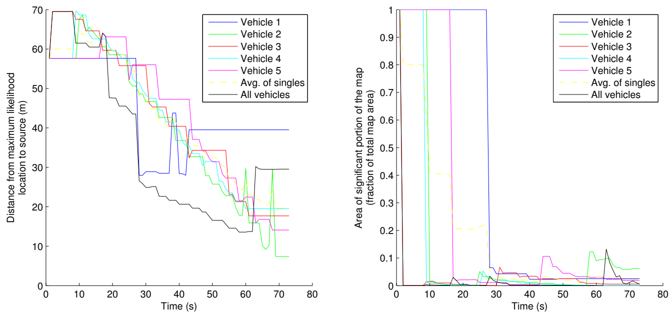

Fig. 17. Evolution of metrics 1 (left) and 2 (right) over time for the 5 AUVs simulation of batch 1. Runs as if the vehicles had ran single are shown, along with their average.

Simulation with 5 AUVs: Figure 15 shows the data gathered for the 5 AUVs simulation of batch 1.

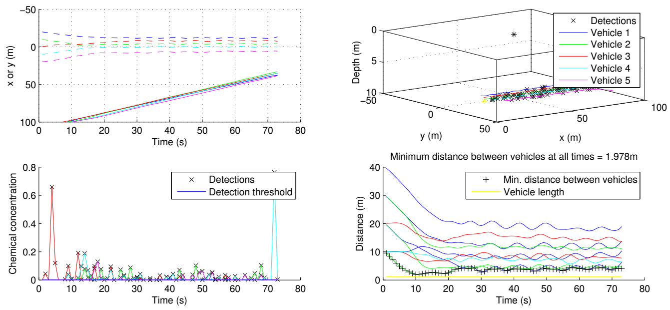

Fig. 15. Plots resulting from analysis of the data gathered with 5 (simulated) AUVs. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left); inter-vehicle distances (bottom right).

The 5 AUVs simulation lead to results very similar to the 4 AUVs simulation. Figures 16 and 17, along with table I, show that the multiple vehicles run yields higher convergence rates but worse final values for metric 1. This results in a final map which is not at all scattered but still with a maximum likelihood location distant from the actual source location.

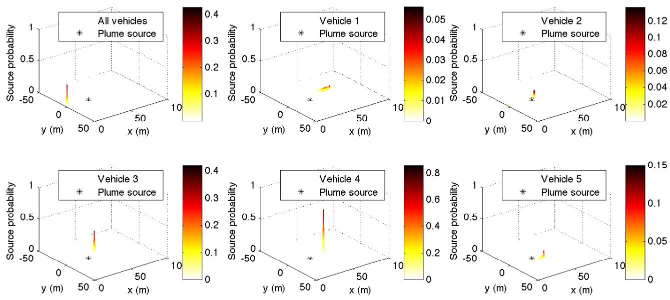

Fig. 16. Final source probability maps for the 5 AUVs simulation of batch 1. The final map for the multiple vehicles run is shown (top left) as well as the final maps for the runs as if the vehicles had ran single (other plots).

Furthermore, the analysis of the evolution of both metrics 1 and 2 done for the 4 AUVs simulation applies directly to the 5 AUVs simulation. So, there is not a clear advantage brought from using multiple vehicles.

C. Batch 2

Simulation with 3 AUVs: Figure 18 shows the data gathered for the 3 AUVs simulation of batch 2.

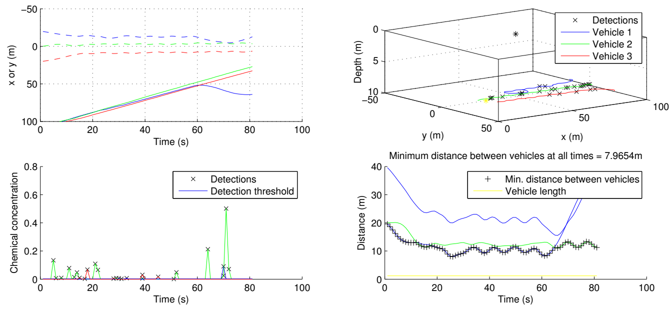

Fig. 18. Plots resulting from analysis of the data gathered with 3 (simulated) AUVs. This simulation is part of batch 2. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left); inter-vehicle distances (bottom right).

In batch 2, the simulation with 3 AUVs returned good results. From figure 19, we can see that the final map for the multiple vehicle run is a very narrow peak located not very distant from the actual source location. The single vehicle runs do not result in such a small area with high probability values although one seems more accurate (vehicle 2).

Figure 20 and table I confirm these observations. For the multiple vehicle approach, metric 2 has an extremely small final value (0.01%) while the single vehicle runs, and its average, resulted in higher significant areas. As for localization accuracy, metric 1 final values show that the multiple vehicle approach only performs worse than one single vehicle run (vehicle 2), as mentioned before. The final value for metric 1 of $11m$, for the multiple vehicle approach, is reasonable.

Furthermore, as for the general case of the already analyzed simulation results, both metrics 1 and 2 seem to converge faster to their final values with multiple vehicles than with a single vehicle.

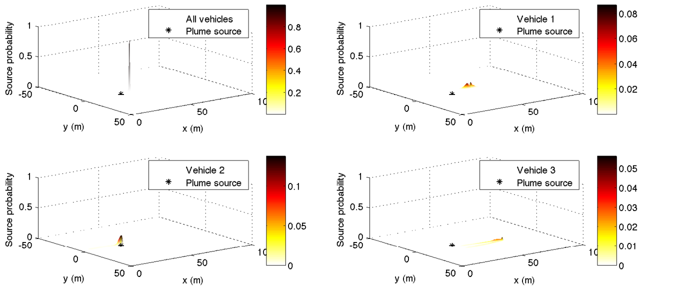

Fig. 19. Final source probability maps for the 3 AUVs simulation of batch 2. The final map for the multiple vehicles run is shown (top left) as well as the final maps for the runs as if the vehicles had ran single (other plots).

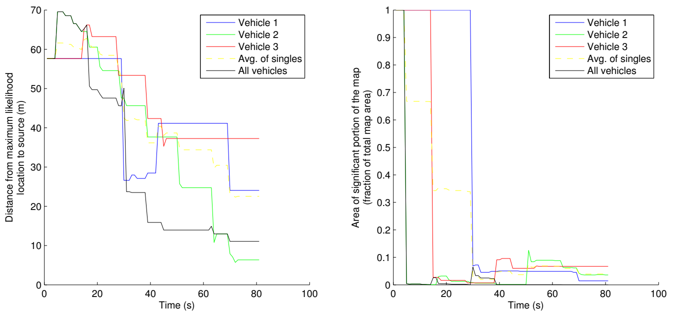

Fig. 20. Evolution of metrics 1 (left) and 2 (right) over time for the 3 AUVs simulation of batch 2. Runs as if the vehicles had ran single are shown, along with their average.

In this case, bringing multiple vehicles to the table can be considered advantageous. When comparing its performance with the average performance of the single vehicle runs, the gain is significant. Still, the comparison with each single vehicle run shows that localization accuracy should be improved in order to firmly prove the superiority of the multiple vehicle approach.

Simulation with 5 AUVs: Figure 21 shows the data gathered for the 5 AUVs simulation of batch 2.

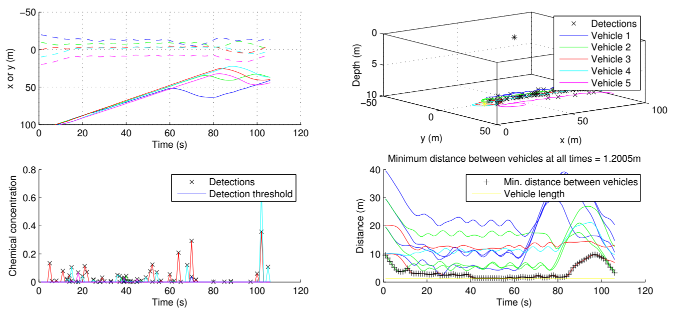

Fig. 21. Plots resulting from analysis of the data gathered with 3 (simulated) AUVs. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left); inter-vehicle distances (bottom right).

Figure 22 shows the final source probability maps for the 5 AUVs simulation of batch 2. We can see that the maps are, in general, similar to what happened with the 3 AUVs simulation (see figure 19). Not only does the final map for the multiple vehicle approach consist of a narrow peak located not very distant from the actual source location but also the single vehicle runs vary, with some appearing more accurate and others more scattered.

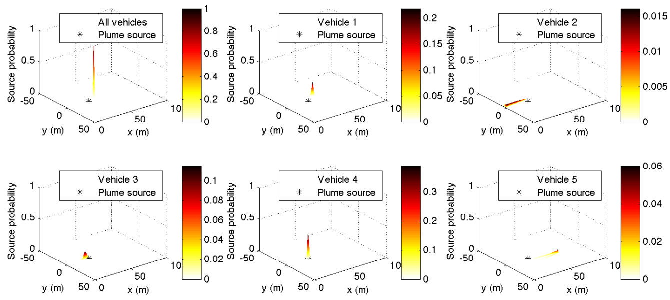

Fig. 22. Final source probability maps for the 5 AUVs simulation of batch 1. The final map for the multiple vehicles run is shown (top left) as well as the final maps for the runs as if the vehicles had ran single (other plots).

Metrics 1 and 2 final values, for the 5 AUVs run, are also the same than for the 3 AUVs simulation (see table I). However, there are now two single vehicle runs that are more precise in localization (vehicles 3 and 4). Furthermore, the multiple vehicle approach yielded final results that are not so distant from the average of the single vehicle runs. The same applies to speed of convergence, which is visible from figure 23.

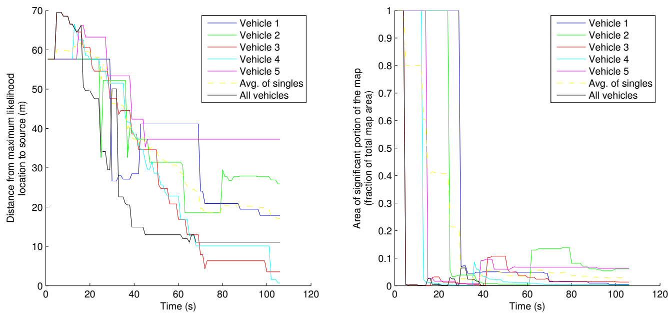

Fig. 23. Evolution of metrics 1 (left) and 2 (right) over time for the 5 AUVs simulation of batch 2. Runs as if the vehicles had ran single are shown, along with their average.

So, as in the 3 AUVs simulation of the same batch, multiple vehicles improve localization performance when comparing to the average of single vehicle performances. However, this enhancement is not pronounced as before, as there are better results within the single vehicle runs.

D. Batch 3

Simulation with 1 AUV: Figure 24 shows the data gathered for the 1 AUV simulation of batch 3.

This run with 1 AUV resulted in a final map that seems more accurate, but equally scattered, than that of the 1 AUV simulation batch 1. This can be seen from figures 4 and 25.

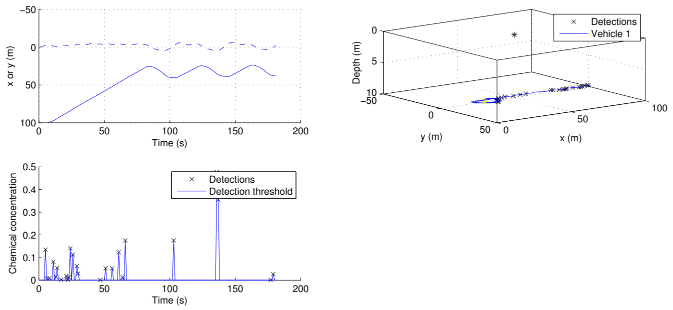

Fig. 24. Plots resulting from analysis of the data gathered with 1 (simulated) AUV. Plots are x and y trajectories (top left); trajectories with chemical detections in 3D (top right); chemical measurements and detections (bottom left).

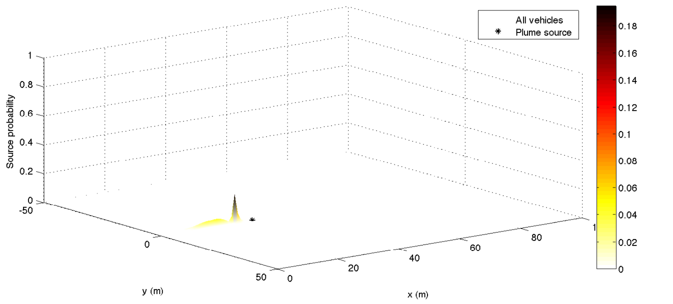

Fig. 25. Final source probability map for the 1 AUV simulation of batch 3.

In fact, metric 1 shows good accuracy in source localization with a value of $5m$. On the other hand, metric 2 is very close to that of the 1 AUV simulation of batch 1. This can be seen from both figure 26 and table I. Moreover, like for the 1 AUV simulation of batch 1, the convergence rate of metric 1 is fairly higher than that of metric 2.

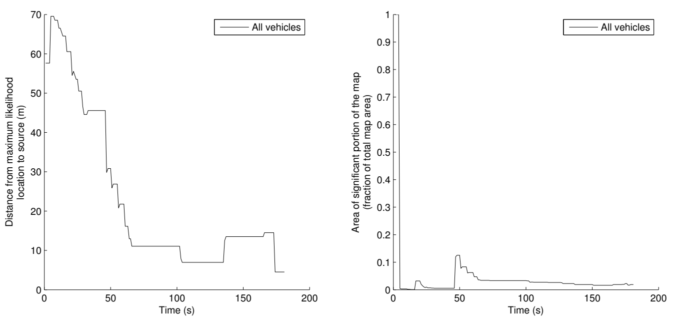

Fig. 26. Evolution of metrics 1 (left) and 2 (right) over time for the 1 AUV simulation of batch 1.

Having discussed the results from the simulations ran, we proceed to draw concluding remarks.

VI. Conclusion

Not all multiple vehicles simulations show a clear advantage on the single vehicle counterparts. However, in general, using multiple vehicles improves, to some extent, either localization accuracy, or certainty, or both. In fact, AUVs have consistent behaviours across the simulations, as they get closer to each other and to the source in spite of the plume's volatility. In summary, *the multiple vehicle approach developed can bring a performance boost, when compared with the single vehicle approach*.

Still, not all the multiple vehicles simulations have shown a clear advantage. The algorithm requires changes in order to definitely show a leading edge from its single vehicle version.

In addition, we consider that the computational effort will not be a major concern in a real experiment. Although an analysis of the processing efforts was not carried out, the fact that the algorithm is computationally centralized means that provisions can be made to have adequate computational power in the ASV.

VII. Future work

This section describes directions for future work as well as problems and/or details that are yet to be addressed. Many of these remarks should be considered before an implementation on real hardware and corresponding experiments are performed.

Relaxing assumptions: The most important directions for the future lie in dealing with the assumptions taken here (see section II-B). Only this way can the algorithm become more robust, in theory and in practice.

The assumption of having independent simultaneous events is a strong one, and it should be tackled with priority. Both in [6, Section 2.1] and [7], forward sensor models are used to deal with this problem. A suggestion for future research is that such approach can perhaps be adapted to a scattered sensor model. This would fit ideally in the scenario outlined in section II-A and bring robustness and performance improvements to the algorithm here presented.

Another important assumption is that of decoupling the mapping and planning problems. Motion planning should definitely be integrated with the source mapping algorithm. To this end, consider simple dynamical models (e.g., unicycle models) for the vehicles and take the source probability map at a given time instant. Such models could be used to predict the next best move for the vehicles (each one or as a group) so that, one time step after, the source probability map improves in localization accuracy and/or certainty. For these developments, perhaps the reach set framework could be of use (e.g., via the Ellipsoidal Toolbox [9]).

Furthermore, the assumption of negligible delays in communication might become less realistic if the experimental setup changes in, e.g., size, or if we require a smaller time step for our algorithm. When delay stops being negligible, there is need to deal with the arrival of delayed measurements.

Still related to communication, we might have to deal with the fact that packet losses are not minor. For this problem, we envision a solution that is based on simulating the losses of packets at the ASV. If we assume a distribution for the Packet Reception Ratio (PRR) at the ASV, we can make up for lost packets by either using heuristics together with the source mapping algorithm or letting the AUVs know that some messages should be resent.

Improving simulations: Future work might also concern improving the quality and breadth of simulations. Options include the integration of the source mapping algorithm with network simulators, e.g., NS-2 [10], via underwater channel models [11]. This would introduce a better, more detailed channel model and hopefully lead to an adaptation of the source mapping algorithm that would improve robustness. Also, other mission scenarios than that described in section II-A should be considered.

Other expansions: Further developments may also consist of trying to turn the algorithm here explained into a distributed one. Moreover, we could follow the ideas in [6] to obtain a multiple source mapping algorithm, thus more complete. Finally, another idea is to expand the algorithm described in [4] to multiple vehicles and compare with the results obtained with our algorithm.

Acknowledgment

The authors would like to thank Prof. Shuo Pang (Embry-Riddle Aeronautical University at Daytona Beach) and Prof. Jay A. Farrell (University of California, Riverside) for their assistance in the development of this work.

References

[1] W. Naeem, R. Sutton, and J. Chudley, “Chemical plume tracing and odour source localisation by autonomous vehicles,” *The Journal of Navigation*, vol. 60, pp. 173–190, 2007.
[2] B. Maciel and J. de Sousa, “Sensor-based problems and techniques for autonomous underwater vehicles,” in *OCEANS 2009 - EUROPE*, May 2009, pp. 1–6.
[3] J. A. Farrell, S. Pang, W. Li, and R. Arrieta, “Chemical plume tracing experimental results with a REMUS AUV,” in *Proc. OCEANS 2003*, vol. 2, Sept. 2003, pp. 962–968 Vol.2.
[4] J. A. Farrell, S. Pang, and W. Li, “Plume mapping via hidden Markov methods,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, vol. 33, no. 6, pp. 850–863, Dec. 2003.
[5] S. Pang and J. Farrell, “Chemical Plume Source Localization,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, vol. 36, no. 5, pp. 1068–1080, Oct. 2006.
[6] M. V. Jakuba, “Stochastic Mapping for Chemical Plume Source Localization with Application to Autonomous Hydrothermal Vent Discovery,” Ph.D. dissertation, MIT/WHOI, 2007.
[7] S. Thrun, “Learning occupancy grid maps with forward sensor models,” *Autonomous Robots*, vol. 15, no. 2, pp. 111–127, 2003.
[8] J. A. Farrell, J. Murlis, X. Long, W. Li, and R. T. Cardé, “Filament-Based Atmospheric Dispersion Model to Achieve Short Time-Scale Structure of Odor Plumes,” *Environmental Fluid Mechanics*, vol. 2, no. 1, pp. 143–169, Jun. 2002.
[9] A. A. Kurzhanskiy and P. Varaiya, “Ellipsoidal Toolbox,” EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2006-46, May 2006. [Online]. Available: http://code.google.com/p/ellipsoids
[10] “The Network Simulator - NS-2 wiki,” http://nsnam.isi.edu/nsnam/index.php/Main_Page.
[11] A. F. Harris, III and M. Zorzi, “Modeling the underwater acoustic channel in ns2,” in *ValueTools ’07: Proceedings of the 2nd international conference on Performance evaluation methodologies and tools*. ICST, Brussels, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2007, pp. 1–8.