



Open Access Articles

Distributed Data Fusion for Multirobot Search

The Faculty of Oregon State University has made this article openly available.
Please share how this access benefits you. Your story matters.

Citation	Hollinger, G., Yerramalli, S., Singh, S., Mitra, U., & Sukhatme, G. S. (2015). Distributed Data Fusion for Multirobot Search. IEEE Transactions on Robotics, 31 (1), 55-66. doi:10.1109/TRO.2014.2378411
DOI	10.1109/TRO.2014.2378411
Publisher	IEEE - Institute of Electrical and Electronics Engineers
Version	Accepted Manuscript
Terms of Use	http://cdss.library.oregonstate.edu/sa-termsofuse

Distributed Data Fusion for Multi-robot Search

Geoffrey A. Hollinger, *Member, IEEE*, Srinivas Yerramalli, *Member, IEEE*, Sanjiv Singh, *Senior Member, IEEE*, Urbashi Mitra, *Fellow, IEEE*, and Gaurav S. Sukhatme, *Fellow, IEEE*

Abstract—This paper presents novel data fusion methods that enable teams of vehicles to perform target search tasks without guaranteed communication. Techniques are introduced for merging estimates of a target's position from vehicles that regain contact after long periods of time, and a fully distributed team planning algorithm is proposed that utilizes limited shared information as it becomes available. The proposed data fusion techniques are shown to avoid overcounting information, which ensures that combining data from different vehicles will not decrease the performance of the search. Motivated by the underwater search domain, a realistic underwater acoustic communication channel is used to determine the probability of successful data transfer between two locations. The channel model is integrated into a simulation of multiple autonomous vehicles in both open water and harbor environments. The results demonstrate that the proposed distributed coordination techniques provide performance competitive with full communication.

Index Terms—path planning for multiple mobile robot systems, networked robots, distributed robot systems, marine robotics, robotic search

I. INTRODUCTION

Communication between networked robotic vehicles is rarely (if ever) perfect. One method for dealing with imperfect communication is to constrain the movements of the vehicles so that they remain within range, line-of-sight, or both. However, any method that depends on connectivity occurring at a fixed time will be brittle if the model of the communication system is inaccurate. For example, if a planning algorithm requires two vehicles to be connected at given positions, a failure would occur if communication is worse than planned. In reality, communication between robots can be highly variable due to environmental factors, particularly in underwater domains with acoustic communication (see Figure 1) [1]. The variability of communication in many real-world applications motivates the development of algorithms capable of operating with any level of shared information.

G. Hollinger is with the School of Mechanical, Industrial & Manufacturing Engineering, Oregon State University, Corvallis, OR 97330 USA, (e-mail: geoff.hollinger@oregonstate.edu).

S. Yerramalli is with Qualcomm Research, San Diego, CA 92121 USA (e-mail: syerrama@qti.qualcomm.com)

U. Mitra is with the Ming Hsieh Department of Electrical Engineering, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA (e-mail: ubli@usc.edu).

S. Singh is with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 15213 USA (e-mail: ssingh@ri.cmu.edu).

G. Sukhatme is with the Computer Science Department, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA, (e-mail: gaurav@usc.edu).

This research has been funded in part by the following grants: ONR N00014-09-1-0700, ONR N00014-07-1-00738, ONR N00014-14-1-0509, NSF CCR-0120778, and NSF CNS-1035866.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Manuscript received June 14, 2013.




Fig. 1. Communication can be highly variable, and it is often difficult to predict whether two vehicles can share information at given locations without knowing environmental factors. Calm wind and low shipping activity results in a connected network of underwater vehicles (left), but high wind and high shipping activity disconnects the same network of vehicles (right). Utilizing such communication systems during multi-robot planning requires fully distributed algorithms capable of operating at any level of communication.

We explore the problem of multi-vehicle coordination with limited shared information through analysis of the moving target search domain. In this scenario, autonomous vehicles need to locate a moving target using spatially-limited sensing. In cases where disturbances in the environment (e.g., ocean currents) affect plan execution, moving target search with communication limitations requires fusing information when vehicles reconnect after being disconnected. For instance, if a vehicle makes a number of observations and then comes into contact with a vehicle that was out of contact for a long time, sharing the entire history would be costly both in terms of communication and in terms of computation required to fold the observations into each vehicles's current information map. Thus, solving this problem requires the development of data fusion techniques and corresponding coordination methods.

The key novelty of this work is the introduction and analysis of data fusion techniques for moving target search tasks. The ability to fuse data from vehicles that have been disconnected for long periods of time enables distributed path planning that operates at varying levels of shared information.

II. RELATED WORK

The study of target search dates back to classical optimal search theory [2]. In this early work, the goal was to perform maritime search operations (e.g. to find a lost nuclear bomb, submarines, shipwrecks, or people lost at sea). Many of these algorithms improved search efficiency and were successfully used by naval search teams. However, classical work in search theory did not consider communication limitations imposed on modern underwater vehicles.

A large body of multi-robot coordination research simplifies the problem by assuming perfect communication. In many cases, it is possible to show performance guarantees on decentralized algorithms for multi-robot search tasks if perfect communication is available [3]. In some domains, communication may be good enough to allow for this simplifying assumption.

However, many communication systems, such as underwater acoustic modems, are extremely noisy and sensitive to a large number of noise sources [4]. Thus, an assumption of perfect communication is often unrealistic.

An alternative to assuming perfect communication is to assume that an “on/off” communication model is available. Such a model assumes that certain configurations are guaranteed to allow two robots to communicate, and other configurations remove all possibility of communication [5], [6]. Algorithms for on/off communication have been implemented on teams of ground robots [7], and it is possible to develop fully distributed approaches [8]. However, these approaches rely on the requirement that the configurations capable of communication are known before execution.

In some cases, communication maps can be built online to determine if connectivity is possible in a given configuration [9]. However, such techniques require training data, which is particularly problematic when environmental conditions are changing. Recent work has also relaxed the requirements for full connectivity by allowing vehicles to lose communication for portions of the task [10]. This prior work still assumes that communication between all vehicles is available at pre-specified points in the plan.

Moving towards more realistic communication modeling, researchers have examined end-to-end bit error rate metrics [11] and bandwidth limitations [12] for maintaining connectivity in mobile robotic networks. Sophisticated communication models have also been utilized to improve teleoperation methods [13]. In our own prior work, we adapted communication channel and transceiver models from the literature to determine their effect on station keeping in a robotics application [1]. We also applied similar models to data collection in underwater robotic sensor networks [14]. In the current paper, we derive a communication system approximation for use in simulating multi-robot coordination in the moving target search domain.

Combining information from team members that have been disconnected requires a method for fusing data between them. The distributed data fusion problem has been previously examined for estimation problems using Gaussian distributions [15]. Our work allows for arbitrary distributions, though we limit the scope to target search problems.

In the related domain of simultaneous localization and mapping (SLAM), researchers have examined the problem of decentralized coordination and estimation [16]. In this prior work, the robots share their history of measurements and odometry, which grows over time as the robots are disconnected. The authors show that a centralized-equivalent estimate can be calculated by taking advantage of the Markov Property in an Extended Kalman Filter. A similar technique might be feasible in the moving target search domain given sufficient communication and computation. As an alternative, we propose a technique that shares belief states, which avoids the communication requirements of sharing the entire knowledge set as well as the computation required to calculate a centralized-equivalent estimate. As a tradeoff, our technique generates an approximation to the centralized-equivalent belief, which we show to be effective in solving the moving

target search problem.

A preliminary treatment of this work appeared in a prior conference paper [17]. The journal version includes additional theoretical results, extended simulations and experiments, and more detail on the proposed methods.

III. PROBLEM SETUP

A. Bayesian Data Fusion

To enable efficient information sharing, we present data fusion techniques for moving target search. Our technique allows vehicles that have been disconnected for a long period of time to fuse information when they later become reconnected. The proposed fusion rule extends prior work in decentralized data fusion [15] by allowing the use of an objective function that is not modeled as a Gaussian distribution.

The general decentralized data fusion framework estimates some feature of interest (e.g., a target’s location) described by a state vector \mathbf{x}_t , where t denotes the current time. The feature is modeled using a probabilistic state transition $\Pr(\mathbf{x}_t|\mathbf{x}_{t-1})$, which is assumed to be Markovian. Observations \mathbf{z}_t are received that provide information about the state \mathbf{x}_t . A model of the sensor likelihood function is also known that provides $L(\mathbf{z}_t|\mathbf{x}_t)$ given the state at the time of the observation. The Bayesian filtering problem is to find a posterior estimate $\Pr(\mathbf{x}_t|\mathbf{Z}^t, \mathbf{x}_0)$ given observations up to and including time t (denoted by \mathbf{Z}^t) and an initial state estimate \mathbf{x}_0 . Using the recursive Bayes’ rule, calculation of the posterior estimate takes the following form:

$$\Pr(\mathbf{x}_t|\mathbf{Z}^t, \mathbf{x}_0) = \eta L(\mathbf{z}_t|\mathbf{x}_t) \sum_{\mathbf{x}_{t-1}} \Pr(\mathbf{x}_t|\mathbf{x}_{t-1}) \Pr(\mathbf{x}_{t-1}|\mathbf{Z}^{t-1}, \mathbf{x}_0), \quad (1)$$

where η is a normalizing constant.

Equation (1) can be separated into a predictive component in which $\Pr(\mathbf{x}_t|\mathbf{x}_{t-1})$ is applied, and an information fusion component in which $L(\mathbf{z}_t|\mathbf{x}_t)$ is applied. The calculation of (1) becomes a decentralized data fusion problem when different vehicles i and j receive different measurement histories \mathbf{Z}_i^t and \mathbf{Z}_j^t , and wish to reconcile them into a common estimate $\Pr(\mathbf{x}_t|\mathbf{Z}_{i \cup j}^t)$. In this case, there is some redundant information between the estimates $\Pr(\mathbf{x}_t|\mathbf{Z}_{i \cap j}^t)$. If the redundant information is known, the fused estimate can be calculated in closed form [15]. However, in many cases the redundant information is not known because it has already been folded into the estimate. Thus, recovering the true distribution can require storing a large number of measurements and reapplying the filtering steps. For increasingly large teams, the combinatorics of such perfect fusion becomes infeasible.

In the case where the redundant information is not known, it is desirable to develop an estimate of the fused distribution that always avoids overcounting the redundant information $\Pr(\mathbf{x}_t|\mathbf{Z}_{i \cap j}^t)$. Such an estimate is known as a *conservative* estimate because it will never become more sure of the target’s position than is warranted by the measurements. For the case of Gaussian distributions, it is possible to achieve a conservative estimate using a weighted combination of the disparate estimates and covariances [18]. However, particularly

in target search applications, the distribution of interest cannot be modeled using these assumptions. We present a method below that provides a conservative fusion method for moving target search using minimal computation.

B. Moving Target Search

Our specific goal is to locate a target of interest in a known environment with a team of autonomous vehicles. This formulation applies to locating a lost target (e.g., a submarine) and to locating features of interest, such as an area of scientific interest at an unknown location. We assume that the environment has been discretized into N cells such that each cell location represents a vertex on a graph, and traversable connections between those locations represent edges (see [3] for a discussion of discretization methods). We also assume that this graphical representation of the environment is known to the searchers.

For the purposes of this paper, we assume that a vehicle located in a given cell has the same sensing capabilities regardless of its exact position in the cell (e.g., it can sense targets in its own cell and perhaps some adjacent cells). The choice of the coarseness of discretization is determined by the sensing capabilities of vehicles in addition to the available computation. Our formulation allows for false negatives (i.e., a target may not be found even if it is in the same cell as a searcher), but we assume that false positives are negligible (i.e., a target will not be identified unless one actually exists in the cell). This is often a reasonable assumption in ocean search scenarios because any potential target can be inspected more closely once it has been identified, and the time scale of such additional inspection is substantially smaller than the time it takes to move to a new location. We also assume that the search game is over when a target is found (i.e., captured).

We are given K vehicles to search for a target that moves between the cells in the graph. The target's state at time t is represented as a belief vector $b(t) = [b_0(t), b_1(t), \dots, b_N(t)]$, where $b_0(t)$ is the probability that the target has been found prior to time t , and $b_1(t)$ through $b_N(t)$ are the probabilities that the target is in cells 1 through N respectively and has not yet been found. We will refer to the state with belief $b_0(t)$ as the capture state.

We note that the capture state represents the probability that the target has been found *in all possible worlds*. For instance, if one were to run the search game a large number of times, the capture state would be equivalent to the number of times the target was found prior to time t divided by the total number of trials. This is different from representing the probability that the target is found in the current world, which, since the false positive rate is considered to be negligible, would always be either zero or one. A key aspect of the proposed method is that the searchers optimize the capture probability over all possible worlds rather than in the current instantiation of the search game. This leads to high performance on average, which is desirable when the search is run many times (e.g., in the persistent underwater search domains of interest).

Given this formulation of the capture state, we can mathematically represent a capture event on the belief vector by

defining a matrix that moves probability from all cells visible from searcher k 's current cell $s_k(t)$ to the capture state. Let β_n^j be the probability that a target is found in cell n given that a searcher is in cell j . When a searcher is located in cell j , the belief update is given by:¹

$$b_n(t+1) = (1 - \beta_n^j)b_n(t) = b_n(t) - \beta_n^j b_n(t), \quad (2)$$

$$b_0(t+1) = \sum_{n \in \{1, \dots, N\}} \beta_n^j b_n(t) + b_0(t). \quad (3)$$

This belief update can be encoded into a capture matrix C_j , which is applied at time t as $b(t+1) = C_j b(t)$. For example, if it is assumed that the searcher has perfect sensing and can only see within the cell it is located, the capture matrix for a searcher in cell one would be the $(N+1) \times (N+1)$ identity matrix with the second row unity value shifted to the first row.

Similarly, we can define dispersion matrices to represent the expected motion of the target in the environment. The discretization of the environment yields an undirected graph of possible target movements between cells. Let α_{ij} be the probability that a target moves from cell i to cell j (note that $\sum_j \alpha_{ij} = 1$, for all i). The dispersion matrix is defined as the matrix that applies the following rule to each cell n :

$$b_n(t+1) = \alpha_{1n}b_1(t) + \dots + \alpha_{Nn}b_N(t). \quad (4)$$

We note that utilizing a dispersion matrix assumes that the target's motion model obeys the Markov Property. The dispersion matrix D at time t can be applied to yield a new target state vector at time $t+1$ as $b(t+1) = Db(t)$. Finally, we can apply both the dispersion and capture matrices to yield an updated belief vector at the next time as $b(t+1) = \prod_{j \in O} C_j Db(t)$, where O is the set of cells occupied by searchers.

The objective function $J(S)$ takes a set of planned paths from K vehicles $S = \{S_1, S_2, \dots, S_K\}$ and returns the expected utility gained. A commonly used objective function for moving target search is the discounted probability of locating the target² or feature of interest [3], [19]:

$$J(S(0), \dots, S(T)) = \sum_{t=0}^T \gamma^t b_0(t), \quad (5)$$

where $S(0), \dots, S(T)$ are the vehicles' planned paths, γ is a discount factor, and $b_0(t)$ is the the probability that the target had been found prior to time t , and T is some arbitrary end time. The vehicles' planned paths determine the β_i^j values applied by the capture matrices (Equation 3), which fully define the value of the capture state.

It is important to note that in the application domains of interest (e.g., underwater search) the planned paths cannot be executed exactly because of disturbances (e.g., wind, ocean currents, etc.), and the vehicles may visit unplanned locations

¹Readjusting the belief distribution to sum to one (Equation 7) could be used in place of Equation 3 for the capture matrix update. However, Equation 3 provides additional intuition about the role of capture events and is also used in the theoretical analysis in Section IV-A.

²Note that the discounted probability of capture objective function possesses useful theoretical properties, including submodularity, which lead to approximation guarantees on a number of efficient algorithms and generally relate to high performance of distributed algorithms [3].

or execute the path more quickly or more slowly than expected. In these cases, the function $J(S)$ provides an estimate of the utility provided by a planned path. Imperfect prediction of future actions limits the performance of a pre-planned coordination method where the vehicles determine their paths beforehand and then simply execute them without requiring additional communication (see Section V-B).

We assume that the target is found (and reward received) regardless of whether that information has been communicated to the entire team. With imperfect communication, the vehicles will not have updated paths for all vehicles, and hence will not be able to compute the objective exactly. Instead, they must approximate the objective based on their own limited information.

In our moving target search formulation, a capture event is an observation of the target's state and corresponds to a measurement \mathbf{z}_t . The dispersion matrices apply the motion model $\Pr(\mathbf{x}_t|\mathbf{x}_{t-1})$. In fact, when false positives are not considered, the application of the dispersion matrices and capture matrices becomes equivalent to the Bayesian update [3]. Thus, we have a Bayesian information fusion problem similar to those presented in the literature.

IV. DISTRIBUTED DATA FUSION AND COORDINATION

A. Data Fusion for Moving Target Search

We now present data fusion rules for locating a moving target. Given the objective function above, we can formulate a data fusion problem for moving target search. The belief vector $b(t)$ describes the probability that the target is in each possible state of \mathbf{x}_t in the decentralized data fusion framework.

We examine a class of fusion rules characterized by element-wise modifications of the probabilities at each cell. At a given time t , a vehicle k is fusing data from I other vehicles. We define a class of fusion rules where for each cell n , we update the belief as $b_n^k(t) = \nu_0 b_n^0(t) + \nu_1 b_n^1(t) + \dots + \nu_I b_n^I(t)$. The fusion rule is defined by the ν_0, \dots, ν_I values, which weight the belief estimates from the different vehicles. Once this update rule has been applied to all beliefs, the capture state is renormalized as in Equation 7. This class of fusion rules is useful to examine because the fused value of each cell n only depends on the corresponding beliefs for that same cell. This property allows for partial beliefs to be transmitted without affecting the fusion rule.

1) *Minimum Fusion Rule:* Before introducing our proposed fusion rule, we will require some additional notation. When vehicles become disconnected, it is assumed that each of them maintains its own belief vector based on the application of its own capture matrices and the capture matrices of other vehicles with which it is in contact. Each vehicle k 's local belief vector will be referred to as $b^k(t)$. The belief vector that would result if all capture matrices were applied will be referred to as $b^*(t)$. If the vehicles have perfect communication, $b^k(t)$ will be equal to $b^*(t)$ for all k .

As long as the vehicles remain connected, they can share their locations, and the appropriate capture matrices can be applied. If the vehicles are disconnected and then regain connectivity, the following rule is introduced to merge any

number of vehicle estimates into a single merged estimate $b^k(t)$:

$$b_n^k(t) = \min_i b_n^i(t), \quad (6)$$

where \min_i is chosen over all vehicles within communication range of vehicle k . Equation (6) is applied over all $n \in \{1, \dots, N\}$. After applying the rule in (6), the belief vector must be renormalized. If a standard renormalization is applied (i.e., divide all elements by the total probability), the value of the capture state could be either increased or decreased. This renormalization is undesirable because the capture state encodes the probability that the target was found prior to time t in all possible worlds. Clearly, this value can only increase as more observations are made. An alternative that avoids this drawback is to adjust the capture state directly to reflect the estimates in each cell after the application of the minimum rule:

$$b_0^k(t) = 1 - \sum_{n=1}^N b_n^k(t). \quad (7)$$

If multi-hop communication is allowed, the resulting distribution $b^k(t)$ will be shared by all vehicles connected to vehicle k (perhaps after some delay due to communication update rates). However, if only single-hop communication is possible, different vehicles may maintain different estimates even after a merge, due to having different neighbors. In this case, each vehicle's merged distribution contains the minimum probability that a target is in each cell given its immediate neighborhood. However, after several subsequent merges, the multi-hop information will be propagated through the network, and all indirectly connected vehicles will eventually share the same distribution [18].

The intuition behind the minimum fusion rule is that each observation will reduce the probability of the target being in one or more cells, which necessarily increases the probability in the capture state. The vehicles will want to regain lost information from missed observations using a fusion rule that decreases the probability in the environment and increases the probability in the capture state. The next section will show that the merged distribution has several desirable properties relative to the true distribution.

2) *Analysis of the Fusion Rule:* It will now be shown that the minimum fusion rule never overestimates the probability that a target is captured, and it never underestimates the probability that a target is in a cell. Overestimating capture would lead to search schedules that avoid areas that would be searched in the optimal schedule, which is undesirable. The following assumptions are made for this analysis.

Assumption 1: The initial belief $b^i(0)$ equals $b^j(0)$ for all vehicles i and j . That is, all vehicles start with the same belief over the target's state.

Assumption 2: The dispersion matrix D is known to all vehicles. That is, all vehicles have the same model of the target's behavior.

Assumption 3: The false positive rate is negligible for all vehicles and all cells. The false negative rate is incorporated into each capture matrix C_j using non-unity values in the corresponding cells.

We now show that the minimum fusion rule will never underestimate the probability that a target is in a given cell.

Theorem 1: The value $\min_i b_n^i(t)$ is greater than or equal to $b_n^*(t)$ for any cell $n \in \{1, \dots, N\}$, any time t , and any number of vehicles included in \min_i .

Proof: The argument is that the capture matrices C_j and the dispersion matrix D are monotone on the appropriate subvectors of the target's belief (i.e., the capture state and non-capture states respectively). Without loss of generality let n be an arbitrary cell in the environment. The proof will be by induction on $b_n(t)$ (i.e., the belief at time t for cell n). By assumption, $b_n^i(0) = b_n^*(0)$ for all vehicles i and cells n . We now show that $b_n^i(1) \geq b_n^*(1)$.

By assumption, each vehicle applies the same dispersion matrix D to the belief vector. Let $b_n^i(0^+)$ be the belief for vehicle i and cell n after the application of the dispersion matrix to $b_n^i(0)$ but before the application of any capture matrices. Similarly, let $b_n^i(1^-)$ be the corresponding belief after the application of the first capture matrix to $b_n^i(0^+)$. Finally, let $b_n^i(1)$ be the corresponding belief after the application of all capture matrices.

After the application of the dispersion matrices, we have $b_n^i(0^+) = b_n^*(0^+)$. If communication is perfect, all capture matrices will be applied to $b_n^i(0^+)$ that are applied to $b_n^*(0^+)$, which leads to $b_n^i(1) = b_n^*(1)$. Due to imperfect communication, one or more capture matrices may not be applied. Let capture matrix C_j be the first capture matrix applied to $b_n^i(0^+)$ that is not applied to $b_n^*(0^+)$. Now, $b_n^i(1^-) = b_n^i(0^+) + \beta_j^i b_n^i(0^+) \geq b_n^*(1^-)$. For each subsequent capture matrix applied or not applied, this inequality continues to hold. Thus, after the application of all capture matrices, $b_n^i(1) \geq b_n^*(1)$.

We now continue with the induction on $b_n^i(t)$. If $b_n^i(t) = b_n^*(t)$, then the argument above holds for $b_n^i(t+1)$. If $b_n^i(t) > b_n^*(t)$ for any o , then the application of the dispersion matrix to $b_n^i(t)$ is a linear combination of smaller values than those used for $b_n^*(t)$. In this case $b_n^i(t^+) > b_n^*(t^+)$. For each capture matrix applied or not applied, the inequality continues to hold as above. Thus, $b_n^i(t+1) \geq b_n^*(t+1)$ for all n . The same argument can be applied to all vehicles i . ■

An immediate corollary is that no vehicle will ever overestimate the belief that the target is found.

Corollary 1: The value of $b_0^k(t)$ is less than or equal to $b_0^*(t)$ at any time t and any vehicle k .

Proof: Immediate from $b_0^*(t) = 1 - \sum_{n=1}^N b_n^*(t)$, Equations 6 and 7, and Theorem 1. ■

The analysis above shows that the vehicles will never overestimate the probability of capture due to using the minimum merging strategy. In addition, the vehicles will never believe that a target is not in a cell when it actually has a high likelihood of being there. Thus, some areas may be searched more often than they would be without the merging, but no area will be neglected due to the merging. As described above, such fusion rules are desirable and are referred to as conservative in the distributed data fusion literature [15].

In addition, the minimum fusion rule can be utilized when a partial map is received from another vehicle. For instance, a vehicle may only share the portion of its belief near its prior path to save communication cost. As long as the information

received is stamped with its location, the partial map can be folded in using the minimum rule on whatever components of the information map that are received. In the context of realistic communication modeling, several packets may be lost, which would correspond to sections of the map. These lost sections would simply not be incorporated into the receiving vehicle's information map.

3) *Alternative Fusion Rules:* The minimum fusion rule is not the only fusion rule possible in this domain that avoids overestimating the probability of the target in a given cell. For instance, a rule that averages over the beliefs of all adjacent vehicles will also avoid overcounting the measurements. However, we will show that the minimum rule provides a more accurate estimate (relative to the L1 error) of the true distribution than the averaging rule. To understand why, recall from the analysis above that missed application of capture matrices will always increase the amount of probability in a given cell. Thus, the average rule throws out more information about capture than the minimum fusion rule.

We now show that the minimum fusion rule provides a smaller approximation error (versus the true belief) than any fusion rule in this class.

Theorem 2: Given the class of fusion rules described by $b_n^k(t) = \nu_0 b_n^0(t) + \nu_1 b_n^1(t) + \dots + \nu_I b_n^I(t)$, the fusion rule with $\nu_{i_{\min}} = 1$ for $i_{\min} = \arg\min_i b_n^i(t)$ and $\nu_i = 0$ for all $i \neq i_{\min}$ minimizes $\sum_{n \in \{0, \dots, N\}} |b_n^*(t) - b_n^k(t)|$, where $|\cdot|$ is the L1-norm.

Proof: The proof follows from the same property of the capture matrices as used in the proof of Theorem 1 (i.e., the capture matrices are monotone on the capture state). As above, we begin with $t = 0$ when by assumption $b_n^k(0) = b_n^*(0)$ for all n and k . By construction, the belief $b^*(1) = \prod_{j \in O} C_j D b^*(0)$, where O is the set of cells occupied by vehicles. Due to imperfect communication, one or more capture matrices may not be applied to $b^k(1) = \prod_{j \in O_i \subseteq O} C_j D b^k(0)$.

We start with the case where the vehicles become reconnected before $t = 2$. By assumption, the same dispersion matrix has been applied by all vehicles at $t = 1$. From Equation 2, we see that the application of any C_j at time $t = 1$ will only decrease the value of $b_n^k(1)$. Thus, the vehicle with minimal $b_n^i(1)$ yields the closest value to $b_n^*(1)$, since $b_n^*(1)$ has applied all relevant capture matrices. Applying this same argument across all n , we get that the sum $\sum_{n \in \{1, \dots, N\}} |b_n^*(1) - b_n^k(1)|$ is also minimized. Finally, since the capture state sums to one, the value $|b_0^*(1) - b_n^k(1)|$ is also minimized.

Now we deal with the case where the vehicles have been disconnected for some time and become reconnected at time $T > 0$. At time T , the value of $b^k(T)$ has been found by the recursive application of $b^k(T) = \prod_{j \in O_i(t)} C_j D b^k(T-1)$. Similarly, $b^*(T)$ would have been calculated through recursively applying $b^*(T) = \prod_{j \in O} C_j D b^*(T-1)$ if all information were available. From Equation 2, we know that the failure to apply any capture matrix C_j yields an increase in the probability at some subset of cells. From Equation 4, we know that the dispersion matrix redistributes existing probability between cells without moving probability to or from the capture state. If more probability exists in a cell n , then more will be redistributed by applying the matrix D one or more times.

Algorithm 1 Distributed coordination with limited communication

```

1: Input: vehicles 1 to  $K$ , mobility graph  $\Psi$ , objective  $J$ ,
   planning horizon  $T$ , initial distribution  $b(0)$ 
2: % Runs in parallel on each vehicle  $k \in [K]$ 
3:  $t \leftarrow 0$ ,  $t_0 \leftarrow 0$ ,  $b^k(t) \leftarrow b(0)$  for all  $k \in [K]$ 
4: while target not found do
5:   % Process any transmissions waiting in the queue
6:   while queue is not empty do
7:     Process received transmission from vehicle  $i$ 
8:     Update  $b^k(t)$  by fusing  $b^i(t)$  and applying  $C_{s_i(t)}$ 
9:     Update path estimate  $S_k^i$  for vehicle  $i$ 
10:  end while
11:  if at replanning location then
12:     $t_0 \leftarrow t$ 
13:    % Generate a subset of informative paths
14:     $\xi = \mathcal{Q}(\Psi_k(t_0, t_0 + T))$ 
15:    % Determine best path given current information
16:     $S_k^* \leftarrow \operatorname{argmax}_{S_k \in \xi} J(b_k(t), S_k, S_k^1, \dots)$ 
17:  end if
18:  Broadcast  $b^k(t)$  and  $S_k^*(t, t_0 + T)$ 
19:  Continue execution of  $S_k^*(t, t_0 + T)$ 
20:   $b^k(t+1) \leftarrow C_{s_k(t)} D b^k(t)$ 
21:   $t \leftarrow t+1$ 
22: end while

```

Thus, the same argument from the case above (where the vehicles reconnect quickly) applies here. ■

B. Coordination with Limited Communication

We now present a distributed technique for coordinating teams of vehicles to locate lost targets under limited communication. Our approach is passive with respect to the communication limitations, in that it does not directly utilize the communication system approximation as part of planning. Instead, our approach uses implicit coordination [3], where vehicles share their plans and information maps to improve the team plan. If information is not available due to communication limitations, each vehicle plans without that information.

We will denote S_k^i as vehicle k 's current estimate of vehicle i 's plan, and we will denote the set $\{1, \dots, K\}$ as $[K]$. All possible feasible paths for vehicle k from times t_0 to time t_1 are denoted as $\Psi_k(t_0, t_1)$. Algorithm 1 gives a summary of the distributed planning approach. The vehicles plan their own paths and broadcast both their intended plans as well as their current estimate of the target's position. The vehicles use a path generation technique represented by function $\mathcal{Q}(\Psi_k(t_0, t_1))$ that takes as input a set of possible paths and returns a subset of those paths. The function \mathcal{Q} can be deterministic (e.g., enumeration of all paths) or stochastic (e.g., random selection of paths) or based on some heuristic. We assume that low-level collision avoidance is used to avoid inter-robot collisions, which is reasonable when the maneuverable space is large relative to the density of the robots (e.g., in ocean search scenarios).

Each vehicle incorporates the plans and estimates received from teammates. If two vehicles are connected, they have

an updated estimate of each other's state and can apply the appropriate capture matrix. If they are not connected and later become connected, they apply the minimum fusion rule as described in Section IV-A. Plans and information maps that are not available are simply not utilized, possibly decreasing the effectiveness of the planned path. Thus, the algorithm can operate at any level of communication. We note that the order of processing transmissions waiting in the queue (line 7) does not affect the outcome of the planner, due to the application of the minimum fusion rule and the path planning being held off until the entire queue is processed.

When running the proposed algorithm, each vehicle optimizes its own plan given the information shared by its teammates. This approach avoids planning in the joint space of plans, represented by the cross product of all the vehicle's planning spaces, which grows exponentially in the size of the team. Such decoupled planning techniques have been shown to perform near-optimally in target search domains and to outperform competing heuristics [3].

V. SIMULATIONS AND EXPERIMENTS**A. Ground Vehicle Search**

Our first simulations test the effectiveness of the proposed fusion rule using ground vehicle search scenarios. This domain provides understanding of the basic behavior of the proposed algorithm. Underwater search, the main focus of this paper, appears in the following section. We ran ground vehicle search simulations in three environments with range and/or line-of-sight communication constraints. These simulations utilized a multi-robot simulation environment implemented in C++ on Ubuntu Linux running on a 3.2 GHz Intel i7 processor with 9 GB of RAM.

In the simulations, a team of autonomous ground vehicles searched for a target that started in a random cell and moved randomly to any adjacent cell at a speed of 0.1 m/s. To adjust for different cell sizes in the dispersion matrices, the probability of the target remaining in its current cell was increased for larger cells (a target moving at a fixed speed would be more likely to remain within the boundaries of a larger cell after any given time step), and the remaining probability was distributed evenly to adjacent cells. The ground vehicles started together in a random cell (cells were big enough to accommodate multiple vehicles) and moved with a constant speed of 1 m/s. A factor of ten difference between the searchers' speeds (1 m/s) and the target's speed (0.1 m/s) was based on relevance to urban search scenarios. Additional simulations (not shown) demonstrated that the proposed methods provide improvement at other relative speeds as well.

The environments were discretized as shown in Figure 2, and the searchers were equipped with simulated sensors capable of detecting targets within the same discrete cell. For these simulations, we used simple models of line-of-sight and range communication that often appear in the literature. The range constraints were set as 1/4 the diagonal of the map, and the obstacles were set to impede line-of-sight communication. In the next section, we will incorporate more sophisticated communication modeling.

The ground vehicles moved holonomically between the centroids of the discretized cells in the environment. To model possible disturbances in the environment (e.g., impassable objects or detours), a random probability (uniformly distributed between 0 and 1) of moving to a cell adjacent to the intended goal was added to the ground vehicle simulation. Complex kinematics and dynamics were not modeled, since the goal of these simulations was to evaluate the data fusion and coordination techniques, but such constraints could be incorporated through the generation of feasible plans during the path generation stage of the coordination algorithm. Conflicts between vehicle paths were handled using low-level collision avoidance.

The searchers utilized Algorithm 1 to plan paths that maximize the discounted probability of finding the target. The vehicles used a planning horizon of four steps and an exhaustive enumeration to that horizon as the path generation function \mathcal{Q} . Searchers were allowed to coordinate with other searchers within the communication constraints. Searchers that could not communicate did not share beliefs and did not integrate measurements received by team members outside of their communication range. When a merging rule was used, searchers that regained communication merged their estimates using the minimum fusion rule proposed in Section IV. When a merging rule was not used, searchers communicated their measurements and paths when connected and incorporated them into their current beliefs and plans. However, the “no merging” case did not attempt to recover measurements from periods of disconnection even after the searchers had reconnected.

Figure 3 shows the effectiveness of using the proposed merging rule to locate a randomly moving target in the indoor and outdoor environments shown in Figure 2. We note that the path lengths and expected capture times are equivalent for a fixed speed because the search game ends when the target is captured. In all environments, the proposed minimum fusion rule decreases the expected time to capture the target. In many cases, the proposed technique yields expected capture times nearly as low as if full communication were available.

Figure 3 also compares to a fusion rule that averages over all available belief estimates for a given cell. This rule fits within the class of fusion rules defined in Theorem 2. These simulations confirm that the minimum fusion rule improves the performance of the search task for line-of-sight and range-limited communication models. The benefit from using the minimum fusion rule over the average fusion rule is particularly pronounced in the indoor environment, which is likely due to infrequent reconnections between the team members in this environment.

B. Underwater Search

In this section, we apply our techniques to underwater search domains where communication is available through acoustic modems. To properly evaluate the proposed approach in these domains, we derive principled estimates of packet error rate based on well-accepted models of acoustic channels, and we validate these models using data from an autonomous




Fig. 2. Indoor and outdoor environments used for simulated coordinated search with ground vehicles. The indoor 60 m \times 60 m SDR (left) map was discretized into 188 free cells using an orthogonal region growing method, and the 200 m \times 150 m outdoor MOUT (right) map was discretized into 227 free cells using a constrained Delaunay triangulation [20]. Green lines denote discretization boundaries.

underwater vehicle (AUV) deployment. We then test the proposed coordination and data fusion techniques through simulations built off these communication models.

1) *Acoustic Communication Modeling:* Underwater acoustic channels are characterized by a path loss that depends not only on the distance between the transmitter and receiver, but also on the signal/carrier frequency. The carrier frequency determines the absorption loss $A(d, f)$ due to the transfer of acoustic energy into heat in the medium. Relying on extensive experimental data, an empirical formula for the path loss for a distance d and frequency f (in kHz) is given in [4].

Noise in underwater acoustic channels is determined by several factors, such as turbulence, the shipping activity in the surrounding region, the surface motion caused by wind-driven waves, and finally thermal noise. The constant surface motion due to wind driven waves are a significant factor contributing to the noise at the operating frequencies of interest for underwater systems (100 Hz - 100 kHz). This noise can be modeled for a given frequency as $N(f)$ using equations from prior work [4].

Overall, the acoustic channel is noise limited at very low frequencies and attenuation limited at high frequencies. For moderate signaling bandwidths B and transmitted power P , the average signal to noise ratio (SNR) at the receiver at a distance d and frequency f is then

$$SNR(d, f) = \frac{P}{A(d, f)N(f)B}. \quad (8)$$

For a Rayleigh fading sub-channel, the probability of error for a moderate to large SNR's can be approximated as

$$P(e) \approx \frac{1}{4SNR(d, f)}. \quad (9)$$

The probability of symbol error on one sub-channel, $P(e)$, is a function of the transmitted power, the frequency of transmission, wind speed, shipping factor, the distance and the bandwidth used. For an uncoded packet with symbols over M sub-channels, the probability of packet error can then be computed as

$$P_{packet} = 1 - (1 - P(e))^M. \quad (10)$$

Given the system approximation described above, we can calculate the probability of error for transmitting an estimate of the target's distribution (i.e., transmitting $b^i(t)$ from vehicle

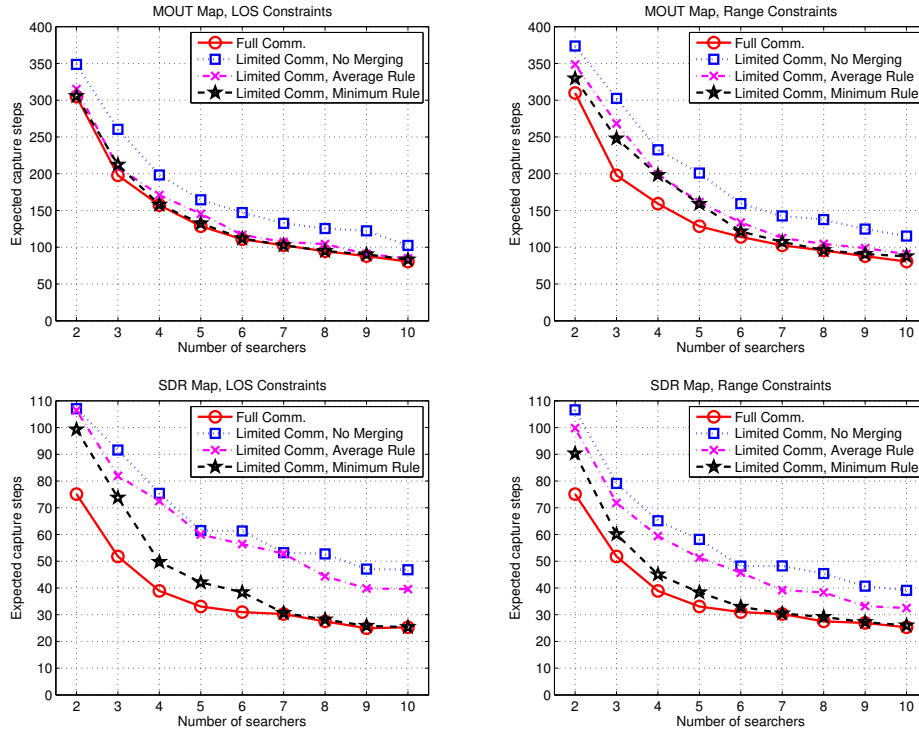


Fig. 3. Comparison of full communication to limited communication with and without the use of data fusion merging rules. The merging rules allow searchers to share information when they regain communication, which improves the performance of the search. Results are averaged over 20,000 trials with a randomly moving target. Error bars are small and therefore omitted for readability.




Fig. 4. Modeled data error rates for varying wind speeds. The curve gives the probability that an error will occur when transmitting the probability distribution of the target's estimated position.

i to vehicle j), given the distance between the vehicles and an estimate of the wind speed and shipping activity. Figure 4 shows the data error rates with the system parameters in Table I and varying wind speeds.

2) *Experimental Validation of Model*: We have validated the model above using data from the SeaBED-class AUV *Lucille* [21], which was equipped with a WHOI Micro-Modem [22] for acoustic communications. In September of 2010, *Lucille* assisted in mapping the submerged portion of the San Andreas Fault off Northern California, at approximately $39^{\circ}50'N, 124^{\circ}W$. During this survey, the AUV's onboard

networking stack periodically transmitted packets to a surface vessel. The experimental results demonstrate the robustness and accuracy of the proposed model. Details on the model validation are available in [14].

3) *Underwater Search Simulations*: We next test our algorithms in a simulated underwater domain that utilizes the acoustic communication models described in the previous section. The simulated underwater vehicles moved at 5 km/hr and had a detection radius of 200 m (motivated by the swath width for a side scan sonar). As in the ground vehicles simulations, we did not impose complex kinematic or dynamic constraints on the underwater vehicles. Such constraints could be incorporated into the coordination algorithm through the generation of feasible paths during each individual vehicle's path optimization stage. As before, the target started in a random cell and moved at 0.5 km/hr to any adjacent cell; the searchers started together in a random cell.

Table I gives the communication system specifications. In these simulations, the vehicles broadcasted their plans, current location, and estimated target distribution when they arrived at each replanning point. The model described above, which incorporates distance, line-of-sight, wind, and shipping activity, determined the probability that each vehicle received the broadcast. Wind and shipping activity were set to vary randomly throughout the map. The wind also affected the vehicles' movement. Higher wind corresponded to a higher probability of being blown off course and searching an area adjacent to the one intended. It was assumed the searchers knew they were blown off course and could adjust their belief update accordingly. Wind speed of 10 m/s caused a 100%

TABLE I
PARAMETERS FOR UNDERWATER ACOUSTIC SIMULATIONS

Parameter	Value
Transmission Frequency	13 kHz
Transmission Power	1 W
Bandwidth	1 kHz
Packet Size	256 symbols
Wind speed	1 m/s - 10 m/s

chance of moving to a cell adjacent to the one intended, and this probability was scaled down linearly to zero with decreasing wind speed. Such disturbances are typical when operating autonomous underwater vehicles, which often surface several kilometers from their intended goals [23].

The proposed coordination and data fusion techniques were validated using a search problem at a fixed depth. A non-adversarial target existed in the environment, and the underwater vehicles searched for its location. The proposed framework allows for both moving and stationary targets. In these simulations, the target was assumed to move randomly. The target's depth was assumed to be known a priori (e.g., it existed on the ocean floor). Known depth reduces the planning problem to 2D; however, the same algorithms could be applied to a 3D problem with an expanded environment graph. Figure 5 shows surface maps of island and harbor environments of varying sizes used for simulated testing. The land masses served as obstacles that prevented both communication and movement. It was assumed that any target that left the map was considered lost (i.e., it could never be captured). Lost targets were not considered in the cost function.

The traversable ocean portions of the maps were discretized into $200 \text{ m} \times 200 \text{ m}$ regular grid cells, and the distributed planning algorithm from Section IV-B was run with perfect communication and then with the communication system model from Section V-B1. In both the perfect communication case and the limited communication case, the underwater vehicles used a planning horizon of $T = 1$ for ease of comparison to less scalable methods. Additional simulations (not shown) showed a small (less than 10%) reduction in capture time for increased horizon lengths. An exhaustive enumeration to the horizon was used as the path generation function \mathcal{Q} . We note that with perfect communication, capture matrices from all vehicles were applied at each time step, and data fusion was not required to estimate the belief vectors.

For comparison, an algorithm was also implemented that maintains full connectivity at all times. The algorithm looks one step in the future and chooses the next location with the highest probability of capture for the team that also ensures that no vehicle is disconnected from the rest of the team. The advantage of this technique is that the searchers can maintain the same belief distribution at all times; however, they must remain connected to do so. The continual connectivity technique requires planning in the joint space, which is exponential in the number of vehicles. Thus, longer planning horizons are not computationally feasible, and the one-step horizon was used for comparison. Based on a worst-case assumption of wind speed and shipping activity, the threshold that would guarantee

connectivity was set to 2 km.

Figure 5 shows quantitative results from the simulations, and Multimedia Extension #1 shows an animation of the search strategy in the harbor environment. The results demonstrate that distributed coordination with the proposed communication model performs almost as well as coordination with perfect communication on both maps. In contrast, constraining the vehicles' paths such that they need to continually maintain connectivity increases the expected time to locate the target. The difference is more significant in the cluttered Long Beach Harbor map, where communication is more difficult to maintain. In this environment, it is beneficial for vehicles to break connectivity, search for the target individually or as sub-teams, and later share information. The distributed coordination and data fusion techniques allow for this behavior, which leads to improved performance. These results demonstrate that the benefits of breaking connectivity outweigh the benefits of maintaining a more accurate estimation of the target's position by constraining the search.

The results in Figure 5 also compare to a pre-planned coordination method where the vehicles used a centralized solver to simulate the target belief forward and generate coordinated paths for the entire mission (using the same algorithm as the full communication case). The vehicles then attempted to execute these paths. Due to disturbances from wind (see Table I), the vehicles were not able to execute those paths exactly, which resulted in the vehicles moving more quickly, more slowly, or searching unexpected regions. This pre-planned coordination method did not provide competitive performance with the proposed method, which highlights the need for data fusion in this domain. The improvement over the pre-planned method is more pronounced in the less cluttered environment due to the more reliable communication available to the proposed method to adjust the teams' paths appropriately.

Figure 6 shows a comparison between the estimate of the capture state for the average fusion and the minimum fusion rule. As predicted by the analysis in Theorem 2, the minimum fusion rule provides an estimate that is closer to the one achieved using full communication. The KL divergence of the estimated belief distributions relative to full communication is also shown in Figure 6. It is interesting to note that there is an initial spike in the KL divergence after the team first becomes disconnected. The KL divergence then settles quickly when AUVs reconnect if the minimum fusion rule is used. This settling takes substantially longer when the average fusion rule is used. In addition, the average fusion rules can lead to a decrease in the value of the capture state, which is undesirable (see Section IV-A). It is also clear from these plots that the benefit of the minimum fusion rule over the average fusion rule increases as the size of the team grows.

The percentage of successful transmissions and percentage of time that the team is fully connected is shown in Figure 7. With two AUVs, nearly 80% of transmissions are successful, and the team is fully connected for the majority of the time. As the number of AUVs increases, the percentage of successful transmissions decreases to near 40% as some vehicles break away from the team. In addition, particularly in the harbor

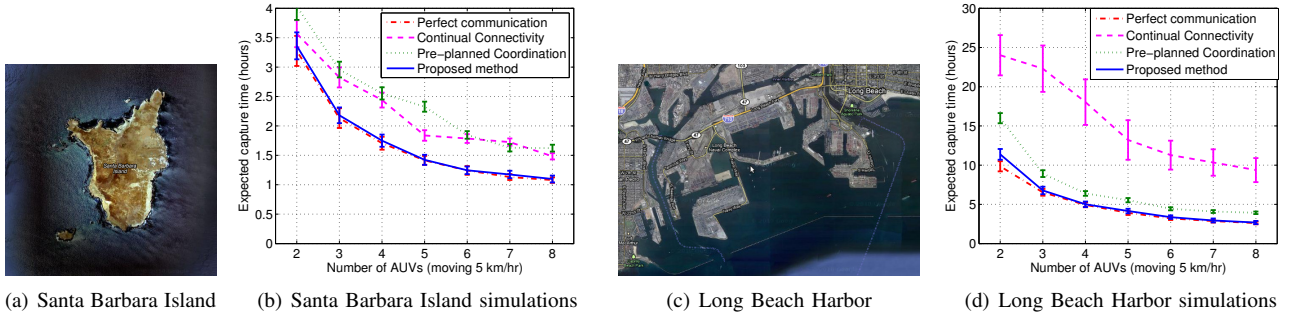


Fig. 5. Simulated results for non-adversarial search in underwater environments with acoustic communications. Santa Barbara Island (left) is $4 \text{ km} \times 4 \text{ km}$ (discretized into 296 free cells) and Long Beach Harbor (right) is $11 \text{ km} \times 8 \text{ km}$ (discretized into 787 free cells). Each data point is averaged over 200 simulations; error bars are one standard error of the mean. The proposed method takes advantage of information as it becomes available and outperforms a method that maintains connectivity at all times. The more cluttered Long Beach Harbor environment increases the advantage of breaking connectivity.

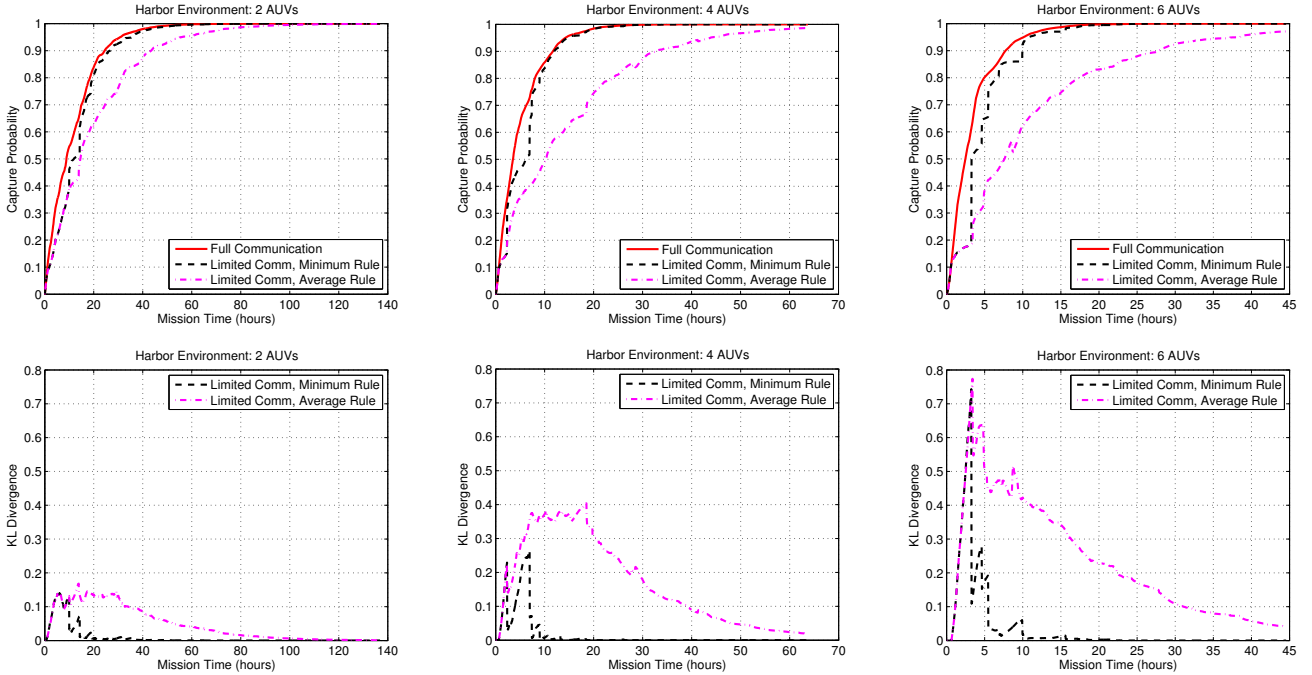


Fig. 6. Comparisons of the belief distribution for a single reference AUV to the belief distribution from full communication in the harbor environment. The minimum fusion rule provides an estimate of the capture state closer to the one found using full communication, and it also achieves a lower KL divergence when compared to the full distribution. The benefit of the minimum fusion rule over the average rule increases as the team size grows.

environment, the AUV team is rarely fully connected (less than 10% of the time in some cases). These results demonstrate that the proposed algorithm yields high performance even when the packet error rate is high and the team is disconnected for a large percentage of the time.

We also examined the mean time the vehicles took to regain full connectivity after becoming disconnected. This value represents the time each vehicle would typically go without a fully updated belief map. We found that the time the team remained disconnected was fairly constant for increasing numbers of vehicles but varied between environments. The mean time without full connectivity was found to be 23 minutes on the Santa Barbara map and 84 minutes on the Long Island Harbor map. Relative to the mission times of 1 to 4 hours on the Santa Barbara map and 2 to 10 hours on the Long Island Harbor map, these values represent substantial periods of time that the team

was not fully connected.

A somewhat surprising observation from these results is that a small amount of opportunistic communication can achieve search performance competitive with full communication. The high performance of data fusion in these scenarios stems from the nature of the search problem where vehicles must coordinate if they are near each other, and hence more likely to be able to communicate. As a result of this beneficial property of the domain of interest, data fusion methods are able to avoid active connectivity maintenance, which is computationally costly and requires additional planning overhead.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a distributed data fusion approach for multi-robot planning with limited communication in the moving target search domain. We proposed and analyzed

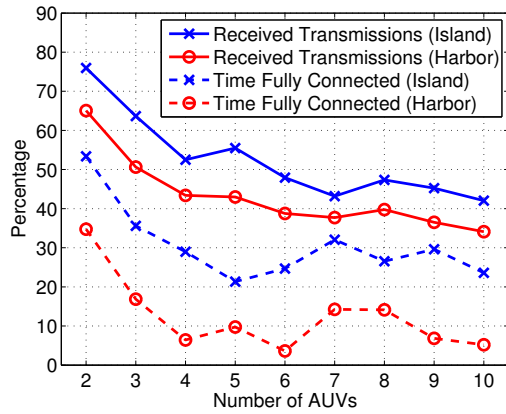


Fig. 7. Percentage of successful transmissions and percentage of time that the AUVs were fully connected when running the proposed solution. The proposed algorithm is able to achieve high performance in the target search task with high packet error rate and with a high percentage of the time disconnected.

data fusion techniques that provide principled methods for incorporating information from robotic vehicles that regain connectivity. Our proposed distributed coordination algorithm utilizes available information to provide solutions robust to changes in communication. Simulated experiments with realistic acoustic communication models demonstrated that it is possible to achieve low average capture times without actively maintaining connectivity between vehicles.

The techniques proposed in this paper move towards fully distributed multi-robot coordination and data sharing. Future work includes further theoretical analysis of performance guarantees with communication limitations, the derivation of more general data fusion techniques, and improved acoustic communication system models. For tasks that require tighter coordination between vehicles, it may be necessary to develop more complex algorithms to remain robust to communication failures. In addition, tasks that require the team to replan often may require new methods to solve when communication is imperfect. Ultimately, this line of research has the potential to enable high-performing multi-vehicle coordination methods that operate at any level of shared information.

ACKNOWLEDGMENT

The authors gratefully acknowledge Joseph Djughash and Benjamin Grocholsky from Carnegie Mellon University and Athanasios Kehagias from Aristotle University of Thessaloniki for their insightful comments.


REFERENCES

- [1] F. Arrichiello, D. Liu, S. Yerramalli, A. Pereira, J. Das, U. Mitra, and G. Sukhatme, "Effects of underwater communication constraints on the control of marine robot teams," in *Proc. Int. Conf. Robot Communication and Coordination*, 2009, pp. 1–8.
- [2] L. Stone, *Theory of optimal search*. New York: Academic Press, 1975.
- [3] G. Hollinger, S. Singh, J. Djughash, and A. Kehagias, "Efficient multi-robot search for a moving target," *Int. J. Robotics Research*, vol. 28, no. 2, pp. 201–219, 2009.
- [4] M. Stojanovic, "On the relationship between capacity and distance in an underwater acoustic communication channel," in *Proc. ACM Int. Workshop on Underwater Networks*, 2006, pp. 41–47.
- [5] M. Zavlanos and G. Pappas, "Distributed connectivity control of mobile networks," *IEEE Trans. Robotics*, vol. 24, no. 6, pp. 1416–1428, 2008.
- [6] F. Arrichiello, J. Das, H. Heidarrsson, A. Pereira, S. Chiaverini, and G. Sukhatme, "Multi-robot collaboration with range-limited communication: Experiments with two underactuated ASVs," in *Proc. Int. Conf. Field and Service Robotics*, 2009, pp. 443–453.
- [7] N. Michael, M. Zavlanos, V. Kumar, and G. Pappas, "Maintaining connectivity in mobile robot networks," in *Proc. Int. Symp. Experimental Robotics*, 2008, pp. 117–126.
- [8] P. Yang, R. Freeman, G. Gordon, K. Lynch, S. Srinivasa, and R. Sukthankar, "Decentralized estimation and control of graph connectivity for mobile sensor networks," *Automatica*, vol. 46, no. 2, pp. 390–396, 2010.
- [9] M. Hsieh, A. Cowley, V. Kumar, and C. Taylor, "Maintaining network connectivity and performance in robot teams," *J. Field Robotics*, vol. 26, no. 1–2, pp. 111–131, 2008.
- [10] G. Hollinger and S. Singh, "Multirobot coordination with periodic connectivity: Theory and experiments," *IEEE Trans. Robotics*, vol. 28, no. 4, pp. 967–973, 2012.
- [11] Y. Yan and Y. Mostofi, "Robotic router formation in realistic communication environments," *IEEE Trans. Robotics*, vol. 28, no. 4, pp. 810–827, 2012.
- [12] Y. Pei and M. W. Mutka, "STARS: Static relays for remote sensing in multi-robot real-time search and monitoring," *IEEE Trans. Parallel and Distributed Systems*, vol. 24, no. 10, pp. 2079–2089, 2013.
- [13] P. R. Giordano, A. Franchi, C. Secchi, and H. H. Bühlhoff, "A passivity-based decentralized strategy for generalized connectivity maintenance," *Int. J. Robotics Research*, vol. 32, no. 3, pp. 299–323, 2013.
- [14] G. Hollinger, S. Choudhary, P. Qarabaqi, C. Murphy, U. Mitra, G. Sukhatme, M. Stojanovic, H. Singh, and F. Hover, "Underwater data collection using robotic sensor networks," *IEEE J. Selected Areas in Communications*, vol. 30, no. 5, pp. 899–911, 2012.
- [15] A. Makarenko, A. Brooks, S. Williams, H. Durrant-Whyte, and B. Grocholsky, "A decentralized architecture for active sensor networks," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2004, pp. 1097–1102.
- [16] K. Y. K. Leung, T. D. Barfoot, and H. H. T. Liu, "Decentralized cooperative SLAM for sparsely-communicating robot networks: A centralized-equivalent approach," *J. Intelligent and Robotic Systems*, vol. 66, no. 3, pp. 321–342, 2012.
- [17] G. Hollinger, S. Yerramalli, S. Singh, U. Mitra, and G. Sukhatme, "Distributed coordination and data fusion for underwater search," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2011, pp. 349–355.
- [18] S. Julier and J. Uhlmann, "General decentralised data fusion with covariance intersection (CI)," in *Handbook of Data Fusion*, D. Hall and J. Llinas, Eds. CRC Press, 2001, pp. 319–343.
- [19] S. Ong, S. Png, D. Hsu, and W. Lee, "Planning under uncertainty for robotic tasks with mixed observability," *Int. J. Robotics Research*, vol. 29, no. 8, pp. 1053–1058, 2010.
- [20] J. Shewchuk, "Delaunay refinement algorithms for triangular mesh generation," *Computational Geometry: Theory and Applications*, vol. 22, no. 1–3, pp. 21–74, 2002.
- [21] H. Singh, A. Can, R. Eustice, S. Lerner, N. McPhee, O. Pizarro, and C. Roman, "Seabed AUV offers new platform for high-resolution imaging," *EOS Trans. AGU*, vol. 85, no. 31, p. 289, 2004.
- [22] L. Freitag, M. Grund, S. Singh, J. Partan, P. Koski, and K. Ball, "The WHOI micro-modem: An acoustic communications and navigation system for multiple platforms," in *Proc. IEEE Oceans Conf.*, 2005, pp. 1086–1092.
- [23] A. Pereira, J. Binney, G. Hollinger, and G. Sukhatme, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *J. Field Robotics*, vol. 30, no. 5, pp. 741–762, 2013.



Geoffrey A. Hollinger is an Assistant Professor in the School of Mechanical, Industrial & Manufacturing Engineering at Oregon State University. His current research interests are in planning, decision making, and learning for robotic systems in marine, air, and ground environments. He has served as PI and co-PI on research grants from ONR, NSF, USDA, and NASA. His past research includes networked marine robotics at the University of Southern California, multi-robot search at Carnegie Mellon University, personal robotics at Intel Research Pitts-

burgh, active estimation at the University of Pennsylvania's GRASP Laboratory, and miniature inspection robots for the Space Shuttle at NASA's Marshall Space Flight Center. He received his Ph.D. (2010) and M.S. (2007) in Robotics from Carnegie Mellon University and his B.S. in General Engineering along with his B.A. in Philosophy from Swarthmore College (2005). He has been a Guest Editor for the *Autonomous Robots* journal, Associate Editor for the *ICRA* and *IROS* conferences, and Area Chair for the *RSS* conference.



Sanjiv Singh is a Research Professor at the Robotics Institute, Carnegie Mellon University and is the CEO and Co-Founder of Near Earth Autonomy, Inc. His recent work has two main themes: perception in natural environments and multi-agent coordination. He has led projects in both ground and air vehicles operating in unknown or partially known environments, in applications such as mining, agriculture, emergency response, surveillance and exploration. He is also actively involved in the automation of complex tasks, such as the assembly of large space

structures, that cannot be addressed by single agents and must necessarily be performed by teams. Prof. Singh received his B.S. in Computer Science from the University of Denver (1983), M.S. in Electrical Engineering from Lehigh University (1985) and a Ph.D. in Robotics from Carnegie Mellon (1995). He is the founder and Editor-in-Chief of the *Journal of Field Robotics*.

Urbashi Mitra received the B.S. and the M.S. degrees from the University of California at Berkeley and her Ph.D. from Princeton University. After a six year stint at the Ohio State University, she joined the Ming Hsieh Department of Electrical Engineering at the University of Southern California, Los Angeles, where she is currently a Professor. Dr. Mitra is a member of the IEEE Information Theory Society's Board of Governors (2002-2007, 2012-2014) and the IEEE Signal Processing Society's Technical Committee on Signal Processing for Communications and Networks (2012-2014). She is a Fellow of the IEEE. Dr. Mitra currently serves on the IEEE Fourier Award for Signal Processing committee and the IEEE James H. Mulligan, Jr. Education Medal committee. She has been/is an Associate Editor for the following IEEE publications: *Transactions on Signal Processing* (2012–), *Transactions on Information Theory* (2007-2011), *Journal of Oceanic Engineering* (2006-2011), and *Transactions on Communications* (1996-2001). Dr. Mitra has held visiting appointments at: the Delft University of Technology, Stanford University, Rice University, and the Eurecom Institute. She served as co-Director of the Communication Sciences Institute at the University of Southern California from 2004-2007. Her research interests are in: wireless communications, communication and sensor networks, detection and estimation and the interface of communication, sensing and control.



Srinivas Yerramalli is currently a Senior Systems Engineer at Qualcomm in San Diego, CA. He received the B.Tech degree from the International Institute of Information Technology in 2007 and joined the M.S. program at University of Southern California in Fall 2007 and the Ph.D. program in Fall 2008. He received his Ph.D. from the Ming Hsieh Department of Electrical Engineering at the University of Southern California in 2013. His research primarily focuses on physical layer communications for highly time-varying channels, especially under-

water acoustic channels, and in understanding the feasibility of cooperation in heterogeneous wireless networks. More broadly, he is interested in problems related to signal processing and optimization for wireless communication systems. He is a recipient of the Ming Hsieh Scholar award for 2011-2012, the IIIT Gold Medal for excellent academic performance in 2007, the Prathibha Scholarship from the Government of Andhra Pradesh from 2003-2007, and a Gold Medal from the President of India for his performance in the All India Engineering Entrance Examination (AIEEE) in 2003.



Gaurav S. Sukhatme is Professor and Chairman of Computer Science (joint appointment in Electrical Engineering) at the University of Southern California (USC). He is the co-director of the USC Robotics Research Laboratory and the director of the USC Robotic Embedded Systems Laboratory which he founded in 2000. His research interests are in robot networks, robot perception and on-body networks. He has published extensively in these and related areas. Sukhatme has served as PI on numerous NSF, ONR, DARPA and NASA grants. He is a fellow

of the IEEE and a recipient of the NSF CAREER award and the Okawa foundation research award. He is one of the founders of the Robotics: Science and Systems conference. He is the Editor-in-Chief of *Autonomous Robots* and has served as Associate Editor of the *IEEE Transactions on Robotics and Automation*, the *IEEE Transactions on Mobile Computing*, and on the editorial board of *IEEE Pervasive Computing*.