In Proceedings, ICRA 2011, May 2011.

Distributed Coordination and Data Fusion for Underwater Search

Geoffrey A. Hollinger, Srinivas Yerramalli, Sanjiv Singh, Urbashi Mitra and Gaurav S. Sukhatme

***Abstract*— This paper presents coordination and data fusion methods for teams of vehicles performing target search tasks without guaranteed communication. A fully distributed team planning algorithm is proposed that utilizes limited shared information as it becomes available, and data fusion techniques are introduced for merging estimates of the target's position from vehicles that regain contact after long periods of time. 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 ocean and harbor search scenarios. The simulated experiments demonstrate that distributed coordination with limited communication significantly improves team performance versus prior techniques that continually maintain connectivity.**

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. These constraints can be maintained continually [1], [2] or periodically [3], [4]. However, any method that depends on connectivity occurring at a fixed time will be brittle if the model of the communication system is incorrect. 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 cannot be predicted by simple heuristics like distance and line-of-sight [5]. The inherent unpredictability of communication in many real-world applications motivates the development of algorithms capable of operating with any level of shared information.

We explore the problem of multi-vehicle coordination with limited shared information through analysis of the underwater moving target search domain. In this scenario, autonomous underwater vehicles (AUVs) need to locate a lost target (e.g., a submarine or enemy underwater vehicle). Moving target search with communication limitations requires both distributed path planning and methods for fusing information gathered when the vehicles were not connected. 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 both coordination and data fusion techniques.

This research has been funded in part by the following grant, ONR N00014-09-1-0700.

G. Hollinger and G. Sukhatme are with the Dept. of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA {geoffrey.hollinger,gaurav}@usc.edu

S. Yerramalli and U. Mitra are with the Dept. of Electrical Engineering, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA {srinivas.yerramalli,ubli}@usc.edu

S. Singh is with the Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 USA ssingh@ri.cm.edu

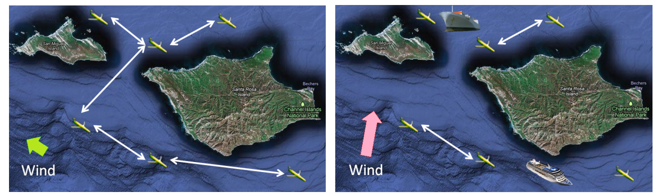

Fig. 1. Communication using underwater acoustic signals is highly imperfect, and it is often difficult to predict whether two vehicles can share information at given locations. Calm wind and low shipping activity results in a connected network of underwater vehicles (left), but high wind and high shipping 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.

In this paper, we examine the problem of moving target search with a team of autonomous underwater vehicles (AUVs) utilizing acoustic communication. We first discuss related work in multi-robot coordination, planning, and search (Section II). We then propose a distributed coordination algorithm in which each AUV optimizes its path given its current information (Section III). In addition, we present and analyze a novel data fusion technique that allows for vehicles to fuse their information (Section IV). Finally, we derive a realistic model for underwater acoustic communication (Section V) and integrate this model into a simulator for validation of the proposed approach versus techniques that require continual connectivity (Section VI). The novelties of this paper include the development of a distributed path planning algorithm capable of operating with varying levels of shared information, the introduction and analysis of data fusion techniques for moving target search tasks, and the incorporation of a system model for underwater acoustic communications into multi-vehicle simulations.

II. RELATED WORK

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 [6], [7]. 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. Thus, an assumption of perfect communication is unrealistic in these cases.

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 [1], [2]. Algorithms for on/off communication have been implemented on teams of ground robots [8], and it is possible to develop fully distributed approaches [9]. However, these approaches rely on the requirement that configurations that are 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 [10]. 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 [3], [4]. This prior work still assumes that communication between all vehicles is available at pre-specified points in the plan.

Arrichiello et al. recently adapted communication channel and transceiver models from the literature to determine their effect on station keeping in a robotic application [5]. We derive a similar communication system approximation in the current paper for use in simulating multi-robot coordination in the moving target search domain.

III. COORDINATION WITH LIMITED COMMUNICATION

We now propose 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 utilizes implicit coordination [7], 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.

Given K vehicles, we will denote the set $\{1, \ldots, K\}$ as $[K]$. All possible feasible paths for robot k from times t_0 to time t_1 are denoted as $\Psi_k(t_0, t_1)$. The function $\mathcal{Q}(\Psi_k(t_0, t_1))$ is a sampling function that takes as input a set of possible paths and returns a sampled set of paths. The function $\mathcal{Q}$ can be deterministic (e.g., enumeration of all paths) or stochastic. In the case of infinite paths, such as when planning in a continuous domain, waypoint-based sampling can be employed. The objective function $J(S)$ takes a set of paths from K vehicles $S = \{S_1, S_2, \ldots, S_K\}$ and returns the expected utility gained by those paths. Let $b_k(t)$ represent vehicle k's information map. We will denote S_i^j as vehicle j's current estimate of vehicle i's plan. The information function can be in the form of information gain, variance reduction, or discounted probability of detection (see Section IV). We assume that $J(S)$ can be evaluated for any number of vehicle paths.

Algorithm 1 Distributed coordination with limited communication

Input: vehicles 1 to K, mobility graph Ψ, objective J, planning horizon T, initial distribution b_0
$t \leftarrow 0$, $t_0 \leftarrow 0$, $b_k(t) \leftarrow b_0$ for all $k \in [K]$
while in parallel on each vehicle $k \in [K]$ **do**
  % Process any transmissions waiting in the queue
  while queue is not empty **do**
    Process received transmission from vehicle i
    Update $b_k(t)$ by fusing $b_i(t)$
    Update path estimate S_i^k for vehicle i
  end while
  if at replanning location **then**
    $t_0 \leftarrow t$
    % Sample informative paths
    $\xi = \mathcal{Q}(\Psi_k(t_0, t_0 + T))$
    % Determine best path given current information
    $S_k^* \leftarrow \operatorname{argmax}_{S_k \in \xi} J(b_k(t), S_k, S_1^k, \ldots)$
  end if
  Broadcast $b_k(t)$ and $S_k^*(t, t_0 + T)$
  Continue execution of $S_k^*(t, t_0 + T)$
  $t \leftarrow t + 1$
end while

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. Each vehicle incorporates any plans and estimates received from teammates, and this new information is fused into each vehicle's current estimate (see Section IV) to improve future plans. 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.

IV. DISTRIBUTED DATA FUSION

To enable efficient information sharing, we now present a data fusion technique for the moving target search problem. 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 [11] 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 the history of observations $\mathbf{Z}^t$ and an initial state estimate $\mathbf{x}_0$. Using the recursive Bayes' rule, calculation of the posterior estimate takes the following form:

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

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 [11]. 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 a conservative estimate of fused distribution. A conservative estimate always overcorrects for the redundant information $\Pr(\mathbf{x}_t|\mathbf{Z}_{i\cap j}^t)$ (possibly undercounting it). For the case of Gaussian distributions, it is possible to achieve a conservative estimate using a weighted combination of the disparate estimates and covariances [12]. However, particularly in target search applications, the distribution of interest cannot be modeled using these assumptions. A method is presented below that provides a conservative fusion method for a number of interesting application using little computation.

A. Data Fusion for Moving Target Search

We now present a data fusion rule for locating a moving target. The goal is to move the probability of the target's location from the environment into a "capture state." The formulation also applies to locating features of interest, such as an area of scientific interest at an unknown location. A common objective function used for moving target search is the discounted probability of locating the target[1] or feature of interest [7], [6]:

$$J(U(1), \ldots, U(T)) = \sum_{t=0}^{T} \gamma^t \Pr(\exists i : s_i(t) = e(t)), \quad (2)$$

where $U(1), \ldots, U(T)$ are the vehicles' planned paths, γ is a discount factor, $s_i(t)$ is the location of vehicle i at time t, and $e(t)$ is the location of the target at time t. $\Pr(\exists i : s_i(t) = e(t))$ can be viewed as the amount of probability in a capture state (i.e., the probability that the target will already have been found at time t). Note that this objective requires that a model of the target's behavior is known (or approximately known) to evaluate the probability term. This model is also assumed to be independent of the positions of the searching vehicles (i.e., the target is not modeled as actively evading or cooperating).

Given the objective function above, we can formulate a data fusion problem for moving target search. The target's state at time t is represented as a belief vector $b(t) = [b_0(t), b_1(t), \ldots, b_N(t)]$, where $b_0(t)$ is the probability that the target has already been found at time t, and $b_1(t)$ through $b_N(t)$ are the probabilities that the target is in cells 1 through N respectively. We will refer to the state with belief $b_0(t)$ as the capture state. The belief vector $b(t)$ describes probability that the target is in each possible state of $\mathbf{x}_t$ in the decentralized data fusion framework.

A "capture" event, or observation that may locate the target, is represented by the application of a matrix C to the belief vector, where C moves some probability from the cells 1 through N to the capture state. A capture event is an observation of the target's state and corresponds to a measurement $\mathbf{z}_t$. The target's movement is modeled using a matrix D that disperses the probability based on a motion model but does not move any probability to the capture state. Subsequent capture and dispersion events are modeled by multiplying the target's belief vector by additional matrices (see reference [7] for a more detailed explanation of capture and dispersion matrices). The dispersion matrices apply the motion model $\Pr(\mathbf{x}_t|\mathbf{x}_{t-1})$.

If vehicles cannot communicate and later become reconnected (e.g., they loose connectivity for a while), they will need to share the information they have gained. In such cases, it is necessary to develop a fusion rule that merges the vehicles' distributions without sharing and incorporating the entire history.

B. Minimum Fusion Rule

Before introducing a fusion rule, we will require some additional notation. When vehicles become disconnected, it is assumed that each of them maintains is own belief vector based on the application of its own capture matrices and the capture matrices of other vehicles with which it is 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, $b^k(t) = b^*(t)$ for all k. The following rule is introduced to merge the estimate of vehicle i with any number of vehicle estimates into a single merged estimate $b^m(t)$:

$$b_n^m(t) = \min_j b_n^j(t), \quad (3)$$

[1]Note 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 [7].

where $\min_j$ is chosen over all vehicles within communication range of vehicle i. Equation (3) is applied over all $n \in \{1, \ldots, N\}$. After applying the rule in (3), 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 reduced. This approach is undesirable since the target will never escape capture. 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:

$$b_0^m(t) = 1 - \sum_{n=1}^{N} b_n^m(t) \tag{4}$$

If multi-hop communication is allowed, the resulting distribution $b^m(t)$ is shared by all vehicles connected to vehicle i. 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 connected vehicles will eventually share the same distribution.

The intuition behind this fusion rule is that each observation will reduce the amount of probability in the environment, which necessarily increases the probability in the capture state. Thus, the vehicles will want to regain that information through the fusion rule if the observation is missed. The next section will show that the merged distribution has several desirable properties relative to the true distribution.

C. 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 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.

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 $\min_k b_n^k(t)$ is greater than or equal to $b_n^*(t)$ for any cell $n \in \{1, \ldots, N\}$, any time t, and any number of vehicles included in $\min_k$.

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). W.l.o.g. let n be an arbitrary cell in the environment, and let $b_n(t)$ be the belief at time t for cell n. The proof will be by induction on $b_n(t)$. By assumption, $b_n^k(0) = b_n^*(0)$ for all vehicles k and cells n. We now show that $b_n^k(1) \geq b_n^*(1)$.

Let α_{ij} be the probability that a target moves from cell i to cell j, which is encoded in the dispersion matrix D. Note that $\sum_i \alpha_{ij} = 1$, for all j. Let β_i^j be the probability that a target is captured in cell i given that capture matrix C_j is applied.

The application of a capture matrix C_j can be rewritten for each cell n as

$$b_n(t+1) = (1 - \beta_n^j) b_n(t) = b_n(t) - \beta_n^j b_n(t). \tag{5}$$

The dispersion rule can be rewritten for each cell n as

$$b_n(t+1) = \alpha_{n1} b_1(t) + \ldots + \alpha_{nN} b_N(t). \tag{6}$$

By assumption, each vehicle applies the same dispersion matrix D to the belief vector. After the application of the dispersion matrices $b_n^k(0^+) = b_n^*(0^+)$. If communication is perfect, all capture matrices will be applied to $b^k(0^+)$ that are applied to $b^*(0^+)$, which leads to $b^k(1) = b^*(1)$.

Due to imperfect communication, one or more capture matrices may not be applied. W.l.o.g. assume that capture matrix C_j is the first capture matrix applied to $b^*(0^+)$ that is not applied to $b^k(0^+)$. Now, $b_n^k(1^-) = b_n^*(1^-) + \beta_n^j b_n^*(1^-) \geq b_n^*(1^-)$. For each subsequent capture matrix applied or not applied, this inequality continues to hold. Thus, in all cases $b_n^k(1) \geq b_n^*(1)$.

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

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

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

Proof: Immediate from the fact that $b_0^*(t) = 1 - \sum_{n=1}^{N} b_n^*(t)$ and Theorem 1. ∎

The analysis above shows that the vehicles will never become "more sure" of the target's 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, the resulting strategy will be conservative. Some areas may be searched more often than they would be without the merging, but no area will be neglected due to the merging.

In addition, the minimum fusion rule can be utilized when a partial map is received from another vehicle. As long 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.

V. UNDERWATER ACOUSTIC COMMUNICATION

In this section, we review key features of underwater acoustic communications and develop a model for the performance of an instantiation of an acoustic modem. The communication models will be used in simulation to determine the probability of communication between two locations. The development of such models allows us to validate the proposed techniques given realistic limitations on communication. In addition, we can analyze the probable effect of environmental factors on the communication characteristics of the multi-vehicle system.

High speed communications over underwater acoustic channels (UAC) have been considered challenging due to range dependent bandwidths, large multipath, time-varying channels, Doppler distortion, and colored ambient noise. The extent of these effects typically depend on several factors, such as local environmental conditions, the movement of transmitters and receivers, and beam patterns of the acoustic transducers. There are fundamental differences that distinguish UAC's from terrestrial radio frequency (RF) channels [13]. The key differences are summarized below.

- Distance and frequency dependent signal attenuation due to spreading and absorption loss resulting in low operating frequencies and range dependent transmission bandwidth.
- Multipath due to low speed of propagation of sound in water, often causing interference of the order of tens of symbols.
- Doppler due to mobility manifests as a signal compression/dilation versus frequency shifts and time-varying multipath in wireless channels.
- Environmental factors such as winds and shipping contribute to noise.

As noted above, UAC's 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 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 [13] as

$$A(d,f) = A_0 d^k a(f)^d, \tag{7}$$

where A_0 is a normalizing constant. For frequencies above a few hundred Hz, the frequency dependent attenuation term is given (in dB/km) as

$$\begin{aligned} 10\log a(f) = {} & 0.11\frac{f^2}{1+f^2} + 44\frac{f^2}{4100+f^2} + \\ & 2.75\cdot 10^{-4} f^2 + 0.003. \end{aligned} \tag{8}$$

The spreading factor k is typically between $k = 1$ for cylindrical spreading in shallow water channels and $k = 2$ for spherical spreading.

Noise in UACs 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). The noise can be modeled [13] as

$$\begin{aligned} 10\log N_t(f) &= 17 - 30\log f \\ 10\log N_s(f) &= 40 + 20(s-0.5) + 26\log f - \\ &\quad 60\log(f+0.03) \\ 10\log N_w(f) &= 50 + 7.5w^{1/2} + 20\log f - \\ &\quad 40\log(f+0.4) \\ 10\log N_{th}(f) &= -15 + 20\log f \end{aligned} \tag{9}$$

where $s \in [0,1]$ models the surface shipping activity and w is the wind speed in m/sec. The effective noise level at a frequency f is then the sum of contributions due to each of these factors:

$$N(f) = N_t(f) + N_s(f) + N_w(f) + N_{th}(f). \tag{10}$$

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

$$SNR(d,f) = \frac{P/A(d,f)}{N(f)B}. \tag{11}$$

A. Signal Model and Probability of Error

We consider a quasi static channel model in which the channel parameters are considered constant over the duration of one packet. To combat inter-symbol interference due to large multipath, we consider Orthogonal Frequency Division Multiplexed (OFDM) signaling. OFDM is emerging as an attractive signaling scheme for UAC's as it can decompose a static frequency selective channel (multipath channels in time are selective in frequency) into a set of independent and interference free frequency domain sub-channels which lead to very low complexity receivers [14]. In this paper, we assume that the Doppler distortion after signal resampling to compensate for the time scaling is negligible. The output x_k of the k^{th} sub-channel for an OFDM symbol can then be modeled as

$$x_k = H_k s_k + n_k, \tag{12}$$

where s_k, H_k and n_k are the unit energy transmitted symbol, channel gain (fading coefficient) and additive noise for the k^{th} sub-channel. The noise is modeled as a complex Gaussian with $n_k \sim \mathcal{N}(0, N(f_0)B)$, where f_0 is the center frequency. We assume that the magnitude of the channel gain is Rayleigh distributed. Using the propagation model previously described, the average SNR can be computed as

$$SNR = \frac{E[|H_k s_k|^2]}{E[|n_k|^2]} = \frac{P/A(d,f_0)}{N(f_0)B}. \tag{13}$$

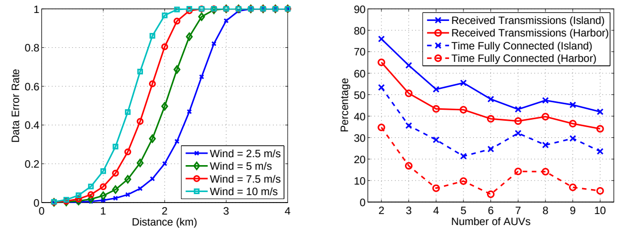

Fig. 2. Left: modeled data error rates for varying wind speeds. The curve gives the probability that an entire probability distribution of the target's estimated position will be transmitted without error. Right: percentage of successful transmissions and percentage of time that the AUVs were fully connected when running the proposed solution.

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

$$P(e) \approx \frac{1}{4SNR}. \tag{14}$$

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 un-coded packet with symbols over N sub-channels, the probability of packet error can then be computed as

$$P_{packet} = 1 - (1 - P(e))^N. \tag{15}$$

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 i to vehicle j), given the distance between the vehicles and an estimate of the wind speed and shipping activity. Figure 2 shows the data error rates with the system parameters in Section VI and varying wind speeds. In most application domains, wind speed can be estimated but not known exactly. It is immediately obvious from these models that a simple step function is insufficient to model the characteristics of the UAC. This motivates the design of algorithms capable of operating at any level of communication without continual connectivity.

VI. SIMULATED RESULTS

A multi-robot simulation environment was implemented in C++ running on Ubuntu Linux to test our proposed coordination and data fusion algorithms. The experiments were run on a 3.2 GHz Intel i7 processor with 9 GB of RAM. The simulated underwater vehicles move at 5 km/hr and have a detection radius of 200 m (motivated by the swath width for a side scan sonar). The communication system specifications are given in Table I. In these simulations, the vehicles broadcast their plans and estimated target distribution when arriving at a replanning point. The probability that the broadcast is received by each robot is determined as described in Section V based on its distance, line-of-sight, and the wind and shipping activity. Wind and shipping activity were set to vary randomly throughout the map.

TABLE I

PARAMETERS FOR UNDERWATER ACOUSTIC SIMULATIONS

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

The proposed coordination and data fusion techniques were validated using a search problem at a fixed depth. A non-adversarial target exists in the environment, and the underwater vehicles must search for its location. The proposed framework allows for both moving and stationary targets. In these simulations, the target is assumed to move randomly. The target's depth is either known a priori or it exists 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 3 shows surface maps of island and harbor environments of varying sizes used for simulated testing. The land masses serve as obstacles that prevent both communication and movement. It was assumed that any target that leaves the map is considered lost (i.e., it can never be captured). Lost targets are not considered in the cost function.

The traversable ocean portions of the maps were discretized into 200 m $\times$ 200 m regular grid cells, and the distributed planning algorithm from Section III was run with perfect communication and then with the communication system model from Section V. In these trials, the vehicles use a planning horizon $T = 4$ and an exhaustive enumeration to that horizon as the path sampling function $\mathcal{Q}$. In addition, a one-step lookahead algorithm was 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 robot is disconnected from the rest of the team. Based on a worst-case assumption of wind speed and shipping activity, the threshold that would guarantee connectivity was set to 2 km. Chain topologies are allowed, since each robot will get the information map from all other robots through the propagation of the minimum fusion rule.

Figure 3 shows the quantitative results from the simulated experiments. 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 proposed distributed coordination and data fusion techniques allow for this behavior, which leads to improved performance.

The percentage of successful transmissions and percentage

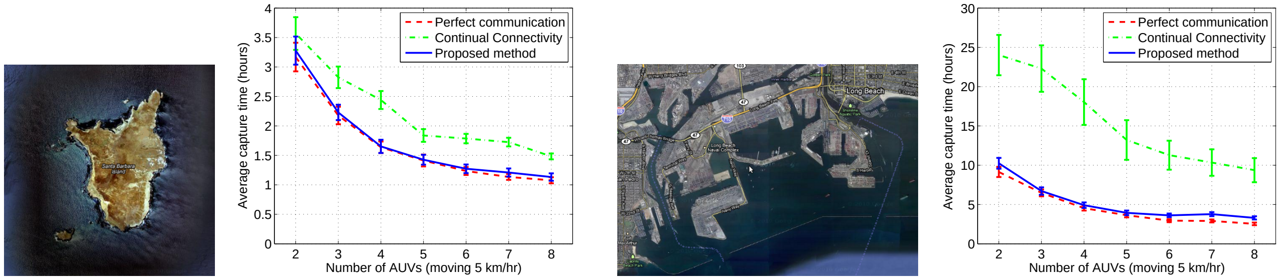

(a) Santa Barbara Island (b) Santa Barbara island simulations (c) Long Beach Harbor (d) Long Beach harbor simulations

Fig. 3. Simulated results for non-adversarial search in underwater environments with acoustic communications. Santa Barbara Island (left) is 4 km × 4 km and Long Beach harbor (right) is 11 km × 8 km. Each data point is averaged over 200 simulations; error bars are one standard error of the mean. The proposed method takes advantage of information at 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.

of time that the team is fully connected (including chain topologies) is shown in Figure 2. 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 robots break away from the team. In addition, particularly in the harbor environment, the AUV team is often not fully connected (less than 10% of the time in some cases). These results demonstrate that the proposed algorithm yields high performance even with poor communication quality.

VII. CONCLUSION

In this paper, we presented a distributed coordination approach for multi-robot planning with limited communication in the underwater search domain. Our proposed coordination algorithm utilizes available information to provide solutions robust to changes in communication. In addition, we proposed and analyzed data fusion techniques that provide principled methods for incorporating information from robots that regain connectivity. Furthermore, the development of realistic models for acoustic modems shows that the simplifying assumptions made in prior work are insufficient for handling the noise characteristics of underwater communication. Simulated experiments with realistic acoustic communication models demonstrate that it is possible to achieve low average capture times with little communication 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 with a large adaptivity gap (i.e., 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.

VIII. ACKNOWLEDGMENTS

The authors gratefully acknowledge Joseph Djugash and Benjamin Grocholsky from CMU and Athanasios Kehagias from Aristotle Univ. of Thessaloniki for insightful comments.

REFERENCES

[1] M. Zavlanos and G. Pappas, "Distributed connectivity control of mobile networks," *IEEE Trans. Robotics*, vol. 24, no. 6, pp. 1416–1428, 2008.

[2] F. Arrichiello, J. Das, H. Heidarsson, A. Pereira, S. Chiaverini, and G. Sukhatme, "Multi-robot collaboraiton with range-limited communication: Experiments with two underactuated ASVs," in *Proc. Int. Conf. Field and Service Robotics*, 2009.

[3] D. Anisi, P. Ögren, and X. Hu, "Cooperative surveillance missions with multiple UGVs," in *Proc. IEEE Conf. Decision and Control*, 2008.

[4] G. Hollinger and S. Singh, "Multi-robot coordination with periodic connectivity," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2010.

[5] 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.

[6] S. Ong, S. Png, D. Hsu, and W. Lee, "Planning under uncertainty for robotic tasks with mixed observability," *Int. J. Robotics Research*, 2010, to appear.

[7] G. Hollinger, S. Singh, J. Djugash, and A. Kehagias, "Efficient multi-robot search for a moving target," *Int. J. Robotics Research*, vol. 28, no. 2, pp. 201–219, 2009.

[8] N. Michael, M. M. Zavlanos, V. Kumar, and G. J. Pappas, "Maintaining connectivity in mobile robot networks," in *Proc. Int. Symp. Experimental Robotics*, 2008.

[9] 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.

[10] M. A. 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.

[11] A. Makarenko and H. Durrant-Whyte, "Decentralized data fusion and control in active sensor networks," in *Proc. Int. Conf. Information Fusion*, 2004.

[12] 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.

[13] M. Stojanovic, "On the relationship between capacity and distance in an underwater acoustic communication channel," in *Proc. ACM Int. Workshop on Underwater Networks*, 2006.

[14] B. Li, S. Zhou, M. Stojanovic, L. Freitag, and P. Willet, "Multicarrier communications over underwater acoustic channels with nonuniform doppler shifts," *IEEE J. Oceanic Engineering*, vol. 33, no. 2, pp. 198–209, 2008.