Cooperative Localization by Factor Composition over a Faulty Low-Bandwidth Communication Channel

Jeffrey M. Walls, Alexander G. Cunningham, and Ryan M. Eustice

***Abstract*— This paper reports on an underwater cooperative localization algorithm for faulty low-bandwidth communication channels based on a factor graph estimation framework. Vehicles measure the one-way-travel-time (OWTT) of acoustic broadcasts to obtain a relative range observation to the transmitting vehicle. We present a method to robustly share locally observed sensor data across the network by exploiting odometry factor composition. Our algorithm calls on approximate marginalization techniques to compute a compact set of informative factors that enable local navigation data to be shared efficiently. We provide results from a real-time implementation of our algorithm using two autonomous underwater vehicles and a surface vehicle.**

I. Introduction

Autonomous underwater vehicles (AUVs) typically integrate body-frame velocities, attitude, and pressure depth to compute a dead-reckoned (DR) navigation solution. Errors in xy horizontal position estimates grow unbounded in time without regular access to an absolute position reference (global positioning system (GPS) is only available at the surface). Bounded error navigation can be achieved with the aid of fixed acoustic beacon systems such as long-baseline (LBL). While these systems can accurately localize an AUV, they do not scale well to large vehicle networks and can be expensive to deploy.

Synchronous-clock acoustic hardware allows communicating vehicles to observe their relative range via the one-way-travel-time (OWTT) of acoustic broadcasts [1]. OWTTs provide a relative range measurement between the transmitting vehicle pose at the time-of-launch (TOL) and the receiving vehicle pose at the time-of-arrival (TOA). The acoustic channel, however, is broadcast and unacknowledged, provides very low bandwidth (typically less than 100 bps), and displays low reception rates (often less than 50%).

We propose a cooperative localization framework in which underwater vehicles act as mobile navigation beacons. We extend ideas originally presented in the origin state method proposed by Walls and Eustice [2] to a novel approach that exploits properties of the constituent factors within a graph-based estimator.

We employ approximate marginalization techniques to enable vehicles to share an informative subset of composed factors. Approximate marginalization has recently become a popular tool within the simultaneous localization and mapping (SLAM) community for reducing graph complexity, e.g., [3–6]. In this work, we exploit similar tools in order to efficiently distribute locally obtained information within a cooperative localization framework.

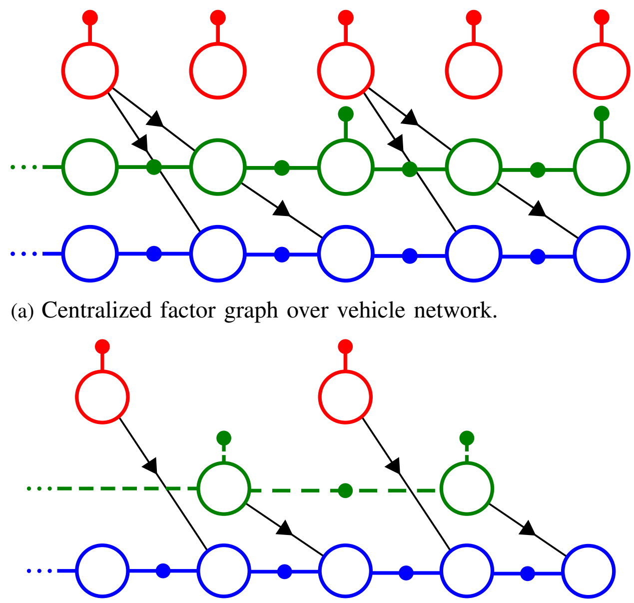

(a) Centralized factor graph over vehicle network.

(b) Factor graph over vehicle network constructed onboard the blue vehicle by our algorithm. Take note that the dashed green factors represent a new set of approximate factors distinct from the original in (a).

Fig. 1: Example cooperative localization factor graph—empty circles represent variable pose nodes, solid dots are odometry and prior factors, and arrows illustrate range-only factors and the direction of communication (TOA–TOL). In this example, red represents a topside ship with access only to GPS, while green and blue represent AUVs.

The specific contributions of this work include:

- We present a cooperative localization algorithm for communication-limited networks that is completely passive (i.e., does not rely on acknowledgments) and employs a *fixed* bandwidth data packet.
- We approximate the true locally constructed factor graph using approximate marginalization techniques to produce a structure that can be more easily broadcast.
- We provide a comparative evaluation of our algorithm and a performance summary from a three vehicle field deployment including two AUVs and a topside ship.

II. Related Work

Within cooperative localization frameworks, teams of communicating vehicles localize each other by sharing local

*This work was supported in part by the Office of Naval Research under award N00014-12-1-0092.

J. Walls, A. Cunningham, and R. Eustice are with the University of Michigan, Ann Arbor, Michigan 48109, USA
{jmwalls, alexgc, eustice}@umich.edu.

sensor data and observing relative vehicle pose. Early work explored sharing information among the team and performing local data fusion [7–11].

More recent methods proposed framing cooperative localization as a multiple vehicle graph-based SLAM problem [12, 13]. These methods implicitly handle correlation that develops between vehicle estimates when relative information is shared; our proposed method falls within this category. The centralized estimator for OWTT cooperative localization shares a similar structure to a pose-graph SLAM formulation with known data association.

Leung et al. [14] presented a decentralized cooperative localization algorithm capable of reproducing the centralized estimate in the presence of dropped communication. The bandwidth required for transmitting knowledge sets, however, may exceed the acoustic channel capacity during periods of lost connectivity. Nerurkar and Roumeliotis [15] similarly targeted cooperative localization in bandwidth constrained scenarios, but require constant connectivity.

Many other researchers in the underwater domain have considered fusing relative range measurements derived from observing the time-of-flight (TOF) of acoustic broadcasts. Previous work [1, 2, 16–23] has considered the limitations of the acoustic channel and the difficulties of fusing range-only observations.

The complexity of graph-based estimation is largely dependent on the size (i.e., number of variable nodes) and sparsity of the graph (i.e., number of edges). Graph sparsification methods have been introduced [3–6] to both reduce the number of variables in the graph and increase the sparsity without greatly affecting the solution. We take advantage of approximate marginalization techniques to refactor a local factor graph over a smaller set of poses to reduce the communication requirement.

The work proposed here is closest to that outlined by Walls and Eustice [2] and Fallon et al. [20]. In our previous work [2], we exploited the structure of the information matrix to cope with unpredictable packet loss. Unfortunately, the algorithm requires periodically shifting an origin pose, requires a recovery mechanism for pathological communication failures, and does not seamlessly allow for a vehicle with no odometry (e.g., a surface vehicle with only GPS) to participate. Fallon et al. [20] proposed constructing a graph-based estimator on each vehicle by broadcasting individual factors to the team. Their method transmits each factor over a large set of pose nodes and requires data requests when communication is lost. We exploit the factor structure of the local graph to robustly communicate information within a completely passive framework.

Paull et al. [23] independently arrived at a similar solution discussed in Section V and Section VI.

III. Problem Statement

We formulate cooperative localization within a factor graph framework [24]. The underlying structure of the factor graph consists of information local to each vehicle and information due to relative vehicle measurements. We will first review the factor graph formulation for a single vehicle and then expand to the full vehicle network.

In the single vehicle setting, the factor graph approach is a smoothing algorithm that estimates the entire trajectory of the vehicle. A factor graph is a bipartite graph with pose (variable) nodes and factor (measurement) nodes representing the joint distribution over the unknown poses. The ith vehicle graph represents the joint distribution over its N poses, $\mathbf{X}_i = [\mathbf{x}_1, \ldots, \mathbf{x}_N]$, as

$$p(\mathbf{X}_i) \propto p(\mathbf{x}_1) \prod_i p(\mathbf{z}_{\text{odo}_i} | \mathbf{x}_i, \mathbf{x}_{i-1}) \prod_j p(\mathbf{z}_{\text{prior}_j} | \mathbf{x}_j), \quad (1)$$

where we assume each vehicle has access to its initial belief $p(\mathbf{x}_1)$. The graph structure is a chain as we only consider unary 'prior' factors, $\mathbf{z}_{\text{prior}}$, (e.g., GPS) and pairwise sequential 'odometry' factors, $\mathbf{z}_{\text{odo}}$, (e.g., integrated velocity).

For convenience, we define a 'link', $\mathcal{L}_i$, associated with the ith pose node, $\mathbf{x}_i$, as a 2-tuple containing the odometry factor to the previous pose node and a prior factor. Note that each link need not have both an odometry and prior factor, for example the initial link only contains a prior factor. The local chain is the set of links which represent the vehicle trajectory corresponding to (1), $\mathcal{C}_{\text{local}} = \{\mathcal{L}_i\}_{i=1}^N$. In Fig. 1a, each uniformly colored subgraph represents a local chain.

We can construct the factor graph over the entire M vehicle network (i.e., all vehicle poses), $\{\mathbf{X}_1, \ldots, \mathbf{X}_M\}$,

$$p(\mathbf{X}_1, \ldots, \mathbf{X}_M) \propto \prod_i \underbrace{p(\mathbf{X}_i)}_{\mathcal{C}_{\text{local}_i}} \prod_k \underbrace{p(\mathbf{z}_k | \mathbf{x}_{i_k}, \mathbf{x}_{j_k})}_{\text{relative factors}}, \quad (2)$$

where each $\mathbf{z}_k$ represents a relative vehicle constraint between poses on vehicles i_k and j_k. In this work, $\mathbf{z}_k$ is a OWTT range constraint between a transmitting vehicle's TOL pose and a receiving vehicle's TOA pose. The factor graph for a three vehicle network is illustrated in Fig. 1a.

The gold-standard would be to compute the maximum a posteriori (MAP) estimate for each vehicle in a centralized estimator as

$$\begin{aligned} \mathbf{X}_i^* &= \arg\max_{\{\mathbf{X}_1,\ldots,\mathbf{X}_M\}} p(\mathbf{X}_1, \ldots, \mathbf{X}_M) \\ &= \arg\min_{\{\mathbf{X}_1,\ldots,\mathbf{X}_M\}} -\log p(\mathbf{X}_1, \ldots, \mathbf{X}_M) \end{aligned}, \quad (3)$$

which results in a nonlinear least-squares problem for Gaussian noise models [24].

The full joint distribution (2) consists of a product of each vehicle's local chain and the relative vehicle factors. Therefore, in order to construct (and perform inference on) the full factor graph, the ith vehicle must have access to the set of local factors from all other vehicles, $\{\mathcal{C}_{\text{local}_j}\}_{j \neq i}$, and the set of all relative vehicle factors. Sharing this information, however, is nontrivial due to the limitations of the acoustic communication channel.

In this paper, we present an algorithm for each vehicle to robustly communicate an *approximation* of its local chain. In turn, each vehicle in the network can construct a graph consisting of its local chain, the set of approximate chains received from other vehicles, and the set of relative vehicle measurements that it observes locally, i.e., range constraints for which it has measured the OWTT (a subset of all ranges)

Algorithm 1 Local chain distribution.

initialize($\mathcal{C}_{\text{local}}, \mathcal{C}_{\text{approx}}$) {Prior factor on initial pose.}
while is_running() **do**
  add_new_link($\mathcal{C}_{\text{local}}$) {New odometry and prior factors.}
  if t is TOL **then**
    append_approx($\mathcal{C}_{\text{local}}, \mathcal{C}_{\text{approx}}$) {Section IV-B.}
    $\mathbf{L}_K$ = choose_links($\mathcal{C}_{\text{approx}}$) {Section IV-C.}
    $\mathbf{L}'_K$ = compose_links($\mathbf{L}_K$) {Section IV-A.}
    broadcast_links($\mathbf{L}'_K$)
  end if
end while

Algorithm 2 Cooperative localization

Require: $\mathcal{C}_{\text{local}}$ {Local chain}
$\mathbf{C} = \emptyset$ {Set of received chains.}
$\mathbf{Z}_r = \emptyset$ {Set of received range observations.}
while is_running() **do**
  if $\mathbf{L}_{\text{rec}}, \mathbf{z}_r, i$ = received_broadcast() **then**
    add_links($\mathbf{C}[i], \mathbf{L}_{\text{rec}}$) {Links from vehicle i.}
    $\mathbf{Z}_r = \mathbf{Z}_r \cup \mathbf{z}_r$
    solve_batch($\mathcal{C}_{\text{local}}, \mathbf{C}, \mathbf{Z}_r$)
  end if
end while

as in Fig. 1b. We will show that our approximation produces an accurate representation that allows robust communication over a faulty and bandwidth-limited channel, and that using a subset of relative vehicle measurements still provides significant improvement over DR navigation.

IV. Cooperative Localization

The key requirement for cooperative localization is the ability to share local chains. Below, we develop an algorithm for robustly broadcasting a local chain across a faulty low-bandwidth communication channel. Our approach is completely passive, in other words, the information each vehicle broadcasts is independent of the rest of the network—a desirable property for an unacknowledged channel. In this section, we outline an effective strategy to *(i)* share a local chain (Algorithm 1) and *(ii)* use the set of received local chains and observed relative range observations to compute a navigation estimate (Algorithm 2), as illustrated in Fig. 1.

A local chain consists of odometry and prior factors. We first show that we can leverage the composition operation over odometry factors to represent each new odometry factor as a transformation relative to an 'origin' (§IV-A). This composition operation is invertible in an equivalent decomposition operation onboard the receiving vehicle. Moreover, both of these operations are robust to communication failure, so that every received odometry factor can be used to reconstruct the local chain.

We compose the odometry factor of each link, $\mathcal{L}_i$, to obtain the broadcast link, $\mathcal{L}'_i$. A receiving vehicle reconstructs the transmitter's chain consisting of decomposed odometry factors and prior factors contained in received links. If a broadcast is missed, upon the next successful reception the receiving vehicle can still reconstruct the chain *without* the prior factor in the missed link. We rebroadcast links with prior factors to help ensure an accurate chain reconstruction, meaning that at each TOL we broadcast a set of links, $\mathbf{L}$.

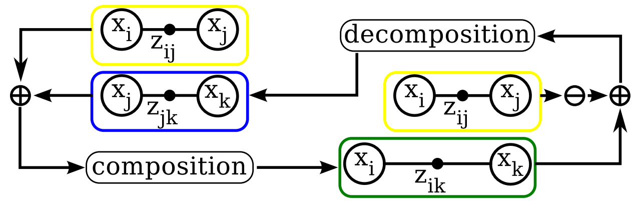

Fig. 2: A sequence of odometry factors can be composed into a single factor. Likewise, a composed factor can be decomposed given another factor.

Since we only use the set of locally observed range measurements, each vehicle need only share TOL pose nodes. However, a local chain, $\mathcal{C}_{\text{local}}$, consists of many additional pose nodes. To reduce the communication burden, we employ approximate marginalization to compute an approximate local chain $\mathcal{C}_{\text{approx}}$ that only contains odometry and prior factors over TOL pose nodes (§IV-B).

The ith receiving vehicle is able to reconstruct the broadcast local chains by decomposing broadcast odometry factors with odometry factors already received. Using the set of reconstructed approximate local chains $\{\mathcal{C}_{\text{approx}_j}\}_{j \neq i}$, its own local chain $\mathcal{C}_{\text{local}_i}$, and the set of observed OWTT range constraints, the vehicle can construct and solve a batch nonlinear least-squares problem (3) to estimate its smoothed trajectory.

A. Odometry Composition

In this section, we show that composition is reversible (decomposition). Moreover, it is equivalent to marginalization over the odometry chain. We leverage this property to robustly broadcast each link $\mathcal{L}_i \in \mathcal{C}_{\text{approx}}$.

Rigid-body transformation (e.g., odometry) observations are full rank relative-pose constraints expressed as

$$\mathbf{z}_{ij} = h_{ij}(\mathbf{x}_i, \mathbf{x}_j) + \mathbf{w}_{ij} \tag{4}$$
$$= \ominus \mathbf{x}_i \oplus \mathbf{x}_j + \mathbf{w}_{ij}, \tag{5}$$

where $\mathbf{x}_i$ and $\mathbf{x}_j$ are two poses with respect to the same reference frame, $\mathbf{w}_{ij} \sim \mathcal{N}(\mathbf{0}, \mathrm{R}_{ij})$ is an independent additive noise vector, and $\oplus$ and $\ominus$ represent generalized compounding and inverse operators. These operators are simple addition and subtraction for poses in $\mathbb{R}^d$ and are nonlinear functions for poses on $SE(2)$ and $SE(3)$ [25].

Per Smith et al. [25], a composite observation $\mathbf{z}_{ik}$ is computed through the sequence

$$\mathbf{z}_{ik} = g(\mathbf{z}_{ij}, \mathbf{z}_{jk}) \tag{6}$$
$$= \mathbf{z}_{ij} \oplus \mathbf{z}_{jk}. \tag{7}$$

The composite relation $\mathbf{z}_{ik}$ is a random variable. A first-order covariance approximation is computed as

$$\mathbf{z}_{ik} \approx g|_{\mathbf{z}_{ij}, \mathbf{z}_{jk}} + \mathrm{J}_{1\oplus}\delta_{ij} + \mathrm{J}_{2\oplus}\delta_{jk} \tag{8}$$
$$\sim \mathcal{N}\big(h_{ik}, \mathrm{J}_{1\oplus}\mathrm{R}_{ij}\mathrm{J}_{1\oplus}^{\top} + \mathrm{J}_{2\oplus}\mathrm{R}_{jk}\mathrm{J}_{2\oplus}^{\top}\big), \tag{9}$$

where $h_{ik} = h_{ij} \oplus h_{jk}$ and $\mathrm{J} = \partial g / \partial(\mathbf{z}_{ij}, \mathbf{z}_{jk}) = [\mathrm{J}_{1\oplus}, \mathrm{J}_{2\oplus}]$.

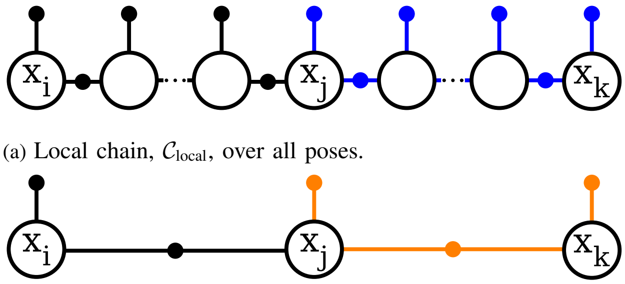

(a) Local chain, $\mathcal{C}_{\text{local}}$, over all poses.

(b) Approximate local chain, $\mathcal{C}_{\text{approx}}$, over TOL poses.

Fig. 3: At the new TOL, $\mathbf{x}_k$, we compute a set of factors (orange) that best represents the distribution induced by the original factors (blue) since the last TOL, $\mathbf{x}_j$, with intermediate nodes (shaded) marginalized out.

Suppose we are given the composed observations $(\mathbf{z}_{ij}, \mathrm{R}_{ij})$ and $(\mathbf{z}_{ik}, \mathrm{R}_{ik})$. It follows directly from (9) that we can invert this operation to recover the transformation between $\mathbf{x}_j$ and $\mathbf{x}_k$ and its covariance as

$$\mathbf{z}_{jk} = \ominus \mathbf{z}_{ij} \oplus \mathbf{z}_{ik} \tag{10}$$

$$\mathrm{R}_{jk} = \mathrm{J}_{2\oplus}^{-1}(\mathrm{R}_{ik} - \mathrm{J}_{1\oplus}\mathrm{R}_{ij}\mathrm{J}_{1\oplus}^{\top})\mathrm{J}_{2\oplus}^{-\top}, \tag{11}$$

where $\mathrm{J}_{2\oplus}$ is guaranteed to be nonsingular by the nature of the compounding operation. We call this operation, 'decomposition'. Composition and decomposition are illustrated in Fig. 2.

In our field trials, each pose is parameterized by the vehicle's xy horizontal position, i.e., $\mathbf{x}_i \in \mathbb{R}^2$. In this case, $\mathrm{J}_{1\oplus} = \mathrm{J}_{2\oplus} = \mathrm{I}$, and composition and decomposition reduce to simple addition and subtraction, respectively. Paull et al. [23] introduced an equivalent composition/decomposition operation applicable in only $\mathbb{R}^2$.

Any link in the local chain can robustly be communicated by broadcasting its composed odometry to an 'origin' pose (in practice, we use the local coordinate frame origin) as in line 7 of Algorithm 1. The receiving vehicle simply decomposes the odometry factor with the set of already received links as in line 5 of Algorithm 2. In this manner, we can exactly distribute the odometry backbone of a local graph. The subset of links each vehicle has received is unknown to the broadcasting vehicles, however, composition/decomposition still allows a receiver to reconstruct a complete chain over received links.

B. Local Chain Approximation

A local chain, $\mathcal{C}_{\text{local}}$, will include both unary prior factors and pairwise sequential odometry factors over the full set of pose nodes. Ideally, the local chain would only include pose nodes at each TOL—the only poses involved in relative ranging events observed by other vehicles—to reduce the size of the chain that must be shared. In this section, we detail a method to incrementally approximate the full local chain with a chain only over TOL pose nodes, $\mathcal{C}_{\text{approx}}$, as illustrated in Fig. 3.

We refer to the marginal distribution over TOL pose nodes computed from $\mathcal{C}_{\text{local}}$ as the target distribution. The marginal cannot be exactly represented by a set of prior and odometry factors (a requirement for robust communication by odometry composition, see §IV-A). However, we can compute the set of factors that most closely models this target distribution. The general framework for representing a Gaussian distribution using a desired nonlinear factor set is outlined by Mazuran et al. [6].

We incrementally construct the approximate local chain by 'refactoring' the full local chain between TOLs as in Fig. 3. The links contained between the last TOL and the current TOL induce a marginal distribution over the TOL nodes $p(\mathbf{x}_j, \mathbf{x}_k) = \mathcal{N}(\boldsymbol{\mu}, \Sigma) = \mathcal{N}^{-1}(\boldsymbol{\eta}, \Lambda)$.

We compute a single odometry factor, $\mathbf{z}_{jk} = h_{\text{odo}}(\mathbf{x}_j, \mathbf{x}_k) + \mathbf{w}_{jk}$, and a prior over each node, $\mathbf{z}_j = h_{\text{prior}}(\mathbf{x}_j) + \mathbf{w}_j$ and $\mathbf{z}_k = h_{\text{prior}}(\mathbf{x}_k) + \mathbf{w}_k$, that most closely induce the target distribution, $p(\mathbf{x}_j, \mathbf{x}_k)$. Let the distribution induced by the approximate factors be

$$q(\mathbf{x}_j, \mathbf{x}_k) = \mathcal{N}^{-1}(\boldsymbol{\eta}', \Lambda') \tag{12}$$

$$\boldsymbol{\eta}' = \mathrm{J}^{\top}\mathrm{R}^{-1}(\mathbf{Z} - h(\mathbf{x}_i, \mathbf{x}_j)) \tag{13}$$

$$\Lambda' = \mathrm{J}^{\top}\mathrm{R}^{-1}\mathrm{J}, \tag{14}$$

where $\mathbf{Z}$ represents the stacked observation vector, J is the stacked measurement Jacobian, and R is the block-diagonal stacked measurement covariance. We choose $\mathbf{Z}$ and R to minimize the Kullback-Leibler divergence (KLD) between $p(\mathbf{x}_j, \mathbf{x}_k)$ and $q(\mathbf{x}_j, \mathbf{x}_k)$. The KLD is minimized when $\mathbf{Z}$ is the expected observation evaluated at the mean, i.e., $h(\boldsymbol{\mu}_i, \boldsymbol{\mu}_j)$. R is then computed as

$$\mathrm{R}^* = \arg\min_{\mathrm{R}\in\mathrm{S}^{++}} D_{\mathrm{KL}}(p(\mathbf{x}_j, \mathbf{x}_k) || q(\mathbf{x}_j, \mathbf{x}_k)) \tag{15}$$

$$= \arg\min_{\mathrm{R}\in\mathrm{S}^{++}} \operatorname{tr}\left(\mathrm{J}^{\top}\mathrm{R}^{-1}\mathrm{J}\Sigma\right) - \ln\det\left(\mathrm{J}^{\top}\mathrm{R}\mathrm{J}\right), \tag{16}$$

where S^{++} is the cone of symmetric positive definite matrices (with appropriate sparsity structure). As shown in Mazuran et al. [6], the above optimization problem is convex and can be solved efficiently. While we have not enforced a consistency constraint on R, this could be expressed as a convex linear matrix inequality as in [26].

In general, Λ will not be full rank and we cannot compute approximate factors as in (16). This will only occur if no prior is available over the time window $i \to j$ (as is often true for AUVs). In this case, however, we can *exactly* represent the target information with a composed odometry factor (see §IV-A) and no prior factors. Similarly, if no odometry is available over $i \to j$ (e.g., for a topside vehicle with only GPS), then the TOL poses are independent and a prior may be computed to exactly represent the target distribution without an odometry factor.

At the TOL, after computing new approximate link factors in $\mathcal{C}_{\text{approx}}$, a vehicle broadcasts the new TOL link and k additional links to rebroadcast, $\mathbf{L}$, as in Algorithm 1. The size of k is predetermined depending on the available bandwidth.

C. Choosing Good Priors

$\mathcal{C}_{\text{approx}}$ can be reconstructed from the set of received links, however, not all broadcasts will be successfully received so that some prior factors will be missing. We use some additional bandwidth to rebroadcast useful or informative links with prior factors. Below, we consider finding the set of prior factors that are most informative about the current vehicle pose so that receiving vehicles are able to best use their observed range measurement.

A simple method is to broadcast the k most recent prior factors (k-last in §V). Intuitively, the most recent prior factors may be most informative about the transmitting vehicle's current pose. However, as a counterpoint, consider a vehicle with perfect (deterministic) odometry. Then, the best prior factors would be the factors with lowest uncertainty, not the most recent. For this reason, we develop a method to identify a more informative set of prior factors (k-best in §V).

Since we broadcast each prior factor over a lossy and unacknowledged channel, we assume that the ith transmission is lost with probability r_i. If a link has been broadcast m_i times, then the probability that it has been received is

$$p_{\text{rec}_i} = 1 - \prod_{i=1}^{m_i} r_i. \tag{17}$$

The results presented in §V assume that r_i is a fixed parameter; however, r_i could incorporate additional knowledge, or even be a learned parameter.

We use a mutual information objective for determining the utility of a prior factor. Mutual information indicates the uncertainty reduction in the current pose of the reconstructed chain, $\mathbf{x}_N$, by adding knowledge of a prior factor. Let $\mathbf{L}_K = \{\mathcal{L}_1, \ldots, \mathcal{L}_k\}$ be the set of k links that we will broadcast. The optimal set is then given by

$$\begin{aligned}\mathbf{L}_K^* &= \arg\max_{\mathbf{L}_K} \mathrm{E}[\mathrm{MI}[\mathbf{x}_N | \mathbf{L}_K]] \\ &= \arg\max_{\mathbf{L}_K} \log \left|\Lambda_{\text{recon}}(\mathbf{x}_N)\right| \\ \Lambda_{\text{recon}} &= \Lambda_{\text{odo}} + \sum_i p_{\text{rec}_i} \Lambda_{\text{prior}_i},\end{aligned}$$

where $\Lambda_{\text{recon}}(\mathbf{x}_N)$ represents the marginal information of $\mathbf{x}_N$ in the reconstructed chain computed from Λ_{recon} via the Schur complement. In general, this is a difficult combinatorial optimization problem. Instead, we follow a simpler approach. We evaluate the objective for each link and then greedily select the single link, $\mathcal{L}_j$, that maximizes the objective, increment m_j, and remove $\mathcal{L}_j$ from the set of potential links to broadcast. We repeat the process of evaluating each link and greedily selecting the best k times.

V. Field Trials

A. Implementation details

For validation, we fielded two Ocean-Server Inc. Iver2 AUVs, termed AUV-A and AUV-B (see Fig. 4), and a topside support ship. The AUVs are each equipped with an advanced dead-reckoning sensor suite including a $600\,\mathrm{kHz}$ RDI Doppler velocity log (DVL), a Microstrain 3DM-GX3-25 attitude heading reference system (AHRS), and a Desert Star Systems SSP-1 digital pressure sensor. Each AUV observed GPS during intermittent surface intervals. The topside support ship had access only to GPS (no measured odometry). All vehicles used the Woods Hole Oceanographic Institution (WHOI) Micro-modem and co-processor board with a synchronous-clock reference for inter vehicle communication and ranging.

Fig. 4: One of the Iver2 AUVs used in the field trials.

Each vehicle pose was parameterized by its xy horizontal position. GPS observations were transformed to a local coordinate frame and constitute a full-rank linear observation. We assumed that the GPS accuracy was constant and therefore we assumed noise with a fixed standard deviation of $5\,\mathrm{m}$. The AUV odometry, $\mathbf{z}_{ij}$, and corresponding covariance, R_{ij}, were computed by Euler integrating DVL and AHRS and performing a first-order covariance approximation [1].

Observed OWTT ranges represent a 3D *slant* range. Since depth is measured with bounded error, we are able to project ranges into the 2D horizontal plane and use the resulting *pseudo* ranges within our estimation framework. We used a fixed OWTT measurement uncertainty since the relative depth between each vehicle's acoustic transducer was small and relatively constant.

Vehicles broadcast navigation messages roughly once per minute according to a fixed time division multiple access (TDMA) schedule. The proposed composition/decomposition allowed a coarse quantization (factors were rounded to $1\,\mathrm{cm}$). Each navigation packet (i.e., set of links $\mathbf{L}$) required a single $64\,\mathrm{B}$ frame of either a Micro-modem Rate 1 or Rate 2 data packet. Previous work [2] required two $64\,\mathrm{B}$ frames, leading to a 50% reduction.

We employed our k-last prior factor selection strategy during real-time experiments (for $k = 2$). In post-process, we computed the prior factors that would have been broadcast with our k-best strategy.

We also implemented the algorithm proposed by Paull et al. [23] for comparison. Their algorithm similarly broadcasts each new TOL pose as a composed factor. There are two large differences in our approach: *(i)* they transmit prior factors directly from the full local chain separately, i.e., they must broadcast each prior link in addition to TOL links, and *(ii)*, they also broadcast received range factors and their corresponding TOA links. The second difference effectively allows for bidirectional ranging at the cost of packets that grow linearly in the size of the network (a TOA pose is broadcast for each vehicle in the network). During the three vehicle post-process evaluation, their algorithm broadcast more links requiring greater bandwidth (4 composed odometry factors + 2 range factors + a prior factor compared to 3 composed odometry factors + 2 prior factors for ours).

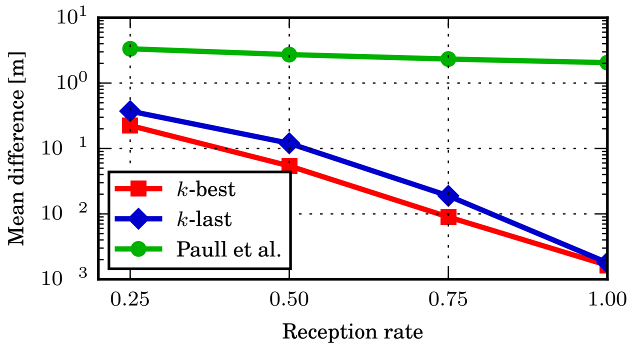

(a) Mean difference in pose between target and reconstructed.

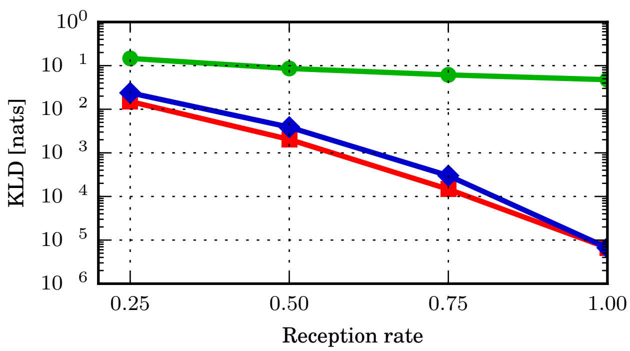

(b) KLD between target distribution and reconstructed.

Fig. 5: Local graph reconstruction quality for k-last, k-best, and Paull et al. for various reception rates averaged over 100 trials for six local chains. Note log scale on ordinate axis.

B. Local graph reconstruction

Here we demonstrate the quality of the local chain reconstruction under varying channel conditions. We ran three trials ranging from $60-90$ min to collect navigation observations and broadcast event times for each of our two AUVs resulting in six local chains. During the trials, each AUV had intermittent access to GPS during brief surface intervals. In post-process, we computed and sampled the set of navigation packets at different reception rates and reconstructed the local chain. For each reception rate, we sampled different sets of received messages 100 times. We then compared the reconstructed local chain to the target chain (the marginal full chain over received TOL pose nodes). Fig. 5 compares the local chain reconstruction for prior factor selection methods, k-last and k-best, as well as the method proposed by Paull et al. [23].

We can see that our proposed algorithm with the k-best selection strategy outperforms k-last by a factor of two or three in most trials. At 100% reception rate (no dropped packets) the full approximate chain was reconstructed and is therefore equivalent for k-best and k-last. Moreover, the full approximate chain closely represents the full local chain.

Our proposed algorithm improves upon Paull et al.'s method by several orders of magnitude in most trials. Our method is able to include more prior information into the broadcast links, and therefore leads to a better reconstruction.

As mentioned in §IV-A, composition is equivalent to marginalization. Therefore, when no priors are present in the local chain, all compared algorithms are exactly equal to the target distribution (up to quantization errors).

C. Cooperative localization

We executed three multiple vehicle trials to demonstrate the ability of our algorithm to provide useful navigation information to an AUV. We varied the number of vehicles and AUV access to GPS in each trial to demonstrate a variety of practical applications. Acoustic reception rates varied between 37–86% up to 500 m relative range. Results are summarized in Fig. 6. Fig. 6j reports each distributed algorithm's ability to produce the centralized uncertainty estimate onboard AUV-B. Although we only show results for AUV-B, note that all vehicles compute a local reconstruction of the centralized estimator.

1) Trial 1: AUV-A and AUV-B performed overlapping orthogonal lawnmower surveys (see Fig. 6a). Both vehicles had only dead-reckoned navigation available. Since there are no GPS priors, our local graph reconstruction is equivalent to Paull et al.'s [23]. Moreover, the local chain reconstruction for all algorithms was exact. k-last and k-best produced the same solution since there are no priors to rebroadcast. Since Paull et al. incorporated range constraints between both vehicles, their algorithm computed a solution that more closely matches the centralized result (see Fig. 6j). Their algorithm did not exactly reproduce the centralized result because AUV-B did not receive all range factors observed by AUV-A. Although we only used the local subset of range factors, we were still able to benefit from relative range observations and the difference compared to the centralized estimator is small.

2) Trial 2: AUV-A executed a narrow diamond trajectory over AUV-B's lawnmower survey. Both vehicles had dead-reckoned navigation available, but AUV-A also intermittently received GPS during short surface intervals (see Fig. 6b). Fig. 6e plots the pose uncertainty onboard AUV-B. Using our proposed algorithm, AUV-B was able to accurately approximate AUV-A's pose-graph, including the many prior factors. k-best here performed slightly better than k-last because of its ability to more accurately reproduce AUV-A's local chain. Although Paull et al. [23] was able to incorporate range information in both directions, our reconstruction of AUV-A was more accurate, leading to a more confident estimate. Due to the diminishing returns of information provided by the range observations, our algorithm was able to closely reproduce the centralized solution despite only using a subset of the range observations.

3) Trial 3: A topside vehicle with constant GPS access supported AUV-A (with intermittent GPS) and AUV-B (see Fig. 6c). The factor graph produced here is illustrated in Fig. 1. AUV-A followed a large diamond over AUV-B's lawnmower survey while the topside vehicle drifted above the survey area. In this case, the centralized estimator used ranges between all three vehicles. Our method only used ranges between the local platform and the other vehicles. Paull et al.'s, however, included ranges in both directions

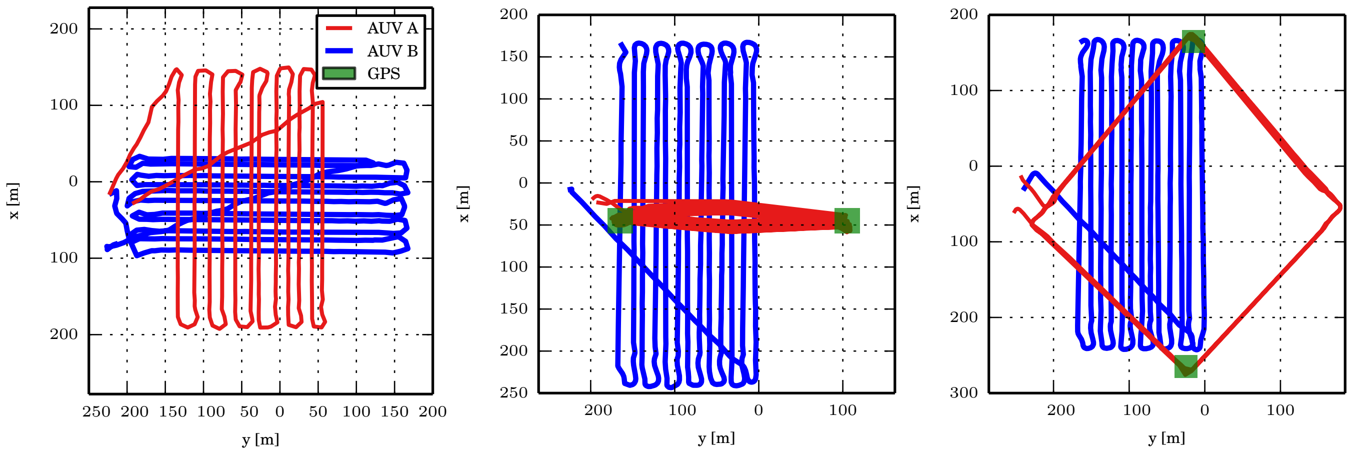

(a) Trial 1: relative paths (1.17 h), no access to GPS.

(b) Trial 2: relative paths (1.50 h), AUV-A has access to GPS.

(c) Trial 3: relative paths (1.55 h), AUV-A, topside (not shown) have GPS.

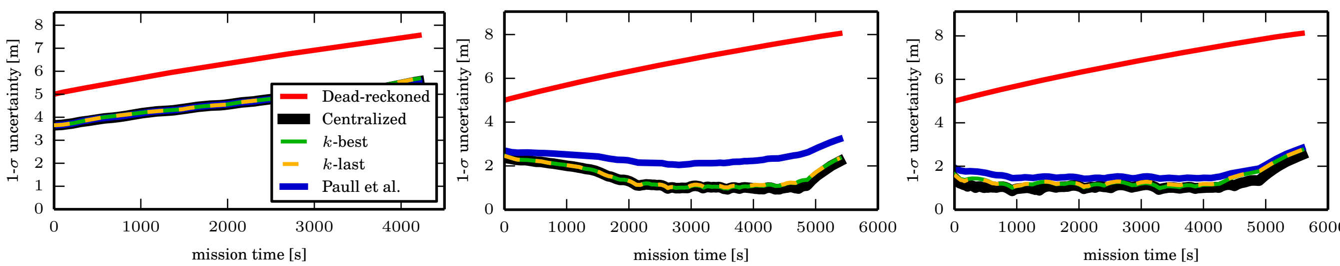

(d) Trial 1: AUV-B's uncertainty.

(e) Trial 2: AUV-B's uncertainty.

(f) Trial 3: AUV-B's uncertainty.

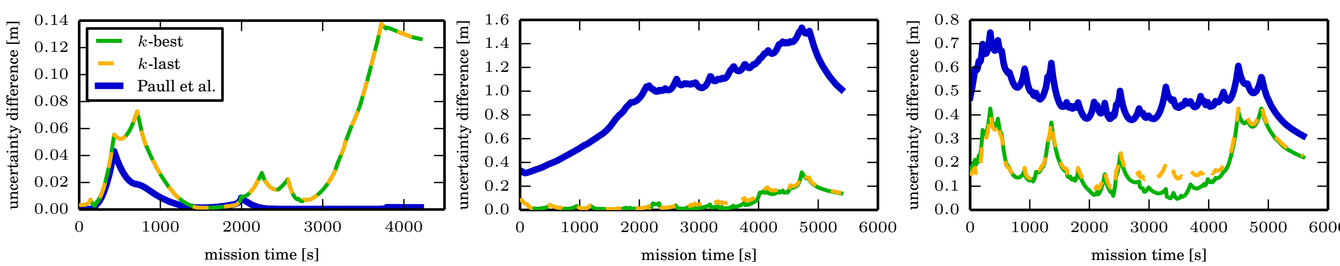

(g) Trial 1: Difference in uncertainty compared to centralized estimator.

(h) Trial 2: Difference in uncertainty compared to centralized estimator.

(i) Trial 3: Difference in uncertainty compared to centralized estimator.

	Trial 1 [m]	Trial 2 [m]	Trial 3 [m]
Paull et al. [23]	**0.005**	0.981	0.463
Proposed (k-last)	0.039	0.086	0.212
Proposed (k-best)	0.039	**0.074**	**0.192**

(j) Average 1-σ uncertainty difference compared to the centralized estimator.

Fig. 6: Summary of field trials and performance comparison. (a–c) shows each vehicle trajectory. (d–f) plots the smoothed uncertainty in each AUV-B pose computed as the fourth root of the determinant of the pose marginal covariance.

and could also use ranges between other vehicles, but only if the TOL and TOA poses in each local chain had been received. Once again, our reconstruction of AUV-A's chain is more informative. As shown in Fig. 6j, we have received the bulk of the benefit in terms of uncertainty reduction from the local set of range observations and achieve a smaller estimate uncertainty compared to Paull et al.

VI. Discussion

Our algorithm is similar in some respects to that independently proposed by Paull et al. [23], however, we believe there are many beneficial differences illustrated in Fig. 7. The most marked improvement in our algorithm is its ability to broadcast a local chain by using an accurate approximation that includes a more informative set of prior factors. The compact approximate chain contains fewer links to broadcast. While we currently only incorporate ranges measured locally, this allows for a fixed bandwidth data packet, regardless of the size of the vehicle network. We can use the additional bandwidth to rebroadcast previous informative links and for other practical purposes (e.g., command, control, vehicle health). Paull et al. [23] are able to incorporate inter-vehicle

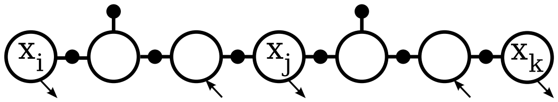

(a) True local chain, $\mathcal{C}_{\text{local}}$, over all poses.

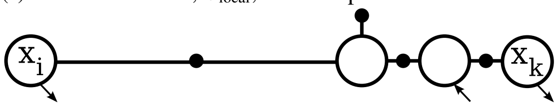

(b) Local chain reconstructed by Paull et al. including a range observed by the other vehicle at a TOA pose.

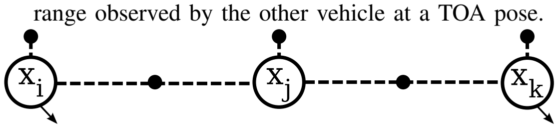

(c) Approximate local chain, $\mathcal{C}_{\text{approx}}$, reconstructed by our method, k-last, with $k \geq 1$.

Fig. 7: Comparison of reconstructed local chains by Paull et al. (b) and our method (c) as if the data transmitted at the second TOL, $\mathbf{x}_j$, was not received. Arrows away from the chain indicate TOLs, arrows toward the chain indicate TOAs.

ranges, but only when the receiving vehicle has received all involved poses—cascaded communication topologies are not supported, as with our method. Finally, the local subset of range observations is sufficient for accurate navigation in many practical situations (as evidenced in §V).

VII. Conclusion

Accurate localization extends the capacity of AUVs to perform ocean science. OWTT underwater cooperative localization promises improved navigation for AUVs over larger area and time scales without additional infrastructure. We exploited the structure of the composition operation and an accurate approximation of the local chain to robustly share locally observed sensor data across a fragile communication channel. We then used the collection of received chains and observed relative pose measurements to compute an improved navigation solution.

Our proposed method can also find application in the broader robotics community where limited communication is an operational factor. Avenues for future work include extending our chain approximation method to include information received from other vehicles. This would allow better localization performance in certain communication topologies, for example, cascaded networks.

References

[1] R. M. Eustice, H. Singh, and L. L. Whitcomb, "Synchronous-clock one-way-travel-time acoustic navigation for underwater vehicles," *J. Field Robot.*, vol. 28, no. 1, pp. 121–136, 2011.

[2] J. M. Walls and R. M. Eustice, "An origin state method for communication constrained cooperative localization with robustness to packet loss," *Int. J. Robot. Res.*, vol. 33, no. 9, pp. 1191–1208, 2014.

[3] J. Vial, H. Durrant-Whyte, and T. Bailey, "Conservative sparsification for efficient and consistent approximate estimation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, San Francisco, CA, Sep. 2011, pp. 886–893.

[4] N. Carlevaris-Bianco, M. Kaess, and R. M. Eustice, "Generic node removal for factor-graph SLAM," *IEEE Trans. Robot.*, vol. 30, no. 6, pp. 1371–1385, 2014.

[5] G. Huang, M. Kaess, and J. Leonard, "Consistent sparsification for graph optimization," in *Proc. European Conf. Mobile Robots*, Barcelona, Spain, Sep. 2013, pp. 150–157.

[6] M. Mazuran, G. D. Tipaldi, L. Spinello, and W. Burgard, "Nonlinear graph sparsification for SLAM," in *Proc. Robot.: Sci. & Syst. Conf.*, Berkeley, CA, Jun. 2014.

[7] R. Kurazume, S. Nagata, and S. Hirose, "Cooperative positioning with multiple robots," in *Proc. IEEE Int. Conf. Robot. and Automation*, 1994, pp. 1250–1257.

[8] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," *Autonomous Robots*, vol. 8, no. 3, pp. 325–344, 2000.

[9] S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," *IEEE Trans. Robot. Autom.*, vol. 18, no. 5, pp. 781–795, 2002.

[10] A. Howard, M. J. Matari, and G. S. Sukhatme, "Localization for mobile robot teams using maximum likelihood estimation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, EPFL, Switzerland, Sep. 2002, pp. 434–439.

[11] F. Dellaert, F. Alegre, and E. B. Martinson, "Intrinsic localization and mapping with 2 applications: Diffusion mapping and marco polo localization," in *Proc. IEEE Int. Conf. Robot. and Automation*, Taipei, Taiwan, Sep. 2003, pp. 2344–2349.

[12] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller, "Multiple relative pose graphs for robust cooperative mapping," in *Proc. IEEE Int. Conf. Robot. and Automation*, Anchorage, AK, May 2010, pp. 3185–3192.

[13] A. Cunningham, V. Indelman, and F. Dellaert, "DDF-SAM 2.0: Consistent distributed smoothing and mapping," in *Proc. IEEE Int. Conf. Robot. and Automation*, Karlsruhe, Germany, May 2013, pp. 5220–5227.

[14] K. Leung, T. Barfoot, and H. Liu, "Decentralized localization of sparsely-communicating robot networks: A centralized-equivalent approach," *IEEE Trans. Robot.*, vol. 26, no. 1, pp. 62–77, 2010.

[15] E. Nerurkar and S. Roumeliotis, "A communication-bandwidth-aware hybrid estimation framework for multi-robot cooperative localization," in *Proc. IEEE Int. Conf. Robot. and Automation*, Tokyo,Japan, Nov. 2013, pp. 1418–1425.

[16] J. Vaganay, J. J. Leonard, J. A. Curcio, and J. S. Wilcox, "Experimental validation of the moving long base-line navigation concept," in *Proc. IEEE/OES Autonomous Underwater Vehicles Conf.*, Sebasco, ME, Jun. 2004, pp. 59–65.

[17] D. K. Maczka, A. S. Gadre, and D. J. Stilwell, "Implementation of a cooperative navigation algorithm on a platoon of autonomous underwater vehicles," in *Proc. IEEE/MTS OCEANS Conf. Exhib.*, Vancouver, BC, Oct. 2007, pp. 1–6.

[18] A. Bahr, J. J. Leonard, and M. F. Fallon, "Cooperative localization for autonomous underwater vehicles," *Int. J. Robot. Res.*, vol. 28, no. 6, pp. 714–728, 2009.

[19] A. Bahr, M. R. Walter, and J. J. Leonard, "Consistent cooperative localization," in *Proc. IEEE Int. Conf. Robot. and Automation*, Kobe, Japan, May 2009, pp. 3415–3422.

[20] M. F. Fallon, G. Papadopoulis, and J. J. Leonard, "A measurement distribution framework for cooperative navigation using multiple AUVs," in *Proc. IEEE Int. Conf. Robot. and Automation*, Anchorage, AK, May 2010, pp. 4256–4263.

[21] S. E. Webster, R. M. Eustice, H. Singh, and L. L. Whitcomb, "Advances in single-beacon one-way-travel-time acoustic navigation for underwater vehicles," *Int. J. Robot. Res.*, vol. 31, no. 8, pp. 935–950, 2012.

[22] S. E. Webster, J. M. Walls, R. M. Eustice, and L. L. Whitcomb, "Decentralized extended information filter for single-beacon cooperative acoustic navigation: Theory and experiments," *IEEE Trans. Robot.*, vol. 29, no. 4, pp. 957–974, 2013.

[23] L. Paull, M. Seto, and J. J. Leonard, "Decentralized cooperative trajectory estimation for autonomous underwater vehicles," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, Chicago, IL, Sep. 2014, pp. 184—191.

[24] F. Dellaert and M. Kaess, "Square root SAM," *Int. J. Robot. Res.*, vol. 25, no. 12, pp. 1181–1204, 2006.

[25] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Proc. Uncertainty AI.* Elsevier, 1986, pp. 267–288.

[26] N. Carlevaris-Bianco and R. M. Eustice, "Learning visual feature descriptors for dynamic lighting conditions," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, Chicago, IL, Sep. 2014, pp. 2769–2776.