

Communication Planning for Cooperative Terrain-Based Underwater Localization

Jacob A. Anderson¹ and Geoffrey A. Hollinger¹

Abstract— We present a decentralized communication planning algorithm for terrain-based navigation (dec-TBN). The proposed algorithm uses forward simulation to approximate the value of communicating at each time step. The simulations are used to build a constrained tree graph that can be searched to provide a minimum cost communication schedule. The algorithm is evaluated in simulations using a real-world bathymetry map from Lake Nighthorse, CO and sensor model derived from the Ocean Server Iver2 vehicle. Results show that the algorithm finds a communication schedule that improves robot localization by up to 27% compared to non-cooperative terrain-based navigation.

I. INTRODUCTION

Marine robots are increasingly being used for underwater data collection [1], [2], [3]. Electromagnetic waves quickly attenuate in water, prohibiting the use of typical localization sensors such as GPS, LIDAR, and cameras. Inertial sensors are available, but they are either cost prohibitive or have too much noise and drift for reliable localization. Acoustics is the primary mode of ranging and communication but similarly tends to be either cost prohibitive or provide low information throughput. Due to these restrictions, terrain-based navigation (TBN) has become the leading means of self-contained underwater localization [4].

Vehicles using TBN for underwater localization are dependent on terrain information to improve their state estimation. Vehicles traveling over areas with distinctive terrain will be able to localize better than those traveling over areas of smooth terrain. Cooperative localization allows vehicles with better state estimation to aid other vehicles. A vehicle with accurate localization can transmit its location and covariance to other vehicles via an acoustic modem. A receiving vehicle can calculate its distance from the transmitting vehicle and use this distance measurement with the information provided by the transmitting robot to improve its own localization [5].

This work focuses on planning when each robot communicates its localization information which is important for two reasons: (1) overlapping communication can cause interference, resulting in failed communication, and (2) acoustic modem bandwidth is extremely limited and often needed to transmit other information, such as scientific data.

This work presents a decentralized communication planning algorithm that determines an optimized communication policy for collaborative underwater TBN localization. The algorithm forward simulates a group of two or more robots following predetermined paths. At each time step a scenario in which the hosting robot communicates is considered and

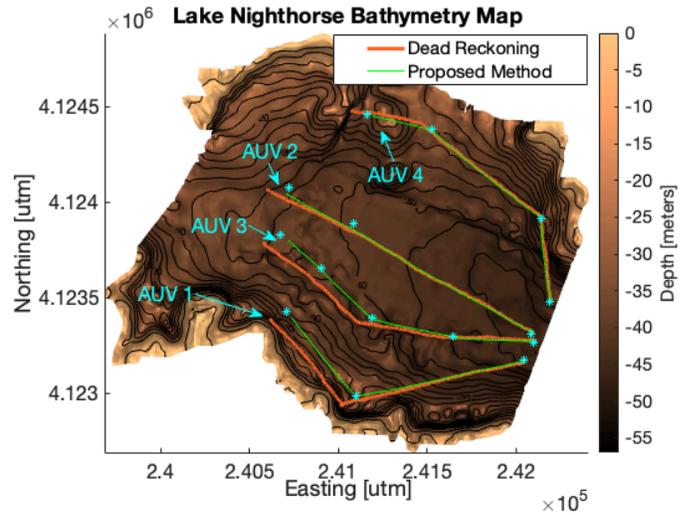


Fig. 1. AUV paths through the simulated environment of Lake Nighthorse, CO. The blue stars are the AUVs’ waypoints. The green lines are the actual paths of the AUVs when using decentralized terrain-based navigation (dec-TBN), and the orange lines are the actual paths of the AUVs when using dead reckoning. Cooperative TBN leads to more accurate path following (closer to the waypoints) even in areas where the terrain is flat. The AUVs travel approximately 1.5km.

then compared to scenarios in which the hosting robot has not communicated. Each communication incurs a cost, and scenarios that result in poor localization are discarded. The resulting communication policy contains a minimum number of communication while limiting the uncertainty in each vehicle’s location. The proposed algorithm enables more accurate localization while conserving energy and allowing opportunities for other types of data to be transmitted. To our knowledge, this is the first algorithm to non-myopically plan the communication actions of underwater vehicles to improve cooperative localization.

We test our approach using simulations leveraging real-world data from an Ocean Server, Iver2 autonomous underwater vehicle (AUV) and a bathymetry map of Lake Nighthorse in Durango, Colorado, USA (See Figure 1). Experimental results are presented for simulation of two and four vehicles following paths through the map. Results show that the communication planning algorithm schedules communications that provide accurate localization while requiring very low bandwidth. In the case of two AUVs, the proposed algorithm schedules 67 communications when 1,058 are possible. The planned communication policy improves localization by 15% over non-cooperative TBN. In the case of four AUVs, 139 communications are scheduled, when 1,048 are possible, providing a 27% improvement in localization over non-cooperative TBN.

¹Collaborative Robotics and Intelligent Systems Institute, Oregon State University, Corvallis, OR 97331., anderja7@oregonstate.edu, geoff.hollinger@oregonstate.edu

“This research has been funded in part by U.S. Naval Facilities Engineering Command (NAVFAC) contract N0002418F8702 and ONR grant N00014-17-1-2581.”

II. RELATED WORK

A. Terrain-Based Navigation

TBN is a technique that originated in 1980 with TERCOM [6] where a flying object, such as a missile, would compare its altimeter readings to a digital elevation map. Recent implementations of TBN have been centered around particle filters that compare a vehicle's altimeter readings to a bathymetry map [7], [4]. TBN can provide an accurate vehicle location if it is in an area of significant bathymetric features. On smooth terrain the particle filters will diverge.

Tan et al. developed a decentralised TBN (dec-TBN) algorithm [5] where multiple vehicles share their location and covariance with each other. Each vehicle hosts its own particle filter and performs regular comparisons of its altimeter readings to a bathymetry map. The vehicles also transmit their particle filter's estimated location and covariance to the other vehicles. The receiving vehicles incorporate this information into their next TBN update. The distance between the vehicles is calculated using the acoustic communication's time of flight which is used with the transmitting vehicle's location and covariance to update the weights of the particle filter. In this work, the vehicles take turns communicating one after another with no consideration given to the timing of communication [5].

Dec-TBN forms the foundation of our work. We use a similar dec-TBN formulation in which each vehicle hosts a particle filter that is informed by altimeter measurements and location data transmitted from other vehicles. Our work builds on dec-TBN by examining communication planning for these vehicles to see if localization accuracy can be retained while reducing the communication overhead.

B. Communication Planning

To date, the majority of communication planning for state-of-the-art distributed localization is focused on choosing what data to share with other robots. One approach is to use metadata from the robots' pose graphs to identify individual scan lines or camera images that may contain loop closures [8], [9], [10]. Another approach is to design linear-quadratic regulators to control data flow [11]. These approaches are intended for terrestrial application where communication bandwidth and reliability are significantly better than underwater applications. In this case the robots are able to transfer large data sets to each other. The bandwidth available on an acoustic modem precluded these approaches.

A more applicable line of research is communication planning for multi-robot coordination. These methods focus on communicating the belief states of the robots for the purpose of deciding what actions they should take. This is analogous to decentralized TBN where the AUVs share their state estimations to help each other localize more accurately.

Williamson et al. applied information theoretics to communication planning for multi-robot cooperation by using KL divergence to quantify the reward of an agent's communication [12]. Their approach then uses this approximation in formulating a decentralized partially observable Markov decision process (DEC-POMDP) to remove reasoning over the value of communication from the POMDP's coordination model. Using a deterministic formula to approximate the value of each communication reduced the search space of the POMDP [12].

Unhelkar and Shah followed the idea of assigning value to communication and proposed a decentralized Markov decision process (DEC-MDP) with a reward function that maximized the expected reward for communication [13]. Marcotte et al. built on the aforementioned DEC-MDP and DEC-POMDP by factoring the planning problem so that each robot could plan independently. This resulted in the algorithm scaling linearly with the number of robots. Additionally, Marcotte et al. forward-simulated the outcomes of message passing to determine the value of each communication. This approach has the added advantage of being able to determine what the message content should be [14]. Similarly, Barcis et al. developed an evaluation model that determines the value of certain types of data. However, rather than using a Markov decision process, Barcis et al. built their evaluation model using domain knowledge of the application [15].

Best et al. considered planning-aware communication [16]. In this work a decentralized planning algorithm is presented in which a group of robots is attempting to complete a task. While the planner is evaluating which actions a robot should perform, it tracks its uncertainty in what actions it expects the other robots to perform. Once the uncertainty of a particular robot exceeds a certain threshold it requests a planning update from that robot. To minimize communication, the algorithm constructs an directed acyclic graph representing the uncertainty in the robots' actions and communication cost. While constructing the graph, every time a communication was requested from a robot the uncertainty in the robot's actions reduced to zero. Subsequently, the base node only had to be evaluated once. The implementation of this directed acyclic graph provides a communication schedule that is optimal with respect to belief and results in an algorithm with polynomial run time complexity [16]. Additionally, finding a communication policy is now equivalent to a longest path search through a directed acyclic graph, which has linear time complexity [17].

The communication planning problem presented here builds on these works in multi-robot coordination. The communication problem is factored so that each robot can plan its own communication separately from the other robots. Forward simulation is used to approximate the value of each communication. The simulations are then used to build a directed acyclic graph in which the robot is required to communicate once the uncertainty in a robot's location exceeds a certain threshold. The graph can then be searched for a minimum cost communication policy. Unlike the aforementioned works, which assumed that the robots have deterministic transition functions, the proposed communication planning applies these techniques to the inherently stochastic problem of TBN.

III. PROBLEM FORMULATION

We are interested in the problem of scheduling communication for a group of AUVs using dec-TBN to localize. The AUVs travel in a mapped environment taking altimeter readings as they move. Each AUV uses the TBN algorithm to estimate its location and the AUVs can communicate their location and uncertainty with each other. We want to plan a communication policy that uses minimal transmissions while limiting the uncertainty in the AUVs' location.

Each AUV r is equipped with an altimeter, a depth sensor, an acoustic modem, and a digital bathymetry map of the environment. The AUVs are also provided the starting locations $X_{t=0}^r$, initial covariances $\Sigma_{t=0}^r$ and paths of all the other AUVs. The AUVs will localize with a dec-TBN algorithm. An individual instantiation of the dec-TBN consists of a particle filter that tracks the hosting vehicle's location. At each time step, the inputs to dec-TBN are the most recent vehicle control inputs and altimeter reading. If available, X_t^i , Σ_t^i , and intra-vehicle range from another vehicle, $i \neq r$, are also provided to the dec-TBN. The dec-TBN algorithm outputs the position estimate X_t^r and covariance Σ_t^r of the hosting vehicle.

The AUVs can communicate with each other via the acoustic modems. If all of the vehicles have synchronized clocks, and assuming isotropic water temperatures, the distance between the transmitting and receiving vehicles can be calculated using the communication's one way time of flight (OWTF). To perform dec-TBN, the AUVs must take turns communicating their localization statistics. The receiving vehicles use the localization statistics and distance measurement from the transmitting vehicle to inform the next dec-TBN update.

The proposed decentralized planning algorithm generates a communication policy $\pi^r = \{\pi_1^r, \pi_2^r, \dots, \pi_T^r\}$ that indicates when the hosting vehicle should communicate its X_t^r and Σ_t^r within the planning horizon T . The variable π^r is a binary sequence, $\pi_t^r = true$ indicates that the vehicle should communicate at time step t and $\pi_t^r = false$ indicates that it should not communicate. If a vehicle communicates at time t , the transmission includes the time of the transmission, X_t^r and Σ_t^r . The communication is received by all the other AUVs and informs their next dec-TBN update.

The objective of the planning algorithm is to minimize the AUVs' communication while maintaining a bound on localization accuracy. Each communication is given a unit cost. The cost of π for N robots is the sum of all communication, i.e.:

$$Cost(\pi) = \sum_{i=1}^N \sum_{t=1}^T 1(\pi_t^i = true) \quad (1)$$

To ensure that the AUVs maintain a certain degree of accuracy, the communication planning algorithm is constrained by the estimated accuracy σ of the AUVs' localization. For a group of N AUVs at time step t , σ_t is defined as:

$$\sigma_t = \sum_{i=1}^N trace(\Sigma_t^i) \quad (2)$$

The proposed planning problem is formulated as a constraint-based optimization to find π^* that has the lowest communication cost while maintaining σ under a user defined threshold σ_{max} :

$$\pi^* = \underset{\pi}{\operatorname{argmin}} Cost(\pi) : \sigma_t < \sigma_{max} \quad \forall T \quad (3)$$

IV. COMMUNICATION PLANNING ALGORITHM

The proposed communication planning method is a decentralized planning algorithm that is intended for vehicles using dec-TBN. The intuition behind the algorithm is that

localization information from a vehicle with an accurate state estimate can be used to improve the localization of other vehicles. It may not be advantageous for vehicles with relatively poor localization to transmit their information either. Additionally, the dec-TBN particle filters use randomly generated noise to disperse the particles. Modeling a particle filter's response to a vehicle's path is impractical. For these reasons the proposed communication planning algorithm involves simulating vehicles traveling through the environments and evaluating the effects of vehicle communication on the group's localization.

The communication planning algorithm is a decentralized algorithm that is run independently on each vehicle. The algorithm builds a directed acyclic graph G . The graph nodes N represent the σ_t of the AUVs and the edges represent the communication cost between nodes. The graph is built by forward simulating a set of vehicles as they follow predefined paths through an environment. Each leaf node hosts a simulated state of the vehicles instantiated as a set of particle filters, e.g., if planning for three vehicles, each leaf node would contain three particle filters. The leaf nodes are expanded by forward simulating the vehicles one time step and localizing via dec-TBN on the bathymetry map. The leaf nodes are expanded twice, once with the hosting vehicle communicating its X_t and Σ_t , and once without communicating, thereby creating two new leaf nodes. The edges between the parent node and the new leaf nodes are given weights of 1 and 0 for communicating and not communicating respectfully.

To reduce computational demands of the algorithm, it is assumed that the state resulting from the host vehicle communicating is the same for all leaf nodes. This is a similar assumption to Best et al. [16] and the resulting structure of the graph is demonstrated in Figure 2. $N_{t,0}$ represents the state in which the host vehicle has communicated at time step t . The rest of the nodes result from not communicating. Constructing the graph as a directed acyclic graph provides polynomial run time with complexity $\mathcal{O}(BNT^2)$ where B is the number of particles used in the particle filters. Additionally, finding a communication policy from the directed acyclic graph can be done with linear time complexity [17].

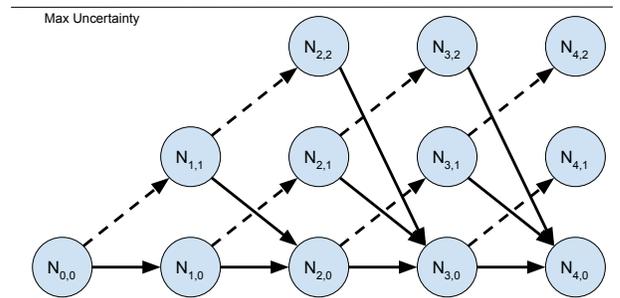


Fig. 2. Communication planning graph. The node N_n represents the n^{th} node in G . Each node holds the value of σ of a simulation state. The dotted edges represent forward simulation without communication. The solid edges represent forward simulation with communication.

In practice, the planning algorithm holds the simulations for each leaf node i in a queue Q_i . At each time step the algorithm cycles through the queue and progresses each

simulation one step forward. The simulations move the robots \mathcal{R} along a path P using a motion model M . Then, they update the dec-TBN algorithm using the simulated vehicles' locations on the bathymetry map to provide depth readings. If communication is indicated, the X_t and Σ_t of the hosting vehicle's simulated state is provided to the other simulated vehicles' dec-TBN updates.

At each time step all of the simulations in Q are forward simulated without communication. If σ_t of the resulting simulation state is less than σ_{max} , then the resulting state is added to Q and a new node is appended to G with the value of σ_t . This node is connected to the leaf node representing the state that was just updated with an edge weight of 0. Q_0 holds the simulation for $N_{t,0}$ in which the hosting vehicle has just communicated. The last update performed at each time step is to forward simulate Q_0 with communication. The resulting state is added to Q as the next Q_0 and a corresponding node is added to G with the simulation's value of σ_t . This node is connected to all of the leaf nodes with an edge weight of 1. Due to dec-TBN's reliance on terrain information to improve localization, a static threshold value for σ_{max} is difficult to determine. Instead, the threshold value is a product of σ_{max} and σ_t of Q_0 . See Algorithm 1.

Algorithm 1 Communication Planning Algorithm

Input: $[X_0, \Sigma_0, M, P]^r, \forall r \in \mathcal{R}, T, \sigma_{max}$
Output: π \triangleright Robot Communication Policy

- 1: \triangleright **Particle filter for each robot**
- 2: $Q_0.sim \leftarrow \text{PF}([X_0, \Sigma_0, M, P]^r, \forall r \in \mathcal{R})$ \triangleright Queue
- 3: $Q_0.parent \leftarrow 0$
- 4: $G.nodes_0 \leftarrow 0$ \triangleright Cost Graph
- 5: $m = 0$
- 6: **for** $t = 1$ to T **do**
- 7: $\mathcal{P} = Q.parent$ \triangleright List of current leaf nodes
- 8: \triangleright **Update Q_i without communication**
- 9: **for** $i = |Q|$ to 1 **do**
- 10: $\eta \leftarrow \text{PF_Update}(Q_i, a = \text{false})$
- 11: **if** $\eta.\sigma < Q_0.\sigma \times \sigma_{max}$ **then**
- 12: $m = m + 1$
- 13: $p = Q_i.parent$
- 14: $Q_{i+1}.sim \leftarrow \eta$
- 15: $Q_{i+1}.parent \leftarrow m$
- 16: $G.nodes_m \leftarrow \eta.\sigma$
- 17: $G.edges_{p,m} \leftarrow 0$
- 18: \triangleright **Update Q_i with communication**
- 19: $m = m + 1$
- 20: $Q_0.sim \leftarrow \text{PF_UPDATE}(N_{0,t-1}, a = \text{true})^r, \forall r \in \mathcal{R}$
- 21: $Q_0.parent \leftarrow m$
- 22: $G.nodes_m \leftarrow Q_0.\sigma$
- 23: $G.edges_{\mathcal{P},m} \leftarrow 1$
- 24: $\pi \leftarrow \text{LOWESTCOSTPATH}(G)$ \triangleright Search over Graph

V. DECENTRALIZED TBN ALGORITHM

Most modern TBN algorithms use particle filters to track the vehicle's position on a digital elevation map or bathymetry map in the case of marine environments [18], [19]. The dec-TBN algorithm used here utilizes an update step that incorporates range measurements to another vehicle

with that vehicle's localization information. First, a vehicle transmits its X_t and Σ_t to the other vehicles via an acoustic modem. A receiving vehicle calculates a range measurement D from the transmitting vehicle via the acoustic communication's OWTF. D provides a measurement that adds information to the particle filter and is used with the location information from the transmitting vehicle in the next particle filter update.

The modified particle filter update propagates the particles via the vehicle's speed S , heading θ , and motion model M . Then it calculates the probability of each particle's location given a depth measurement z from an altimeter. A probability density function pdf_{alt} using the altimeter's mean and standard deviation is used to compare the depth measurement to the expected depths \mathcal{Z} from the bathymetry map. The probability of each particle's location is also calculated by creating a multivariate normal distribution pdf_{coms} using the location and covariance received from the other vehicle. The locations of the particles are then shifted towards the transmitting vehicle's location by the distance calculated from the acoustic modem's time of flight. The resulting particle locations are used to sample the aforementioned normal distribution. The probabilities resulting from the bathymetry measurement and the communication measurement are multiplied together with the previous particle weights to create the new particle weights. This process is illustrated in Algorithm 2.

Algorithm 2 Particle Filter Update with a Received Communication

- 1: **function** PF_UPDATE($X_0, T, S, \theta, z, D, X^r, \Sigma^r$)
- 2: $particles \leftarrow X_0$ \triangleright Initialize particles
- 3: $w \leftarrow 1/|particles|$ \triangleright Initialize particle weights
- 4: **for** $t = 0$ to T **do**
- 5: $particles \leftarrow \text{MOVE_PARTICLES}(S, \theta, t, M)$
- 6: $\mathcal{Z} \leftarrow \text{BATHYMETRYMAP}(particles)$
- 7: $w_{bathy} \leftarrow \text{PDF}_{alt}(z - \mathcal{Z}, \mu_{alt}, \sigma_{alt})$
- 8: $w_{comms} \leftarrow \text{PDF}_{coms}(particles + D, X^r, \Sigma^r)$
- 9: $w \leftarrow w \times w_{bathy} \times w_{comms}$
- 10: $w \leftarrow w / \sum(w)$
- 11: $X \leftarrow \sum(weights \times particles)$
- 12: $\Sigma \leftarrow \text{Covariance}(particles)$
- return** X, Σ

VI. RESULTS

To evaluate the proposed communication planning algorithm, simulations were run to determine how well AUVs using dec-TBN could localize given the communication policy produced by the algorithm. The simulations leverage real-world data by using a bathymetry map of Lake Nighthorse near Durango, Colorado, USA. The map was created from an extensive survey of the reservoir. Depth readings from the survey were corrected for temporal changes in the height of the reservoir. Then, the depths were combined into a digital elevation model via a sliding window Kalman filter (See Figure 1). The depth readings from this survey were also used to build a sensor model of the altimeter on an Ocean Server Iver2 AUV that was used for part of the data collection. This sensor model was used in the simulator to provide depth readings to the simulated vehicles.

Each simulation involved two or more AUVs following predefined paths across the bathymetry map. Each AUV was given the initial locations of all the vehicles with a corresponding uncertainty of three meters. The AUVs were also given paths that each vehicle would follow. The AUVs performed the communication planning algorithm independently before departing on their paths. Once underway, the AUVs attempted to follow the prescribed paths by using dec-TBN. The AUVs communicated their locations and covariances at the time steps indicated by the communication policy. The AUVs did not share their communication policies, so it was possible that communications would overlap. In this case, the communications were assumed to interfere with each other and were not received by any of the vehicles.

Each AUV used its dec-TBN state estimate to compute control inputs. Gaussian white noise was introduced into the vehicle's true movements to emulate the navigational errors that occur in real underwater vehicles. The sensor model derived from the Iver2's altimeter was used to introduce noise into the depth readings used for the AUVs' dec-TBN. Figure 1 shows the simulated paths of four AUVs. The blue stars are waypoints that the AUVs followed. The green lines are the actual paths that the AUVs traveled when using dec-TBN. The orange lines are the actual paths that the AUVs traveled when using dead reckoning. The AUVs traveled a little more than 1.5 km in these simulations.

The AUVs' localization accuracy in the simulation is evaluated by comparing the AUVs' state estimations to their true locations. The error δ_t^r for AUV r at time step t is calculated as $\delta_t^i = \|X_t^i - \hat{X}_t^i\|$. In this case, X is the true location of the AUV and \hat{X} is the TBN estimate. The joint error Δ_t of all N AUVs is:

$$\Delta_t = \sum_{i=1}^{i=N} \delta_t^i \quad (4)$$

The particle filter in the dec-TBN algorithm uses stochastically generated noise when moving the particles. Additionally, the simulations add Gaussian white noise to the ground truth movements of the vehicles and the vehicles' sensor measurements. To determine the relative average joint error, every simulation was run 100 times.

To the authors' knowledge, previous dec-TBN works only use a full communication scheme where the AUVs take turns communicating at each time step. The planning algorithm results are compared to simulations with full communication and communication that happens on an incremented schedule. The comparison schedules varied the amount of bandwidth used. Lower bandwidth schedules involved the vehicles communicating in evenly spaced blocks. Each block of communication involved all of the vehicles taking a turn to communicate. The blocks were spaced so that the total number of communications used equaled a percentage of the full communication. The bandwidths are listed in the left columns of Table I and II with the communication policy generated by the planning algorithm. The center columns list the corresponding number of communications. The right column shows the total error for each communication scheme. The total error is calculated as the area under the curve of the relative average joint errors which can be seen in Figures 3, 4, 5 and 6. Note that dead reckoning error is

not present on the figures because it is much greater than the errors being shown.

TABLE I
COMMUNICATION BANDWIDTH AND TOTAL JOINT ERROR FOR 2 AUVS

Communication Bandwidth	Number of Communications	Total Error [m · s]
Policy (Proposed)	67	11,004
Dead Reckoning	0	73,644
TBN - No Comms.	0	12,911
Full (100%)	1058	19,694
80%	816	17,645
60%	612	16,523
40%	408	15,088
20%	204	12,853
10%	102	12,044
5%	50	11,669
2.5%	26	11,782

TABLE II
COMMUNICATION BANDWIDTH AND TOTAL JOINT ERROR FOR 4 AUVS

Communication Bandwidth	Number of Communications	Total Error [m · s]
Policy (Proposed)	139	8,929
Dead Reckoning	0	73,514
TBN - No Comms.	0	12,180
Full (100%)	1048	18,227
80%	840	15,489
60%	621	13,052
40%	416	11,066
20%	208	9,730
10%	108	9,341
5%	52	9,820
2.5%	28	10,607

Figures 3 and 5 show an overview of the joint errors for some of the communication policies evaluated in simulations with two and four AUVs respectively. The presented policies include the policy produced by the planning algorithm, the full communication policy, some of the proportional communication policies, and non-cooperative TBN. Figures 4 and 6 provide closer looks at the joint errors for the planned communication and lower bandwidth communication policies for simulations with two and four AUVs. Note that the 5% communication policy is present on all graphs for continuity.

Figures 3, 4, 5 and 6 show that the proposed algorithm finds a communication schedule that provides more accurate localization. In all graphs the joint error experienced by the planned communication policy is less than the other communication policies. For the simulations with two AUVs the planning algorithm schedules a total of 67 communications which is similar to low bandwidth policies of 5% and 2.5% which use 50 and 26 communications respectively. For the simulations with four AUVs the algorithm schedules 139 communications which is similar to the 10% bandwidth policy. This is nearly twice as many communications as prescribed by the planning algorithm for the two AUV simulations and achieves notably more accurate localization. This is due to the greater amount of information that is available from the two additional AUVs.

The lower bandwidth communication schemes perform better in general because they reduce the number of noisy measurements being transmitted. Additionally, the higher communication schemes can cause the TBN particle filters

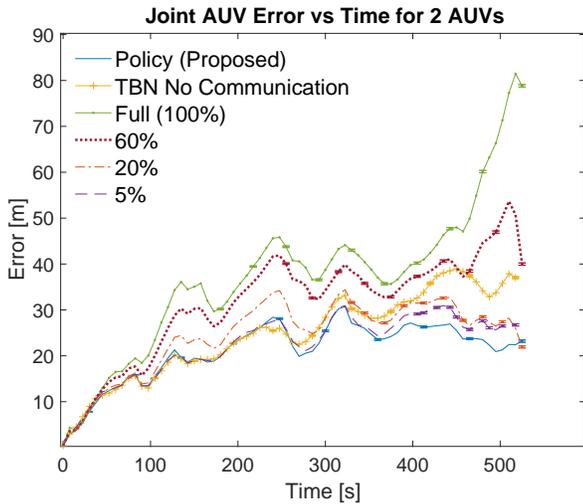


Fig. 3. Overview of simulation results for 2 AUVs using varying communication bandwidths. Error bars show standard error of the mean.

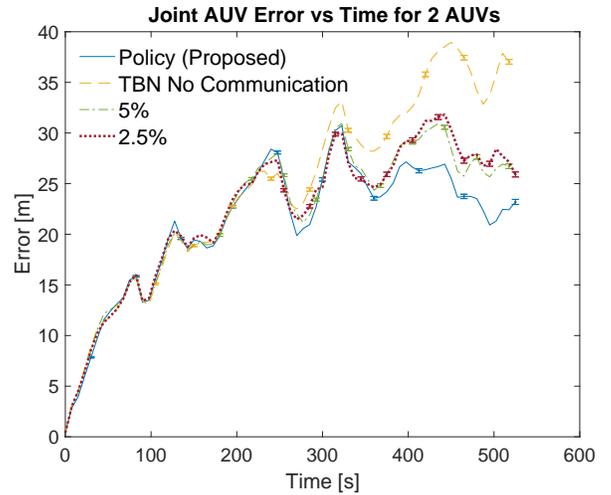


Fig. 4. Simulation results for 2 AUVs using the planned policy, TBN without communication, and low bandwidth communication policies. Error bars show standard error of the mean.

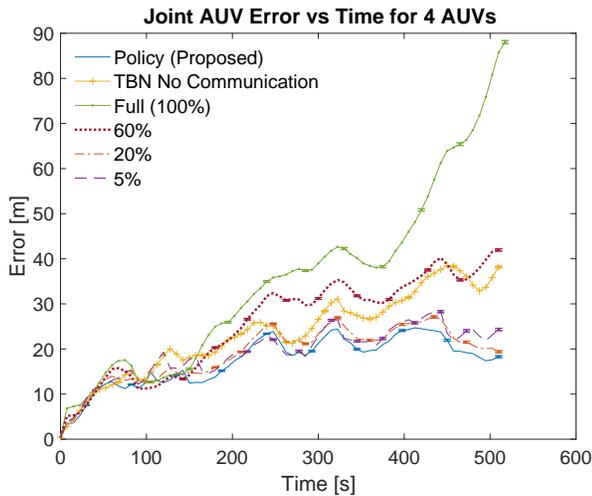


Fig. 5. Overview of simulation results for 4 AUVs using varying communication bandwidths. Error bars show standard error of the mean.

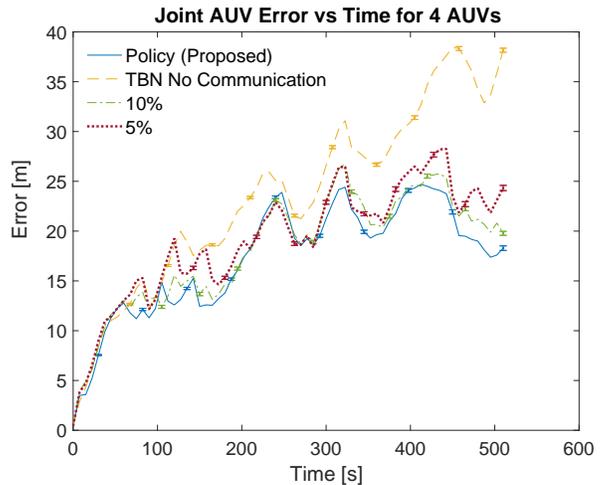


Fig. 6. Simulation results for 4 AUVs using the planned policy, TBN without communication, and low bandwidth communication policies. Error bars show standard error of the mean.

to become overly confident in their state estimate. This over confidence can result in the filter diverging, especially in these scenarios where at least one AUV is traveling through an area with minimal terrain features.

VII. CONCLUSION AND FUTURE WORK

We have proposed a communication planning algorithm for AUVs using dec-TBN. The algorithm uses forward simulation to estimate the accuracy of each vehicle's state estimation at a given time step and, subsequently, the value of that vehicle communicating. By using the simulations to build a communication tree graph, a policy can be determined that will provide accurate localization with minimal communication. We present results showing that the planned communication policies provide more accurate localization than non-cooperative TBN and regularly incremented communication policies.

Future work for this algorithm includes adding an acoustic modem model and field trials. An acoustic modem model will provide an effective range for acoustic communication. This will enable the algorithm to account for vehicles being out of communication range. In this case a vehicle transmitting its localization statistics may not be as effective since it won't help vehicles that are out of range. At the same time, overlapping communications may not interfere with each other if the transmitting vehicles are out of range.

Field trials will provide an empirical evaluation of the planning algorithm. Ideally the field trials will be held in areas with varying bathymetric features. This will allow for the algorithm to be evaluated in areas where rich terrain features are available, as well as flat areas where TBN tends to fail. Aside from validating the performance of the algorithm, such trials would give insight into how the number of communications scale depending on the amount of information that is available from the environment.

REFERENCES

- [1] M. E. Furlong, D. Paxton, P. Stevenson, M. Pebody, S. D. McPhail, and J. Perrett, "Autosub long range: A long range deep diving AUV for ocean monitoring," in *Proc. IEEE/OES Autonomous Underwater Vehicles (AUV)*, 2012, pp. 1–7.
- [2] R. B. Wynn, V. A. Huvenne, T. P. Le Bas, B. J. Murton, D. P. Connelly, B. J. Bett, H. A. Ruhl, K. J. Morris, J. Peakall, D. R. Parsons *et al.*, "Autonomous Underwater Vehicles (AUVs): Their past, present and future contributions to the advancement of marine geoscience," *Marine Geology*, vol. 352, pp. 451–468, 2014.
- [3] T. O. Fossum, G. M. Fragoso, E. J. Davies, J. E. Ullgren, R. Mendes, G. Johnsen, I. Ellingsen, J. Eidsvik, M. Ludvigsen, and K. Rajan, "Toward adaptive robotic sampling of phytoplankton in the coastal ocean," *Science Robotics*, vol. 4, no. 27, 2019.
- [4] G. Salavasidis, A. Munafò, C. A. Harris, T. Prampart, R. Templeton, M. Smart, D. T. Roper, M. Pebody, S. D. McPhail, E. Rogers *et al.*, "Terrain-aided navigation for long-endurance and deep-rated autonomous underwater vehicles," *Journal of Field Robotics*, vol. 36, no. 2, pp. 447–474, 2019.
- [5] Y. T. Tan, M. Chitre, and F. S. Hover, "Cooperative bathymetry-based localization using low-cost autonomous underwater vehicles," *Autonomous Robots*, vol. 40, no. 7, pp. 1187–1205, 2016.
- [6] J. P. Golden, "Terrain contour matching (TERCOM): a cruise missile guidance aid," in *Image processing for missile guidance*, vol. 238. International Society for Optics and Photonics, 1980, pp. 10–18.
- [7] F. Maurelli, S. Krupinski, Y. Petillot, and J. Salvi, "A particle filter approach for AUV localization," in *Proc. IEEE OCEANS 2008*, 2008, pp. 1–7.
- [8] Y. Tian, K. Khosoussi, M. Giamou, J. P. How, and J. Kelly, "Near-optimal budgeted data exchange for distributed loop closure detection," *arXiv preprint arXiv:1806.00188*, 2018.
- [9] M. Giamou, K. Khosoussi, and J. P. How, "Talk resource-efficiently to me: Optimal communication planning for distributed loop closure detection," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 1–9.
- [10] Y. Tian, K. Khosoussi, and J. P. How, "A resource-aware approach to collaborative loop closure detection with provable performance guarantees," *arXiv preprint arXiv:1907.04904*, 2019.
- [11] A. Kassir, R. Fitch, and S. Sukkarieh, "Communication-efficient motion coordination and data fusion in information gathering teams," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 5258–5265.
- [12] S. Williamson, E. Gerding, and N. Jennings, "A principled information valuation for communications during multi-agent coordination," *Proc. AAMAS Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains*, pp. 137–151, 2008.
- [13] V. V. Unhelkar and J. A. Shah, "Contact: Deciding to communicate during time-critical collaborative tasks in unknown, deterministic domains," in *Proc. Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [14] R. J. Marcotte, X. Wang, D. Mehta, and E. Olson, "Optimizing multi-robot communication under bandwidth constraints," *Autonomous Robots*, vol. 44, no. 1, pp. 43–55, 2020.
- [15] M. Barciś, A. Barciś, and H. Hellwagner, "Information distribution in multi-robot systems: Utility-based evaluation model," *Sensors (Basel, Switzerland)*, vol. 20, no. 3, 2020.
- [16] G. Best, M. Forrai, R. R. Mettu, and R. Fitch, "Planning-aware communication for decentralised multi-robot coordination," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 1050–1057.
- [17] G. Best, W. Martens, and R. Fitch, "Path planning with spatiotemporal optimal stopping for stochastic mission monitoring," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 629–646, 2017.
- [18] S. Carreno, P. Wilson, P. Ridaou, and Y. Petillot, "A survey on terrain based navigation for AUVs," in *Proc. IEEE OCEANS 2010*, 2010, pp. 1–7.
- [19] J. Melo and A. Matos, "Survey on advances on terrain based navigation for autonomous underwater vehicles," *Ocean Engineering*, vol. 139, pp. 250–264, 2017.