

DRACo-SLAM: Distributed Robust Acoustic Communication-efficient SLAM for Imaging Sonar Equipped Underwater Robot Teams

John McConnell, Yewei Huang, Paul Szenher, Ivana Collado-Gonzalez and Brendan Englott

Abstract—An essential task for a multi-robot system is generating a common understanding of the environment and relative poses between robots. Cooperative tasks can be executed only when a vehicle has knowledge of its own state and the states of the team members. However, this has primarily been achieved with direct rendezvous between underwater robots, via inter-robot ranging. We propose a novel distributed multi-robot simultaneous localization and mapping (SLAM) framework for underwater robots using imaging sonar-based perception. By passing only scene descriptors between robots, we do not need to pass raw sensor data unless there is a likelihood of inter-robot loop closure. We utilize pairwise consistent measurement set maximization (PCM), making our system robust to erroneous loop closures. The functionality of our system is demonstrated using two real-world datasets, one with three robots and another with two robots. We show that our system effectively estimates the trajectories of the multi-robot system and keeps the bandwidth requirements of inter-robot communication low. To our knowledge, this paper describes the first instance of multi-robot SLAM using real imaging sonar data (which we implement offline, using simulated communication). Code link: <https://github.com/jake3991/DRACo-SLAM>.


I. INTRODUCTION

Underwater robotics has proliferated over the past decade, supporting a variety of tasks, including ship hull inspection, harbor security, maintenance-inspection-repair (MIR), and intelligence surveillance reconnaissance (ISR). However, autonomous underwater vehicles (AUVs) are limited in their sensing capabilities. Firstly, the global positioning system (GPS) cannot be used due to the attenuation of GPS signals in the water. While AUVs can surface to collect GPS measurements, this is inefficient and impossible during under-ice missions or certain tactical situations. Further, if a small, lightweight, low-power perceptual package is desired, underwater lidar may be infeasible, and the use of sonar may be required, especially in turbid or dark environments. While the capabilities of sonar-based systems have expanded to include fully autonomous operation, the primary area of contribution is still focused on single agents executing their missions. However, real-world problems often benefit from a team of robots, a multi-robot system. When considering tasks ranging from autonomous exploration to ISR, a multi-robot system can reduce time, energy expenditure or create redundancy when operating in an adversarial environment.


A fundamental capability for both single and multi-robot systems is state estimation. When operating in unknown en-

J. McConnell, Y. Huang, P. Szenher, I. Collado-Gonzalez and B. Englott are with Stevens Institute of Technology, Hoboken, NJ, 07030, USA {jmconn1,yhuang85,pszenher,icollado,benglott}@stevens.edu


This research was supported in part by ONR Grant N00014-21-1-2161 and NSF Grants IIS-1652064 and IIS-1723996.



(a) Our experimental “AUV”: a custom-instrumented BlueROV2.



(b) Satellite image of the SUNY Maritime marina.



(c) An example multi-robot SLAM run from the environment in (b), from the perspective of one robot. This robot's poses are shown as black dots connected by several factors: sequential scan matching factors (SSM) in green, non-sequential scan matching factors (NSSM, intra-robot loop closures) shown in red, inter-robot (IR) loop closures in blue and partner robot (PR) factors in purple. This robot's map is shown in red with blue points merged from the other robot in the mission.

Fig. 1: Distributed multi-robot SLAM overview: (a) the robot used in our experiments, (b) one of the settings for our experiments, and (c) a representative experimental result are shown.

vironments, simultaneous localization and mapping (SLAM) is utilized to estimate vehicle location and provide situational awareness. For underwater robots, SLAM has typically been implemented to localize a single vehicle. Recent works on SLAM for underwater multi-robot systems consider both direct encounters (e.g., inter-robot range measurements) and indirect encounters (e.g., commonly observed targets in the survey area) as a means for fusing robot state estimates [1], [2]. There are some fundamental challenges to implementing such systems. Firstly, acoustic communication links have notably low bandwidth, limiting information exchange between robots. Another issue is the nature of sonar as a tool for perception. Sonar has a low signal-to-noise ratio,

low resolution, and often lacks 3D information. Moreover, there are no (or poor) initial conditions relating the robots in a multi-robot system without GPS.

In this work, we propose a system for distributed multi-robot SLAM that uses an imaging sonar's perceptual data as a basis for state estimation, and we validate it over real-world datasets. In addition to utilizing real perceptual data, we do not provide any initial conditions relating the robots a priori, and we do not rely upon inter-robot ranging, permitting an exclusive reliance on indirect encounters between robots. Lastly, we consider the bandwidth limitations of wireless acoustic communications and design our system for compatibility with realistic constraints (we note that beyond bandwidth, there are additional limitations associated with real-world acoustic comms [3]). Our key contributions are:

- The first underwater multi-robot SLAM system developed for use with real imaging sonar data.
- A high-performance pipeline for robustly identifying, registering, and rejecting outliers of indirect encounters using real sonar data with no initial conditions.
- A strategy to manage bandwidth usage, making real world application with acoustic modems a possibility.
- A validation of our system on two real-world data-sets.

In the following sections, we discuss related work, mathematically define our problem, present our system and results.

II. RELATED WORK

A. Underwater Multi-Robot SLAM

In the underwater space, inter-robot constraints generally come in two forms [2]: *direct* encounters where robots observe one another via acoustic ranging, and *indirect* encounters, where robots observe the same targets in the environment and may derive inter-robot measurement constraints relating one another. We note that direct encounters require synchronized clocks, which may not be practical over long periods without GPS clock corrections due to clock drift.

We first consider [4], where a mobile base station is used to localize a team of low-cost vehicles lacking perceptual sensors using an acoustic beacon. A similar concept is considered in [5], except instead of a team of robots, a leader-follower arrangement is used. [6] considers inter-robot ranging without fixed acoustic beacons and performs a simulation study comparing the use of fixed beacons to a cooperative localization solution. [7] considers an algorithm for processing inter-robot acoustic pulses in a distributed manner. [1] proposes a pose-graph-based method for cooperative localization of a team of robots using dead-reckoning, GPS, and inter-robot ranging. The outcome of this work is a system where a robot maintains an understanding of its state, and the team's state. Moreover, when GPS measurements occur at the surface, their effect is shared across robots. [2] integrates perception into the above, and commonly observed features are shared and integrated into robot pose-graphs. [2] assumes that robots share point-landmark observations in their survey area, but the framework is only tested in simulation. Further, due to GPS, each robot is effectively

localized in a common frame, enabling the sharing of range-bearing measurements to commonly observed targets without a need to solve the complex data association problem. Recall, GPS may not always be available due to under-ice operations or tactical situations (e.g., jamming). [8] considers the robot map merging problem using only similarity in feature space, in this case, ship hull curvature. While this work is able to merge trajectories and lower the data transmission requirements between robots, it requires a highly descriptive feature vector. It is only validated over a single dimension, whereas we consider a 3DOF system.

There are some notable examples of underwater multi-robot SLAM using cameras [9]–[12] and cameras with other navigation sensors [13]. Cameras, however are not robust in all water conditions. While these works are informative, because we operate in turbid conditions with low visibility, we do not consider them further.

Another area to consider is the concept of multi-session SLAM, where a robot is provided with a prior SLAM run and merges the map it builds with the prior run as the mission progresses. This is examined for ship hull inspection [14], [15], bathymetric mapping [16], ship-wreck reconstruction [17] and environmental monitoring [18]. However, multi-session does not account for communicating information between robots, as it is a single agent system with prior information. Additionally, it may be the case that the robot has some knowledge of its location in the prior map, constraining the inter-robot loop closure search space.

In contrast to the above, we utilize no notion of an initial guess relating robot reference frames. Second, we consider the bandwidth limitations of real acoustic modem hardware and take steps to manage network utilization. Lastly, we implement a fully functioning system to detect and estimate indirect encounters in sonar data: inter-robot loop closures.

B. Place Recognition With Sonar

Place recognition (loop closure) is fundamental for a perception-driven multi-robot system. Place recognition has been widely studied in LIDAR sensing [19]–[21]. These works assemble a 2D bird's eye view image of a 3D LIDAR scan with coarse discretization to compare scenes. They also derive a 1D descriptor to support scene search and retrieval.

Place recognition has also been studied with sonar-equipped AUVs. Firstly, [22] considers building scene graphs to compare scenes and evaluate rigid-body transformations but requires at least fifteen objects in each scene to run. Machine learning has also been used to support this task in [23], [24]. However, few public sonar datasets exist, and this is often a research area unto itself. Iterative closest point (ICP) based loop closure is used in our prior work [25] to support single-agent active SLAM. This work uses sonar derived point clouds with ICP; when ICP provides a transform between keyframes, we estimate overlap between the point clouds. Point cloud overlap, then pairwise consistent measurement set maximization (PCM) [26], are used to reject outlier loop closures. Inliers are integrated into the graph-based pose SLAM solution. We note the extension of the

PCM algorithm [15], but due to its additional complexity, we will not utilize it in our work.

C. Acoustic Modems

To communicate, AUVs generally use acoustic modems, which are low-bandwidth compared to in-air systems. One of the most well-known devices is the WHOI Micromodem [27], which can transfer a maximum of 5400 bits/s over a long-range. For shorter-range transmission (300 meters), higher bandwidth (62.5 kbits/s) units are available [28]. Moreover, recent research has demonstrated the feasibility of achieving even higher bandwidth in real-world conditions [29]. In our work, we consider bandwidth to be limited, and we will study the network utilization of our multi-robot system as a critical parameter. Further, we will take steps to minimize the transmission of large data structures.

III. PROBLEM DESCRIPTION

In this work we consider a team of robots N . Each robot $n \in N$ maintains its own SLAM solution, in its own reference frame $\mathcal{I}_n$. Each robot receives a set of observations $\mathbf{z}_{n,t}$ at a given pose $\mathbf{x}_{n,t}$ across discrete time steps t . Each robot pose is defined in the plane (fixed depth) as

$$\mathbf{x}_{n,t} = (x_r \ y_r \ \theta_r)^\top. \quad (1)$$

Each set of observations consists of sonar returns in spherical coordinates with ranges $r \in \mathbb{R}_+$, bearings $\theta \subseteq [-\pi, \pi)$, elevation angles $\phi \subseteq [-\pi, \pi)$, and associated intensity values $\gamma \in \mathbb{R}_+$. These sonar observations can be mapped into Cartesian coordinates by

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = r \begin{pmatrix} \cos \phi \cos \theta \\ \cos \phi \sin \theta \\ \sin \phi \end{pmatrix}. \quad (2)$$

Each robot moves through the environment according to the dynamics

$$\mathbf{x}_{n,t} = \mathbf{g}(\mathbf{u}_{n,t}, \mathbf{x}_{n,t-1}) + \sigma_{n,t}, \quad (3)$$

where $\mathbf{x}_{n,t-1}$ is the previous timestep's pose, $\mathbf{u}_{n,t}$ is the control command and $\sigma_{n,t}$ is process noise. The posterior probability over the time history of poses is defined as

$$p(\mathbf{x}_{n,1:t}, \mathbf{m}_n | \mathbf{z}_{n,1:t}, \mathbf{u}_{n,1:t}), \quad (4)$$

with map $\mathbf{m}_n$. While we consider a distributed multi-robot SLAM solution, there is no initial transform between each of the reference frames $\mathcal{I}_n$. This work aims to use each robot's observations to derive inter-robot loop closures and estimate team member trajectories in each robot's reference frame.

IV. ALGORITHM

This section will provide the technical details of our distributed multi-robot SLAM system. An overview of this system from the perspective of one robot is shown in Fig. 2. Each robot logs messages from team members, using a k-d tree to search for potential inter-robot loop closures. We then attempt registration and reject outliers. Inter-robot loop closures are added to the factor graph and broadcast to the rest of the team. The pose graph is optimized, and significant state changes are communicated to the rest of the team.




Fig. 2: **System diagram for a single robot.** Each robot maintains a message pool of received scene descriptors. Based on k-d tree search, registration is attempted and outliers rejected. Inlier inter-robot loop closures are added to the pose graph and sent to team members. Large SLAM corrections are also sent to team members.

A. Sonar image processing

At each pose, $\mathbf{x}_{n,t}$, the sonar observations $\mathbf{z}_{n,t}$ consist of a sonar image. The 2D image is populated with acoustic intensity values. However, not all the pixels in the image represent contact with structures in the environment. Our first step is to identify which pixels constitute a sonar contact and which do not. We use constant false alarm rate (CFAR) detection, [30] which is derived from radar processing and has supported our previous work [31].

Once contacts are identified, the pixels are mapped to meters as an in-plane point cloud. Note that while imaging sonar observes a 3D volume of water, the sensor does not return 3D information, only a 2D projection with ϕ as zero. The consequence is that our system is confined to fixed depth, in-plane pose estimation. The point cloud is subject to voxel down-sampling, with each voxel's output being the medoid of the contained points. Example point clouds are shown in Fig. 1c.

B. Point Cloud Compression

In this work, each point in a point cloud is a pair of 32-bit floating points (x,y), which may overwhelm the data link. For this reason, we consider a simple voxel-based compression algorithm. We take a point cloud, cast it onto a planar voxel grid, and only retain centers of occupied voxels. This compression is performed first by discretizing each point, $p_{x,y}$, to a designated resolution, $\Delta_{compression}$.

$$p_{i,j} \approx \frac{p_{x,y}}{\Delta_{compression}} \quad (5)$$

Note that the $p_{i,j}$ are discrete. Each cell in the voxel grid is populated according to Eq. (5). If a voxel contains a non-zero number of points, then that voxel's center is recorded for transmission. In our implementation, each voxel center is a pair of 8-bit unsigned integers, offering significant savings over a pair of 32-bit floating points. This compression, though, results in a loss of some geometric information. Our experiments will evaluate compression efficacy by characterizing SLAM error and network utilization in Section V.

C. Single robot SLAM

For multi-robot SLAM, first, each robot must estimate its state. We utilize the vehicles' onboard Doppler velocity log (DVL) and inertial measurement unit (IMU) with the uncompressed point clouds from Section IV-A. We formulate this as a graph-based pose SLAM problem and use the GTSAM [32] implementation of iSAM2 [33]. We use odometry factors f^0 from the vehicle dead reckoning system between sequential poses. Next, we add sequential scan matching (SSM) factors. f^{SSM} are derived from calling ICP [34] between sequential frames. Lastly, we consider non-sequential scan matching factors (i.e., intra-robot loop closures). f^{NSSM} factors are derived by calling ICP between non-sequential frames, outlier factors are rejected using a minimum required point cloud overlap, then PCM [26].

$$f(\Theta) = f^0(\Theta_0) \prod_i f_i^0(\Theta_i) \prod_j f_j^{SSM}(\Theta_j) \prod_q f_q^{NSSM}(\Theta_q)$$

Note that poses (keyframes) are instantiated if the dead reckoning system indicates a distance or rotation larger than a threshold compared to the previous pose. When each robot passes its state to the team members, it incrementally passes its optimized poses as keyframes are instantiated.

D. Distributed Multi-robot SLAM

Now that each robot has a system to estimate its state, we extend it to include the rest of the team's states in a distributed multi-robot SLAM solution. We use the existing pose graph to integrate inter-robot loop closures relating team members' trajectories. Note that each robot is responsible for its SLAM solution while passing the required information to build a multi-robot state estimate that includes each robot.

Each robot will incrementally share its pose estimates via the datalink as keyframes are added. We add any detected inter-robot factors, f^{IR} , as in Section V-F. Next, we will add the entire trajectory for each robot in the team if that robot has at least one f^{IR} . The trajectory will be added as a series of sequential factors between robot poses, denoted f^{PR} , *partner robot* factors. This completes the factor graph as shown in Fig. 3.

E. Inter-robot Loop Closure Search

We formulate the search for inter-robot loop closure candidates as a retrieval problem. Given a candidate and a database, we wish to search the database for possible matches. To do so, we build a scene descriptor using a range-based histogram similar to [35] for each scene. These histograms encode a scene according to the number of measurements at each range bin, making it rotationally invariant. Note that range bins are discrete, the size and range of which are provided a priori. Using the point cloud, each histogram bin is the number of points in its respective range bin.

Each robot encodes and communicates its scene descriptors to the rest of the team incrementally as the keyframes are instantiated. Note that when this message is shared, it also includes the pose estimate of the keyframe in the sender robot's reference frame. Additionally, each robot maintains




Fig. 3: **SLAM Factor Graph.** Robot poses x for three robots, α , β , δ are considered with several factors: sequential scan matching factors (SSM) in green, non-sequential scan matching factors (NSSM, intra-robot loop closures) shown in red, inter-robot (IR) loop closures in blue and partner robot (PR) factors in purple.

a k-d tree of its scene descriptors. As a robot receives scene descriptors from team members, it compares them to its own using a nearest neighbor search, with a designated maximum neighbor distance. Once nearest neighbors are identified, registration is attempted.

F. Registration

When performing registration between two scenes, we estimate a 3DOF rigid body transformation

$$T = (x_{\Delta} \ y_{\Delta})^T \quad (6)$$

$$R = \begin{pmatrix} \cos \theta_{\Delta} & -\sin \theta_{\Delta} \\ \sin \theta_{\Delta} & \cos \theta_{\Delta} \end{pmatrix}, \quad (7)$$

in order to derive inter-robot factors, f^{IR} . In the registration problem we minimize the squared Euclidean distance between two sets of points A and B , that have been associated according to nearest neighbor.

$$E(R, T) = \sum \|a_i - Rb_i + T\|^2. \quad (8)$$

Note we do not have any notion of an initial guess between robot reference frames, and therefore have no notion of initial guess when solving the registration problem. It is for this reason we utilize Go-ICP [36]. Go-ICP finds a globally optimal registration result, showing good performance under partial overlap and poor initial guess conditions. For Go-ICP initial alignment, we apply Euclidean mean subtraction to the point clouds and set R and T as identity and zero, respectively. Once Go-ICP completes, we call standard ICP to refine the transformation using the Go-ICP result as an initial guess.

G. Outlier Rejection

Once registration is complete, the result must be evaluated as legitimate or erroneous. In our implementation, we consider outlier rejection before and after registration. Before a registration call is made, we check that the point cloud has a minimum number of points, to avoid attempting registration with a cloud containing insufficient information

to provide a reasonable transformation estimate. Next, we consider the ratio between the point clouds; if the point ratio is too large or small, the scenes are not likely to be of the same content and therefore discarded. Next, we cast sonar points onto a coarse grid to enable rapid scene-to-scene comparison, which we term the *scene image*. When comparing loop closure candidates, the two relevant scene images are compared, and if the sum of absolute differences between them is high, the scenes are considered different and do not warrant registration. Registration is attempted if the point clouds have the minimum number of points, comply with our ratio requirements, and have a small-enough sum of absolute differences between scene images.

Once registration is completed, we evaluate the overlap between the two clouds using the resultant transform estimate. Overlap is computed by evaluating the percentage of points with the nearest neighbor inside 0.5 meters. If the overlap is sufficient, we pass the inter-robot loop closure to PCM [26] for geometric verification. Once a loop closure is validated by PCM, it is added to the pose graph and communicated to the rest of the robots in the team.

H. Communication Strategy

As noted in Section II-C, communication between robots is bandwidth limited. For this reason, we formulate our communications strategy to minimize the transmission of large data structures by first exchanging small ones. After the robot has completed its pose graph optimization at each step, it shares the newest scene histogram, as in Section IV-E, with the associated pose. When another member of the multi-robot system receives this message, the histogram is used to query a k-d tree. If any of the nearest neighbors comply with the maximum tree distance, then that point cloud is requested. Upon receipt of this cloud, registration is attempted as per Sections IV-F and IV-G. This model aims to eliminate the exchange of raw sensor data at each time step. Raw sensor data is only exchanged if scenes are close in feature space, potentially resulting in an inter-robot loop closure.

Note that we incrementally share pose estimates of robot keyframes as keyframes are instantiated. AUV state estimation, however, is prone to high SLAM drift and subsequent drift correction when loop closures are detected. These corrections mean that robots may need to share updated pose estimates with their team; otherwise, they will be operating with deprecated knowledge of other robots in the system. We handle this by setting a change threshold for a robot's state; if any pose changes significantly, that pose is re-sent to the team. The team then uses this new information to implement the f^{PR} .

V. EXPERIMENTS

A. Hardware Overview

In this work, we utilize our customized BlueROV2-Heavy as shown in Fig. 1a. This vehicle features depth and attitude control implemented with an onboard PixHawk. The sensor payload includes a Rowe SeaPilot DVL, VectorNav VN-100 MEMS IMU, Bar30 pressure sensor, and a twin sonar arrangement. The vertical sonar is a Blueprint Subsea Oculus

M1200d, and the horizontal sonar is an Oculus M750d. We use the horizontal sonar as the SLAM perceptual input; with a max range of 30 meters, the sensor has a 20-degree vertical aperture and a 130-degree horizontal field of view. In order to manage sensor data and our SLAM system, we utilize the robot operating system (ROS). The computer used to manage all instances of SLAM in our multi-robot system is equipped with an Intel i9-9880H 2.30GHz CPU.

In order to perform multi-robot experiments, we record several datasets and concurrently play them back on a single computer. We use $|N|$ instances of our system during playback to create a multi-robot system. Please note that we play back *real data* (sonar, DVL, IMU) and simulate acoustic communication between robots. To simulate communications between robots, we use the existing ROS message passing system to pass the relevant messages between instances of our system.

B. Metrics

In order to consider different parameterizations of our system, we use several metrics. Note that there is no ground truth in our data, nor is there enough publicly available sonar data with ground truth information to perform a study of this type. However, since the goal of our system is to provide awareness of the team relative to a single robot, we formulate the following. Each dataset has a stable single robot SLAM solution; this solution is transformed into the frame of each team member. The location of each team member in a common reference frame is determined using manual alignment. We then compare the single-agent SLAM outcome, transformed into each team member's reference frame, to the estimate each robot has built of its team members using our system. We denote this as inter-robot error, characterized with mean absolute error (MAE) and root mean squared error (RMSE).

Our second metric is network utilization, measured by monitoring the communication channel between robots and measuring total traffic. Note that we exclude network overhead and only report the cost of message contents, as network overhead may be hardware-specific. Network utilization is reported as a time series plot and considering the total usage divided by the mission duration. Note that we perform ten trials for each test case, and results for inter-robot error and network utilization are aggregated across all ten trials.

C. Multi-robot SLAM Ablation Study

To validate our system, we consider two environments: SUNY Maritime College in The Bronx, New York, shown in Fig. 1b, and the United States Merchant Marine Academy (USMMA) in Kings Point, New York. At SUNY Maritime, we use a system of two robots crossing paths from opposite sides of the environment, with a single opportunity for inter-robot loop closure. At USMMA, we consider a system of three robots traveling in similar directions, with several opportunities for inter-robot loop closure. At SUNY Maritime, the robots start on opposite sides of the environment, while at USMMA all robots start close together.

Case	Entire Trajectory				Poses With Inter-robot Loop Closures					
	Euclidean Dist. (m)		Theta (deg.)		Euclidean Dist. (m)		Theta (deg.)		Network Utilization (bits/second)	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	Average	Min/Max
SM 1	2.83	2.89	5.12	5.12	2.2	2.2	5.12	5.12	-	-
SM 2	2.36	2.39	3.47	3.47	2.06	2.06	3.47	3.47	-	-
SM 3	1.92	1.98	3.5	3.51	1.5	1.97	3.27	3.28	337.62	217.05 / 468.05
SM 4	2.63	2.68	4.74	4.74	2.03	2.04	4.75	4.75	161.07	134.93 / 208.55
SM 5	7.25	7.58	20.27	26.76	6.84	6.88	14.47	15.83	213.65	124.26 / 486.54
USMMA 1	2.28	2.66	3.24	3.89	1.22	1.44	1.78	2.38	-	-
USMMA 2	2.09	2.41	2.21	2.63	1.13	1.39	1.38	1.78	-	-
USMMA 3	1.17	1.29	1.93	2.27	0.66	0.73	1.27	1.47	3176.61	2556.64 / 5579.54
USMMA 4	1.44	1.62	2.62	3.07	0.82	0.97	1.71	2.09	1244.84	1050.19 / 1284.24
USMMA 5	3.09	3.59	17.76	22.66	3.26	3.31	11.51	13.33	2126.56	2079.58 / 2173.16

TABLE I: **Real world multi-robot SLAM results.** This table summarizes inter-robot error and network utilization for our five test cases. We analyze inter-robot error in two ways; the entire robot trajectory and only poses with inter-robot loop closures. We report the average network utilization and the min/max of the 10 trials for each test case. Note SM denotes SUNY Maritime.

To test the algorithmic components of our multi-robot SLAM architecture, we present an ablation study in which we progressively add components of our system until the system is complete. Specifically, we vary the contents of our outlier rejection system and specific components of the communication strategy. We note that there is no difference in system parameterization between datasets. For each dataset, we test each of the following cases:

- **Case 1:** min. no. pts., point ratio, and overlap conditions
- **Case 2:** min. no. pts., point ratio, overlap, and scene image conditions
- **Case 3:** min. no. pts., point ratio, overlap, scene image conditions, and PCM
- **Case 4:** min. no. pts., point ratio, overlap, scene image conditions, PCM, and point cloud compression (**The complete proposed system**)
- **Case 5:** min. no. pts., point ratio, overlap, scene image conditions and PCM without re-sending updated pose information

Our results are summarized in Table I, and qualitative examples are shown in Fig. 4. We also provide playback of our experiments (for case 4) in our video attachment. We require a minimum number of 75 points, a maximum point ratio of 2.0, a scene image max. of 0.8, minimum overlap of 55% and compression voxel size of 0.25 meters. When computing SLAM metrics, we consider Euclidean distance in meters and yaw (theta) in units of degrees. We break our SLAM metrics into two categories: (1) the entire trajectory and (2) only the poses with inter-robot loop closures. We also report network utilization in units of bits/second for each case. Note that cases 1 and 2 do not have network utilization reported. Their communication configuration does not differ from case 3, nor does their difference in outlier rejection change any form of communication logic.

When considering the results of our multi-robot SLAM ablation study, there are several important takeaways. Firstly there is added value in using a scene image as a method of outlier rejection (case 2) at both USMMA and SUNY Maritime in terms of error. Secondly, the addition of PCM (case 3) yields additional value in terms of inter-robot error. Next, when considering the use of point cloud compression

(case 4), there is a slight increase in inter-robot SLAM error in exchange for cutting network utilization in half compared to uncompressed clouds (case 3). Time series plots showing network utilization are given in Fig. 5. We note the reduction in network usage when using the compression method of Sec. IV-B, as well as the spike at the beginning of the USMMA mission, which is due to heavy point cloud exchange. Lastly, when removing the sending of significant updates in a robot’s state estimate to the rest of the team (case 5), a greatly deteriorated inter-robot SLAM result is observed with some savings in network utilization relative to case 3.

We note that our system runs faster than keyframes are added; an aggregate summary of runtimes from USMMA and SUNY Maritime is shown in Table II. We note that our registration time in Table II *includes* our outlier rejection system (without PCM, which is captured separately in the table). It is important to note that network utilization with three agents is reasonable when considering the limitations of commercial off-the-shelf hardware options. We want to underscore the significance of this result. This work is the first instance of multi-robot SLAM using real imaging sonar data to utilize indirect encounters, to the best of our knowledge. Not only does our system result in inter-robot loop closures, but these loop closures are also leveraged to build estimates of team members in the frame of each robot in a distributed manner. Further, it does so with an efficient exchange of information, as shown by the network utilization results. Lastly we note the success rate, i.e., how often inter-robot loop closures are successfully detected. For cases 1-4 in both data sets and USMMA case 5 we note a 100% success rate. For SUNY Maritime case 5, we note a 70% success rate, but given this is a data point meant to show the value of an otherwise included feature, we consider it no further.

D. Analysis of Perception Message Overhead

Some readers may note that using a computer vision paradigm is common when attempting to register a pair of keyframes with no initial guess. This general paradigm extracts salient point features, associates them, and then computes a transform with RANSAC or similar. We note the use of KAZE [37] and AKAZE [38] features on sonar imagery [39], [40] and the use of ORB [41] in RADAR




Fig. 4: **Example trajectories from USMMA.** This dataset includes three robots, α (red), β (blue) and δ (orange). Solid lines show trajectory estimates with dotted lines showing ground truth used to compute Table I. Note that there is only robot relative ground truth, so there is no true (dotted) line for a robot's own trajectory. We also note that estimated trajectories (solid) may be shorter than true (dotted) due to the varying shutdown times of robots. Stars show starting locations for robots. We avoid using γ since it refers to sonar intensity.

Algorithm	Mean (ms)	Standard Deviation (ms)
PCM	0.32	0.33
Kd-tree Search	0.88	4.49
Registration (w/ rejection)	258.88	265.43

TABLE II: **Runtime.** PCM is used for loop closure verification, k-d tree search is for scene comparison, and registration refers to the Go-ICP based registration method in Section IV-F, including the outlier rejection computations.

Method	Mean Count	Std. Dev.	Mean KBits
Scene Descriptor	1	1	.128
Point cloud - float32	138.69	142.94	9.14
Point cloud - compressed	138.69	142.94	2.28
KAZE	48.68	56.73	116.20
AKAZE	37.05	41.14	20.07
ORB	168.62	322.43	82.54

TABLE III: **Perception message overhead.** Mean count and std. dev. refer to the average number of points extracted with the given method. Mean Kbits are computed by taking the overhead of a single point and multiplying it by the mean number of points.

imagery. However, in the multi-robot case, that would require broadcasting feature points *and descriptors*. We provide a summary of some relevant feature extraction tools for sonar imagery in Table III. Here we consider the data-overhead per feature and the number of features per sonar image in our datasets. We use the OpenCV implementations of KAZE, AKAZE, and ORB to perform this comparison.

As shown in Table III, our scene descriptor is a data-efficient means of summarizing content. Moreover, even passing simple 32-bit floating-point-based point clouds is significantly less costly than KAZE, AKAZE, and ORB. Lastly, our compression method, while simple, shows a considerable reduction in data requirements, and as we have shown in Section V-C and Table I, only results in a minor increase in inter-robot error.

E. Multi-agent Mapping

While the purpose of our system is to provide each robot with knowledge of the states of the other agents in the system, it can also be used to derive merged maps. Consider a system of robots tasked with mapping an area in a




Fig. 5: **Network utilization.** A comparison of network utilization with and without the voxel compression method described in Section IV-B. The y-axis shows network utilization using a sliding window (window of 100) average. The x-axis shows time. A representative run is shown from the three-robot USMMA dataset at top, and from the two-robot SUNY Maritime dataset at bottom.

collaborative modality, a core task for a multi-robot system. This section showcases an example of a map that was merged offline, using the multi-robot SLAM results from Section V-D. We specifically consider the SUNY Maritime dataset, which has two robots. Using the multi-robot state estimate, we can transform the sonar point clouds in the frame of each robot, shown in Fig. 1c. We note the coverage increase with the use of a multi-robot system and the drift in our SLAM system. This drift is primarily due to the high drift rate of our low-cost MEMS IMU. We also note that the inter-robot drift is comparable to that of a single robot sweeping the entire workspace shown in Fig. 1c.

VI. CONCLUSIONS

In this work, we have proposed a system to find and integrate inter-robot loop closures in a distributed, graph-based pose SLAM architecture. We have demonstrated the real-time viability of our system, the accuracy of accepted inter-robot constraints, and the efficacy of our communications

system. When considering the potential shortcomings of our system, while we take steps to prevent perceptual aliasing from corrupting the SLAM solution, this effect can overwork the communications link. If scene descriptors are similar, we exchange point clouds, meaning that perceptual aliasing can result in the over-exchange of information that our outlier rejection system will cull. Future work will consider geometric verification in scene descriptors to ensure only useful perceptual data is exchanged. However, this is the first example of imaging sonar-based underwater multi-robot SLAM, to the best of our knowledge. There is much potential for future work, including achieving a wide variety of co-operative tasks, and multi-robot active SLAM. Furthermore, our system enables multi-agent autonomous operations in settings where GPS is unavailable or clock synchronization is impractical, which are relevant constraints in many real-world applications.

REFERENCES

- [1] L. Paull, M. Seto and J.J. Leonard, "Decentralized cooperative trajectory estimation for autonomous underwater vehicles," *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Syst.*, 184-191, 2014.
- [2] L. Paull, G. Huang, M. Seto and J.J. Leonard, "Communication-constrained multi-AUV cooperative SLAM," *Proc. IEEE Int. Conf. Robotics Automation*, 509-516, 2015.
- [3] M. Stojanovic and P.P. Beaujean, "Acoustic Communication," *Springer Handbook of Ocean Engineering*, M.R. Dhanak and N.I. Kiros, Eds., Springer, 2016.
- [4] E.M. Fischell, N.R. Rypkema and H. Schmidt, "Relative Autonomy and Navigation for Command and Control of Low-Cost Autonomous Underwater Vehicles," *IEEE Robotics Automation Lett.*, 4(2): 1800-1806, 2019.
- [5] Y. Yao, D. Xu and W. Yan, "Cooperative localization with communication delays for MAUVs," *Proc. IEEE Int. Conf. Intelligent Computing and Intelligent Syst.*, 244-249, 2009.
- [6] Y. Li, Y. Wang, W. Yu and X. Guan, "Multiple Autonomous Underwater Vehicle Cooperative Localization in Anchor-Free Environments," *IEEE J. Oceanic Engineering*, 44(4): 895-911, 2019.
- [7] T. Halsted and M. Schwager, "Distributed multi-robot localization from acoustic pulses using Euclidean distance geometry," *Proc. Int. Symp. Multi-Robot and Multi-Agent Systems*, 104-111, 2017.
- [8] J.G. Mangelson, R. Vasudevan and R.M. Eustice, "Communication Constrained Trajectory Alignment For Multi-Agent Inspection via Linear Programming," *Proc. IEEE/MTS OCEANS Conf.*, 2018.
- [9] A. Elibol, J. Kim, N. Gracias and R. Garcia, "Efficient image mosaicing for multi-robot visual underwater mapping," *Pattern Recognition Lett.* 46: 20-26, 2014.
- [10] M. Pfingsthorn, A. Birk and H. Bülow, "An efficient strategy for data exchange in multi-robot mapping under underwater communication constraints," *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Syst.*, 4886-4893, 2010.
- [11] F. Bonin-Font and A. Burguera, "Towards Multi-Robot Visual Graph-SLAM for Autonomous Marine Vehicles," *J. Marine Science and Engineering*, 8(6): 437, 2020.
- [12] M. Xanthidis et al., "Towards Multi-Robot Shipwreck Mapping," *IEEE Int. Conf. Robotics Automation (ICRA) Workshop on Underwater Active Perception*, 2021.
- [13] R. Campos, N. Gracias and P. Ridao, "Underwater Multi-Vehicle Trajectory Alignment and Mapping Using Acoustic and Optical Constraints," *Sensors*, 16(3): 387, 2016.
- [14] P. Ozog, N. Carlevaris-Bianco, A. Kim and R.M. Eustice, "Long-term Mapping Techniques for Ship Hull Inspection and Surveillance using an Autonomous Underwater Vehicle," *J. Field Robotics*, 33(3): 265-289, 2016.
- [15] H. Do, S. Hong and J. Kim, "Robust Loop Closure Method for Multi-Robot Map Fusion by Integration of Consistency and Data Similarity," *IEEE Robotics Automation Lett.*, 5(4): 5701-5708, 2020.
- [16] E. Galceran, S. Nagappa, M. Carreras, P. Ridao and A. Palomer, "Uncertainty-driven survey path planning for bathymetric mapping," *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Syst.*, 6006-6012, 2013.
- [17] S.B. Williams O. Pizarro, B. Foley, "Return to Antikythera: Multi-session SLAM Based AUV Mapping of a First Century B.C. Wreck Site," *Proc. Int. Conf. Field and Service Robotics*, 2015.
- [18] M. Bryson, M. Johnson-Roberson, O. Pizarro and S. Williams, "Automated registration for multi-year robotic surveys of marine benthic habitats," *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Syst.*, 3344-3349, 2013.
- [19] G. Kim and A. Kim, "Scan Context: Egocentric spatial descriptor for place recognition within 3D point cloud map," *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Syst.*, 4802-4809, 2018.
- [20] H. Wang, C. Wang and L. Xie, "Intensity Scan Context: Coding Intensity and Geometry Relations for Loop Closure Detection," *Proc. IEEE Int. Conf. Robotics Automation*, 2095-2101, 2020.
- [21] G. Kim, S. Choi and A. Kim, "Scan Context++: Structural Place Recognition Robust to Rotation and Lateral Variations in Urban Environments," *IEEE Trans. Robotics*, DOI: 10.1109/TRO.2021.3116424, 2021.
- [22] M.M. Santos, G.B. Zaffari, P. Ribeiro, P.L. J. Drews-Jr. and S.S.C. Botelho, "Underwater place recognition using forward-looking sonar images. A topological approach," *J. Field Robotics*, 36(2): 355-369, 2018.
- [23] P.O.C.S. Ribeiro et al., "Underwater Place Recognition in Unknown Environments with Triplet Based Acoustic Image Retrieval," *Proc. IEEE Int. Conf. Machine Learning and Applications*, 524-529, 2018.
- [24] M. Larsson, N. Bore and J. Folkesson, "Latent Space Metric Learning For Sidescan Sonar Place Recognition," *Proc. IEEE/OES Autonomous Underwater Vehicles Symp.*, 2020.
- [25] J. Wang, F. Chen, Y. Huang, J. McConnell, T. Shan and B. Englot, "Virtual Maps for Autonomous Exploration of Cluttered Underwater Environments," *IEEE Int. Conf. Robotics Automation (ICRA) Workshop on Underwater Active Perception*, 2021.
- [26] J.G. Mangelson, D. Dominic, R.M. Eustice and R. Vasudevan, "Pairwise Consistent Measurement Set Maximization for Robust Multi-Robot Map Merging," *Proc. IEEE Int. Conf. Robotics Automation*, 2916-2923, 2018.
- [27] Woods Hole Oceanographic Institute, "WHOI Micromodem," <https://acomms.whoi.edu/micro-modem/>
- [28] EvoLogics, "HS Devices," <https://evologics.de/acoustic-modem/hs>
- [29] E. Demirors, G. Sklivanitis, G.E. Santagati, T. Melodia and S.N. Batalama, "A High-Rate Software-Defined Underwater Acoustic Modem With Real-Time Adaptation Capabilities," *IEEE Access*, 6: 18602-18615, 2018.
- [30] M. Richards, *Fundamentals of Radar Signal Processing*, McGraw Hill, 2005.
- [31] J. McConnell and B. Englot, "Predictive 3D Sonar Mapping of Underwater Environments via Object-specific Bayesian Inference," *Proc. IEEE Int. Conf. Robotics Automation*, 6761-6767, 2021.
- [32] F. Dellaert, "Factor Graphs and GTSAM: A Hands-on Introduction," Georgia Institute of Technology, Technical Report No. GT-RIM-CPR-2012-002, 2012.
- [33] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental Smoothing and Mapping using the Bayes Tree," *Int. J. Robotics Res.*, 31(2): 216-235, 2012.
- [34] P.J. Besl and N.D. McKay, "A Method for Registration of 3-D Shapes," *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2): 239-256, 1992.
- [35] S. Thrun, W. Burgard, D. Fox, "Probabilistic Robotics" MIT Press, 2005.
- [36] J. Yang, H. Li and Y. Jia, "Go-ICP: Solving 3D Registration Efficiently and Globally Optimally," *Proc. IEEE Int. Conf. Computer Vision*, 1457-1464, 2013.
- [37] P. Alcantarilla, A. Bartoli and A. Davison, "KAZE Features," *Proc. European Conf. Computer Vision*, 214-227, 2012.
- [38] P. Alcantarilla, "Fast Explicit Diffusion for Accelerated Features in Nonlinear Scale Spaces," *Proc. British Machine Vision Conference*, 2013.
- [39] E. Westman, A. Hinduja and M. Kaess, "Feature-Based SLAM for Imaging Sonar with Under-Constrained Landmarks," *Proc. IEEE Int. Conf. Robotics Automation*, 3629-3636, 2018.
- [40] J. Wang, T. Shan, and B. Englot, "Underwater Terrain Reconstruction from Forward-Looking Sonar Imagery," *Proc. IEEE Int. Conf. Robotics Automation*, 3471-3477, 2019.
- [41] E. Rublee, V. Rabaud, K. Konolige and G. Bradski, "ORB: an efficient alternative to SIFT or SURF," *Proc. IEEE Int. Conf. Computer Vision*, 2564-2571, 2011.