

iMESA: Incremental Distributed Optimization for Collaborative Simultaneous Localization and Mapping

Daniel McGann and Michael Kaess

The Robotics Institute, Carnegie Mellon University, Pittsburgh PA, USA

Email: {danmcgann, kaess}@cmu.edu




Fig. 1: An illustration of the iMESA algorithm. When robots observe shared variables, they enforce equality using "biased priors" in their local factor-graph. Biased priors control constraint tightness with dual variables (λ) which are shown for robot A as purple in the robot's Bayes Tree. Over time, as communication is available, robots tighten equality constraints with dual gradient ascent to provide consistent solutions. Meanwhile, robots incorporate new measurements efficiently using the iSAM2 algorithm. Through this process, iMESA is able to accurately and efficiently solve incremental distributed C-SLAM problems.

Abstract—This paper introduces a novel incremental distributed back-end algorithm for Collaborative Simultaneous Localization and Mapping (C-SLAM). For real-world deployments, robotic teams require algorithms to compute a consistent state estimate accurately, within online runtime constraints, and with potentially limited communication. Existing centralized, decentralized, and distributed approaches to solving C-SLAM problems struggle to achieve all of these goals. To address this capability gap, we present Incremental Manifold Edge-based Separable ADMM (iMESA) a fully distributed C-SLAM back-end algorithm that can provide a multi-robot team with accurate state estimates in real-time with only sparse pair-wise communication between robots. Extensive evaluation on real and synthetic data demonstrates that iMESA is able to outperform comparable state-of-the-art C-SLAM back-ends.

I. INTRODUCTION

Collaborative Simultaneous Localization and Mapping (C-SLAM) is a fundamental capability for multi-robot teams [1]. A key component of the C-SLAM system is the back-end algorithm responsible for estimating the state of the robot team from their distributed, noisy measurements [2]. However, existing C-SLAM back-end algorithms struggle to handle the practical conditions experienced by multi-robot teams deployed in the real-world. During field deployments (e.g. search and rescue, forestry inspection, or scientific exploration [3–5]), multi-robot teams cannot assume the presence of communication infrastructure. Rather, teams can only assume an ad-hoc network that permits sparse communication between robots. Additionally, we expect multi-robot teams will need to re-plan frequently as new information is gathered. Such re-planning requires that the team have accurate and up-to-date state estimates. Existing C-SLAM back-end algorithms struggle to operate under these conditions either requiring

reliable network connectivity, needing impractical amounts of computation time, or struggling to achieve accurate results.

In this paper, we work to bridge the gap between existing C-SLAM back-end capabilities and the requirements for real-world multi-robot teams. We do so by developing an incremental distributed C-SLAM back-end that can tolerate sparse communication between robots and can provide accurate state estimates to the robot team in real-time.

The rest of the paper is structured as follows. First, we concretely define the incremental, communication limited, and computationally constrained C-SLAM problem we strive to address. Next, we review prior works and discuss their applicability to solve this task. We then provide a summary of the prior work MESA [6] before introducing our novel incrementalization of that algorithm – iMESA (Fig. 1). We conclude by rigorously evaluating the proposed algorithm on synthetic and real-world data to demonstrate that it achieves high quality results and is able to outperform existing algorithms. Finally, to aid future research, we make the source code for iMESA publicly available online¹.

II. PROBLEM DEFINITION

We seek to solve the generic C-SLAM problem in which a team of robots \mathcal{R} estimates variables of interest Θ using noisy measurements \mathcal{M} . In this multi-robot case, each robot takes a subset of the measurements and observes a subset of the total variables. Importantly, some measurements taken by robots will be inter-robot measurements. These measurements may come from direct observation of other robots or from a distributed place recognition system [7]. Inter-robot measurements enable the multi-robot team to collaborate. This

¹Source Code: <https://github.com/rpl-cmu/imesa>

can improve their individual state estimates and can ensure all robots' solutions remain in the same global frame. Due to inter-robot measurements, multiple robots may observe the same variable. Therefore, let $\Theta_i \subset \Theta$ denote the non-disjoint subset of variables observed by robot $i \in \mathcal{R}$. Additionally, let $\mathcal{M}_i \subset \mathcal{M}$ denote the robot's disjoint subset of measurements.

The de-facto standard for SLAM is to formulate the problem as a factor-graph and solve for the optimal variables via Maximum A-Posteriori (MAP) inference [2]:

$$\Theta_{MAP} = \arg \max_{\Theta \in \Omega} P(\Theta) \prod_{i \in \mathcal{R}} P(\mathcal{M}_i | \Theta_i) \quad (1)$$

where Ω is the product manifold constructed by the manifold of each variable in Θ . When we assume that each measurement $m \in \mathcal{M}$ is affected by Gaussian noise with covariance Σ_m this problem can be solved by Nonlinear Least Squares (NLS) optimization [8]:

$$\Theta_{MAP} = \arg \min_{\Theta \in \Omega} \sum_{i \in \mathcal{R}} \sum_{m \in \mathcal{M}_i} \|h(\Theta_i) - m\|_{\Sigma_m}^2 \quad (2)$$

where $h(\Theta_i)$ is the measurement prediction function that computes the expected measurement from an estimate of the state. However, we are not interested in just solving this optimization problem once. For practical robotic applications we need to solve this problem at every timestep, with the advantage that we have access to the previous solution. We refer to this problem as the incremental C-SLAM problem to distinguish it from the "batch" problem above [9]:

$$\Theta_{MAP}^t = \arg \min_{\Theta^t \in \Omega^t} \sum_{i \in \mathcal{R}} \sum_{m \in \mathcal{M}_i^t} \|h(\Theta_i^t) - m\|_{\Sigma_m}^2 \quad (3)$$

given: Θ_{MAP}^{t-1}

where (3) must be solved online to report results with minimal delay to downstream tasks like planning.

Remark 1 (Generic C-SLAM vs. PGO): *It is common in state estimation literature for C-SLAM back-ends to solve only Pose-Graph Optimization (PGO). PGO is a subset of C-SLAM in which each variable is a pose that lives on the $SE(N)$ manifold and all measurements are relative poses [10–15]. PGO is useful in many applications and provides structure to exploit in algorithm design. However, it limits map representations (e.g. landmarks) as well as measurement sources (e.g. bearing and range), the latter of which is particularly useful in specific domains (e.g. underwater [16] and space [17]). To develop the most useful C-SLAM back-end across all applications, we focus on the generic C-SLAM problem.*

In the incremental C-SLAM setting, information is distributed across robots in the team and robots must communicate to solve (3). In real-world applications, the team may not have access to existing communication infrastructure. As such, we can only reliably assume that robots can communicate over an ad-hoc network built from the robots themselves. Due to the scale of the environment, hardware constraints, motion of the robots, and sources of interference, this network will be bandwidth limited, time varying, and frequently disconnected.

Let us represent the communication network at a specific time t by a undirected graph $\mathcal{G}^t = (\mathcal{R}, \mathcal{E}^t)$ with nodes made up of the robots \mathcal{R} and edges \mathcal{E}^t defined by the currently

available connections between robots. Moreover, we assume that at each timestep any robots connected in \mathcal{G}^t are only able to communicate pairwise with one other robot during that timestep due to bandwidth and computation constraints.

This communication model is very restrictive and represents a worst-case scenario. However, it is representative of what multi-robot teams may encounter in remote wilderness, disaster areas, or on extraterrestrial bodies. Additionally, an algorithm that can operate effectively under this restrictive model can also operate in more optimistic scenarios. If the team has use of communication infrastructure, then \mathcal{G}^t is simply more densely connected at each timestep. If network bandwidth and robot hardware permits communication with multiple teammates, the algorithm can be run multiple times per timestep to use the additional resources and converge faster. Therefore, algorithms that can tolerate this model will be applicable even when communication is less limited.

III. RELATED WORK

Existing works struggle to provide effective solutions to incremental, communication-limited C-SLAM problems.

The simplest approach to the task is to use a centralized back-end. In this method, robots communicate their measurements to a central server that performs all computation and sends solutions back to the team. These approaches can make use of existing incremental solvers like iSAM2 or MR-iSAM2 to provide very accurate results [18, 19], and have been used successfully under a variety of conditions [20–30]. However, they require that \mathcal{G}^t is connected at each timestep and can support significant network traffic to return solutions to all team members. Additionally, a central server introduces a single point of failure into the system and centralization does not scale well as the number of robots increases [31].

Similar to centralized algorithms are decentralized back-ends. In these approaches each robot holds a copy of the global problem and independently computes the solution [32–37]. Like centralized methods, decentralized back-ends can use existing solvers to provide accurate results. Furthermore, decentralization removes any single point of failure from the system. However, decentralized back-ends require even greater bandwidth than centralized methods to send all measurements to every other robot. Additionally, decentralization necessitates that each robot performs expensive and redundant computation. These issues with bandwidth and computation can be mitigated, but not removed, by sparsifying the global problem [33] and sharing computation within robot clusters [37].

Another major class of C-SLAM back-ends are distributed algorithms. Choudhary et al. proposed distributed Multi-Block ADMM (MB-ADMM) for this task [38]. MB-ADMM can solve general C-SLAM problems accurately under some conditions. However, it is not guaranteed to converge for even linear problems and it requires synchronous communication over a connected network at each timestep [39].

Choudhary et al. also proposed the Distributed Gauss Sidel (DGS) algorithm for solving distributed batch PGO [10]. DGS optimizes a chordal relaxation of the PGO problem to provide

robust but approximate solutions [40]. Despite recovering only approximate solutions and requiring reliable communication, DGS has seen success in real-world systems [41–43].

To avoid the approximations made by DGS, Cristofalo et al. proposed the GeoD algorithm which optimizes the "Geodesic" PGO formulation enabling better accuracy [11]. GeoD treats each pose as an independent agent and optimizes the objective using gradient descent. GeoD has been shown to provide more accurate results than DGS at the cost of slower convergence.

Also targeting distributed batch PGO, Fan and Murphy proposed a Majorization-Minimization algorithm (MM-PGO) to optimize geodesic PGO [14, 15]. This approach performs optimization via gradient descent and can guarantee convergence to first order critical points. While this optimization has slow convergence it can be sped up using Nesterov’s method.

Going one step further than guaranteeing convergence to first order critical points, Tian et al. proposed an algorithm that optimizes the semi-definite relaxation of the batch PGO problem using distributed Riemannian gradient descent [12, 13]. The algorithm (DC2-PGO) and its asynchronous variant (ASAPP) can both compute accurate solutions to PGO and can even certify correctness of the solution under some conditions [44]. These algorithms have been used with success on real world systems [45–47]. However, they demonstrate slow convergence due to their use of first order optimization.

Recently, McGann et al. proposed MESA, a distributed optimization algorithm for generic batch C-SLAM problems based on Consensus Alternating Direction Method of Multipliers (C-ADMM) [6]. In addition to handling generic C-SLAM problems, MESA demonstrates significantly faster convergence and improved or comparable accuracy to other works.

However, MB-ADMM, GeoD, MM-PGO, DC2-PGO, ASAPP, and MESA are all insufficient to solve incremental, communication-limited C-SLAM problems. All of these distributed algorithms target batch C-SLAM problems. Naively, these methods could be applied to the incremental case (3) by solving each timestep as a batch problem. However, doing so cannot meet the requirements outlined in Sec. II. These methods require hundreds to thousands of iterations to converge, significantly exceeding the communication and computation requirements of our problem. Further, while some of these methods can tolerate asynchronous communication and thus communication dropout between algorithm iterations, they would all require a connected network during the current timestep, which is not guaranteed in real world scenarios.

Due to the challenges faced by distributed batch back-ends, some prior works have proposed distributed algorithms specifically targeting incremental C-SLAM. Cunningham et al. proposed Distributed Data Fusion Smoothing and Mapping (DDF-SAM and DDF-SAM2) which are custom distributed optimization algorithms designed around sharing marginals without double counting information [48, 49]. DDF-SAM2 is able to achieve online performance. However, it does not constrain robots to maintain equal linearization points which can significantly degrade solution consistency and accuracy.

Recently, Murai et al. proposed a new incremental dis-

tributed C-SLAM back-end based on Distributed Loopy Gaussian Belief Propagation (DLGBP) [50]. DLGBP has been shown to recover very accurate results under many problem conditions, but struggles to converge in some settings. Additionally, DLGBP achieves online performance with a windowed approach in which only the k most recent factors are updated. This windowing can significantly degrade performance, particularly when robots observe loop-closures.

Concurrent to this work, Matsuka and Chung also proposed a distributed incremental factor-graph optimizer based on C-ADMM called iDFGO [51]. While iDFGO was designed for optimizing convex multi-agent factor-graphs with linear constraints, it can be applied to nonlinear problems like C-SLAM. While similar in spirit to the algorithm we propose, there are a number properties of this concurrent work (e.g. required synchronous communication) that make it difficult to apply to incremental, communication limited C-SLAM problems. For details on these properties see Remark 10.

Within the context of the problem defined in Sec. II we can see that existing works do not provide a solution that is sufficient. Centralized and decentralized approaches cannot tolerate the desired communication model and scale poorly with the number of robots in the team. Similarly, batch distributed algorithms like MB-ADMM, DGS, GeoD, MM-PGO, DC2-PGO, ASAPP, and MESA cannot meet either the communication or computational requirements set forth. Even iDFGO, which shares a theoretical basis with our algorithm, cannot operate within bounds of our problem definition. While DDF-SAM2 and DLGBP are able to function within the bounds of our problem, their theoretical drawbacks often result in inaccurate solutions. To fill the gap left by prior works, we propose a novel incremental C-SLAM back-end – iMESA.

IV. MANIFOLD EDGE-BASED SEPARABLE ADMM (MESA)

In this section we review the batch distributed C-SLAM back-end algorithm MESA as we build off this prior work to design our proposed algorithm iMESA. For additional details on MESA, refer to the original publication [6]. MESA models batch C-SLAM (2) as a constrained optimization problem:

$$\begin{aligned} \Theta_{MAP} = \arg \min_{\substack{\Theta_i \in \Omega_i \forall i \in \mathcal{R} \\ Z \in \mathcal{Z}}} & \sum_{i \in \mathcal{R}} \sum_{m \in \mathcal{M}_i} \|h_m(\Theta_i) - m\|_{\Sigma_m}^2 \\ \text{s.t.} & \text{Log} \left(\theta_{s_i} \ominus z_{(i,j)_s} \right) = 0 \\ & \text{Log} \left(\theta_{s_j} \ominus z_{(j,i)_s} \right) = 0 \\ & \forall \theta_s \in \mathcal{S}_{(i,j)} \quad \forall (i, j) \in \mathcal{R} \times \mathcal{R} \end{aligned} \quad (4)$$

where Ω_i is the product manifold to which Θ_i belongs, $\mathcal{S}_{(i,j)}$ is the set of variables (e.g. robot poses or landmarks in the environment) shared between robots i and j , θ_{s_i} is the copy of shared variable s owned by robot i . For each shared variable, the problem is augmented with edge variables $z_{(i,j)_s}$ and $z_{(j,i)_s}$ which are held by robots i and j respectively. \mathcal{Z} is the set of all $z_{(i,j)_s}$ and \mathcal{Z} is the appropriate product manifold that is further constrained such that $z_{(i,j)_s} = z_{(j,i)_s}$.

Constraints are introduced, as inter-robot measurements result in some state variables that are shared between robots. To ensure the multi-robot teams produces a consistent solution, we must constrain the local solutions to these variables to be equal between all robots that share them. Additionally, variables in C-SLAM problems are typically robot poses on the Special Euclidean manifold. Thus, constraints may be nonlinear and are written using Lie group notation. A tutorial on Lie groups and this notation is provided by Solà et al. [52].

Remark 2 (Constraint Functions): *The authors of MESA originally proposed a general formulation of (4) using a generic constraint function. Given existing results, we opt to use Geodesic constraints [6, Table 1]. However, any valid constraint function could be used.*

MESA solves this constrained optimization problem via separable C-ADMM using an edge-based communication model and taking special considerations for on manifold variables. Concretely, this approach solves the Augmented Lagrangian for problem (4) by alternating between optimizing the primal variables Θ and performing gradient ascent on the dual variables Λ . The augmentation with edge variables $z_{(i,j)_s}$ permits the C-ADMM iterates to be solved in an asynchronous distributed fashion. The iterates are as follows:

$$\begin{aligned} \Theta_i^{k+1} = \arg \min_{\Theta_i \in \Omega_i} & \sum_{m \in \mathcal{M}_i} \left\| h_m(\Theta_i^k) - m \right\|_{\Sigma_m}^2 \\ & + \sum_{j \in \mathcal{R}} \sum_{s \in \mathcal{S}_{(i,j)}} \frac{\beta_{(i,j)}^k}{2} \left\| \text{Log} \left(\theta_{s_i}^k \Theta z_{(i,j)_s}^k \right) + \frac{\lambda_{(i,j)_s}^k}{\beta_{(i,j)}^k} \right\|^2 \end{aligned} \quad (5)$$

$$z_{(i,j)_s}^{k+1} = \text{SPLIT} \left(\theta_{s_i}^{k+1}, \theta_{s_j}^{k+1}, 0.5 \right) \quad (6)$$

$$\lambda_{(i,j)_s}^{k+1} = \lambda_{(i,j)_s}^k + \beta_{(i,j)}^k \text{Log} \left(\theta_{s_i}^{k+1} \Theta z_{(i,j)_s}^{k+1} \right) \quad (7)$$

$$\beta_{(i,j)}^{k+1} = \alpha \beta_{(i,j)}^k \quad (8)$$

where k is the iteration number, $\lambda_{(i,j)_s} \in \Lambda$ is the dual variable added for the constraint between θ_{s_i} and $z_{(i,j)_s}$, $\beta_{(i,j)}$ is the penalty parameter and learning rate for gradient ascent of the dual variables, α is a parameter to scale this penalty parameter, and $\text{SPLIT}(a, b, \gamma)$ interpolates a and b by factor $\gamma \in [0, 1]$ such that any linear component is interpolated linearly, and any rotational component is interpolated spherically.

The bulk of computation in MESA occurs in iterate (5). This optimization problem is the same as the classic SLAM NLS optimization problem (2) with additional terms:

$$\frac{\beta}{2} \left\| \text{Log}(\theta \Theta z) + \frac{\lambda}{\beta} \right\|^2 \quad (9)$$

These terms are referred to as "biased priors" [38]. Biased priors are the mechanism by which MESA constrains all robots' local solutions to be equal for any shared variables – thus ensuring a consistent state estimate for the team. Importantly, biased priors can be represented in standard "factor" form from $\|h(\Theta) - m\|_{\Sigma}^2$ meaning that (5) can be solved with existing SLAM optimization methods.

Iterates (6), (7), and (8) are computationally inexpensive and occur when a pair of robots communicate. The pair first communicates their local solution to any shared variables. Then,

each robot updates their edge variables as the interpolation of the local estimates. Edge variables are the values to which biased priors will pull a robot's local solution. Importantly, a pair of robots' edge variables for a shared variable are always equal ($z_{(i,j)_s} = z_{(j,i)_s}$). Therefore, each robot's biased priors will pull their solutions towards a consistent estimate, even as their local solutions change between communications. Next, the pair updates dual variables. The dual variables are an adaptive value by which MESA tightens the equality constraints for variables shared between robots. Finally, robots can optionally update β to change the step size for dual variable gradient ascent and further tighten constraints.

It has been shown that via this process, MESA converges to consistent solutions on generic batch C-SLAM problems [6]. Moreover, MESA is able to converge faster and more accurately than other prior works while permitting asynchronous communication between robots.

Remark 3 (MESA Convergence): *Prior work has demonstrated empirically that MESA converges for typical SLAM problems [6]. However, there is no formal proof of its convergence. Variants of C-ADMM has been proven to converge for non-convex problems [53, 54], on-manifold problems [55], problems with nonlinear constraints [56], asynchronous problems [57], and separable problems [58]. However, to the best of the author's knowledge, it has not been shown to converge for problems like C-SLAM that exhibit all of these traits and have coupled nonlinear constraints.*

V. INCREMENTAL MESA

In this section we introduce iMESA, an incrementalization of the MESA algorithm developed to solve the the incremental, communication-limited problem outlined in Sec. II. As discussed in Sec. III, batch C-SLAM back-end algorithms cannot be naively applied to the incremental problem. MESA, however, provides an algorithmic structure that permits us to design an effective incrementalized version.

A. iMESA Theoretical Overview

An iteration of MESA can be broadly broken down into two steps – a local optimization step (5) and a communication step (6), (7), and (8). In the local optimization step, robots compute their local estimate, pushing this solution to equal that of the other robots via biased priors. In the communication step, robots share relevant portions of their estimates and tighten shared variable constraints to provide better consistency. It is over many iterations that shared variable constraints are fully realized and a final consistent solution is reached.

The key idea of iMESA is to amortize this process (i.e. the tightening of constraints) over time as communication is available between robots. This amortization is possible with MESA for three reasons. Firstly, and most importantly, intermediate results provided by MESA (i.e. from iterations in which shared variable constraints are not tight) are still practically useful. The solution to (5) is effectively lower bounded by that found by a robot operating independently of the team and improved with even partially tight constraints. Secondly, MESA demonstrates fast convergence and therefore consistency and solution quality improve quickly even when the algorithm is amortized over time. Finally, MESA is fully

agnostic to the optimization process used to solve (5) as it only requires that we have a solution when communication occurs. This permits us to use existing incremental SLAM optimization methods to achieve online performance.

Remark 4 (Batch C-SLAM Amortization): *This structure of amortization is not compatible with all batch C-SLAM algorithms. Algorithms based on distributed gradient descent [10–13] do not provide useful results at intermediate steps of the algorithm. MESA, therefore, provides a unique basis for an incremental C-SLAM back-end.*

From this key idea, we develop a two stage algorithm consisting of a local update step in which a robot incrementally updates (5) as new measurements are observed and a communication step in which a pair of robots update their edge and dual variables to tighten shared variable constraints. These steps, combined with some practical bookkeeping, construct the iMESA algorithm and are described in detail below.

B. iMESA Implementation Details

To realize the idea proposed above, we must handle the practicalities of implementing the algorithm. Firstly, let each robot i maintain the following internal state:

<code>isam2</code>	An instance of an iSAM2 optimizer [18].
Θ	The local solution.
$S_j \forall j \in \mathcal{R}$	Sets variables shared with other robots.
$S'_j \forall j \in \mathcal{R}$	Like S_j but variables are known only to robot i due to limited communication.
$\Lambda = \{\lambda_{(i,j)_s}\}$	The set of all dual variables.
$Z = \{z_{(i,j)_s}\}$	The set of all edge variables.
$B = \{\beta_{(i,j)_s}\}$	The set of all penalty parameters.
C	Cache of biased priors to be added the optimizer.
\mathcal{K}	Cache of variables affected by recent communication.

Secondly, let each robot i be configured with the following hyper-parameter values:

β_{uninit}	$1e^{-4}$	Penalty parameter for uninitialized biased priors.
β_{init}	1.0	Penalty parameter for biased priors.

The purpose of these state variables and parameters are explored in detail in the following sub-sections.

1) *Bookkeeping*: Unlike batch settings, in incremental C-SLAM, the underlying problem is changing at every timestep as robots take new measurements. We must therefore, efficiently track information required to perform local updates and communicate information between robots.

To facilitate communication, each robot i tracks two sets for all other robots: $S_j \forall j \in \mathcal{R}$ which holds the set of variables shared between the pair (i, j) and $S'_j \subset S_j$ which holds the set of variables shared between the pair but are currently known only to robot i . Such variables exist as robot i may have detected a loop-closure to a variable in robot j 's factor-graph in the period of time since their last communication. When a robot i makes a new inter-robot measurement that observes a shared variable not already marked in S_j for the respective robot j we must add it to both S_j and S'_j .

For each variable s shared with another robot j each robot i must also maintain a corresponding dual variable $\lambda_{(i,j)_s}$, edge variable $z_{(i,j)_s}$, penalty parameter $\beta_{(i,j)_s}$, and biased prior (9).

Thus, when observing a new shared variable a robot must extend Λ , Z , and B as well as construct a corresponding biased prior that references these new values. The bookkeeping process is summarized in Alg. 1.

Algorithm 1 iMESA: Bookkeep (Local to robot i)

```

1: In: New shared variables  $S_{new}$  and their initial estimates  $\Phi$ 
2: for each variable  $s$  in  $S_{new}$  shared with robot  $j$  do
3:   Extend  $S_j$  and  $S'_j$  for new variable  $s$ 
4:   Extend  $\Lambda$  for  $j, s$  with value  $\mathbf{0}$ 
5:   Extend  $Z$  for  $j, s$  with value  $\Phi_s$ 
6:   Extend  $B$  for  $j, s$  with value  $\beta_{uninit}$ 
7:   Extend  $C$  with a biased prior on  $s$  referencing  $\lambda_{(i,j)_s}$ ,  $z_{(i,j)_s}$ , and  $\beta_{(i,j)_s}$ 
8: end for

```

Remark 5 (Bookkeeping Initial Values): *All dual variables are constructed with zero vector of appropriate dimension to match the MESA algorithm [6]. All edge variables are constructed with the local initial estimate as there has not yet been opportunity to communicate with the robot j to interpolate estimates. Additionally, since we know that $z_{(i,j)_s}$ is constructed to an incorrect value, we construct the penalty parameter with a very small value (β_{uninit}) so that the biased prior does not yet have an impact on the robot's local solution.*

2) *Update*: When new factors are observed, we must not only bookkeep the new information, but also update our local solution. The process of doing so is simple in the context of iMESA as prior work, namely that of Kaess et al., has already developed algorithms for efficient incremental updates to a SLAM problem using the Bayes Tree [18]. Encapsulating the iSAM2 algorithm, iMESA's update is summarized in Alg. 2

Algorithm 2 iMESA: Update (Local to robot i)

```

1: In: New factors  $\mathcal{F}$  and new initial estimates  $\Phi$ 
2:  $S_{new} \leftarrow$  any new shared variable observed in  $\mathcal{F}$ 
3: Bookkeep( $S_{new}, \Phi$ )
4:  $\Theta \leftarrow$  isam2.update( $\mathcal{F} \cup C, \Phi$ , reelim= $\mathcal{K}$ )
5: Reset  $C \leftarrow \emptyset$  and  $\mathcal{K} \leftarrow \emptyset$ 

```

▷ Alg. 1
▷ [18, Alg. 8]

Remark 6 (Biased Prior Cache): *During Alg. 2, we incorporate any biased priors in C that may have been cached during communication(s) since the last update. Biased priors are cached and included in the update here to avoid repeated work that would be required by updating iSAM2 explicitly after communication. This delays effects of communication from appearing in the local solution, but the delay is minor given we expect updates to be frequent. If effects from communication are required immediately, the iSAM2 update can be called both here and below in Alg. 3 at a computational cost.*

Remark 7 (iSAM2 Implementation Details): *By default iSAM2 will only re-linearize factors for variables which have changed by more than a set threshold. However, since edge and dual variables are not considered variables by iSAM2 (as they are not optimized by iSAM2), we must specifically indicate that biased priors who's edge and dual variables have been updated must be re-eliminated. We use the iSAM2 implementation provided in `gtsam` [59] that provides this functionality during iSAM2's update. Specifically, we mark for re-eliminated all variables \mathcal{K} for which corresponding dual and edge variables have been modified since the last update.*

3) *Communication*: When communication is available between two robots, MESA only requires that the pair share their current solution for all shared variables. This requires minimal bandwidth on the communication link between the pair. However, as noted above, the set of shared variables may have changed since the pair's last communication and each

robot may be only partially aware of these changes.

Therefore, iMESA makes use of a two stage communication process. First, the robots must update each-other on their set of shared variables by communicating S'_j . Upon receiving S'_i from the other robot, each must then update their bookkeeping. After this, both robots can share their local estimates for the, now jointly known, set of shared variables and update Λ , Z , and B . This communication process is summarized in Alg. 3.

Algorithm 3 iMESA: Communication (Robot i communicating with robot j)

```

1: Send  $S'_j$  to robot  $j$  and receive  $S'_i$  from robot  $j$            ▷ Comm. Step 1
2: Bookkeep( $S'_i$ ,  $\Theta$ ) and reset  $S'_j = \emptyset$                        ▷ Alg. 1
3: Send  $\theta_{s_i}$  and receive  $\theta_{s_j} \forall s \in S_j$                      ▷ Comm. Step 2
4: for  $s$  in  $S_j$  do
5:   Update  $z_{(i,j)_s} \in Z$  using Eq. 6
6:   Update  $\lambda_{(i,j)_s} \in \Lambda$  using Eq. 7
7:   If  $\beta_{(i,j)_s} == \beta_{\text{uninit}}$  then update  $B$  using  $\beta_{(i,j)_s} \leftarrow \beta_{\text{init}}$ 
8:   Extend  $\mathcal{K}$  with  $s$ 
9: end for

```

Remark 8 (Two Stage Communication): *A two stage communication process, is required for efficient communication by any incremental C-SLAM back-end, including DDF-SAM2 and DLGBP in scenarios with sparse communication. All incremental C-SLAM back-ends communicate some information pertaining to shared variables. Therefore, before communicating, robots must first agree upon a joint set of shared variables which may have changed since the last communication. Alternatively, robots could simply send their entire solution, but doing so would induce a much greater communication cost.*

4) *Penalty Parameters & Biased Priors*: A keen reader will notice that the penalty parameter update in Alg. 3 differs significantly from that used in MESA (8). Updating the penalty parameter along with updating dual variables can be used to tighten the shared variable constraints. However, unlike dual variables, changes to the penalty parameter are not adaptive and once tightened, they cannot be loosened. This can cause challenges in incremental settings when loop-closures are observed. If constraints are enforced with a large penalty parameter, it can prevent a robot’s local optimization process from updating its solution with that new information. Dual variables on the other hand, adapt in both magnitude and direction as local estimates change. Therefore, we use only dual variable updates to tighten constraints and use a constant penalty parameter of $\beta = 1$ once biased priors are initialized. Note that this requires we maintain a unique penalty parameter $\beta_{(i,j)_s}$ for each shared variable s .

However, while very general, the selection of $\beta = 1$ can cause stability issues in some edge cases. With an initial dual variable ($\lambda = \mathbf{0}$), a biased prior is effectively a standard prior factor with an isotropic noise model ($\Sigma = I$). Therefore, biased priors that affect pose variables provide a weak constraint on the rotational component. This can cause numerical stability issues when inter-robot measurements don’t observe rotations (e.g. bearing+range). The crux of the challenge is that pose translation (t) and rotation (r) components live in different domains with different scales $t \in [-D, D]$ and $r \in (-\pi, \pi]$ where D is the size of the team’s operational area. To improve stability, we also provide biased priors on poses with a noise model so they have similar effect between the two domains.

We specifically use a noise model with $\sigma_r = 0.1$ and $\sigma_t = 1$.

5) *iMESA Summary*: The iMESA algorithm consists entirely of running Alg. 2 and Alg. 3 on demand when measurements and communications are available. These algorithms are very efficient and allow iMESA to run online using sparse, bandwidth-limited communication as outlined in Sec. II. A visual depiction of iMESA in operation can be found in Fig. 1.

Remark 9 (iMESA Convergence): *Given that MESA does not come with any convergence guarantees (See Remark 3), iMESA also does not guarantee convergence. This, however, is standard among incremental distributed SLAM algorithms. As noted in Sec. III, DDF-SAM2 does not enforce consistency between shared variables and, in-turn, will not converge to a single solution [49]. Additionally, being an extension of Loopy Belief Propagation, DLGBP is not guaranteed to converge for even some convex problems, let alone non-convex C-SLAM problems [60].*

Remark 10 (Relationship to iDFGO): *Concurrent to this work Matsuka and Chung also proposed an incremental factor-graph optimization algorithm based on C-ADMM that can be applied to C-SLAM problems [51]. This algorithm (iDFGO) and iMESA share a similar spirit – to use C-ADMM for distributed factor-graph optimization and utilize existing incremental optimizers for efficient local optimization. However, these algorithms differ significantly.*

Unlike iMESA, iDFGO assumes that agents perform multiple iterations per-timestep where each iteration requires synchronized communication over a connected network. This design prevents iDFGO from meeting the communication requirements outlined in Sec. II. Additionally, iDFGO was primarily designed as a convex algorithm, and in-turn formulates constraints linearly. This is equivalent to using Chordal constraints which were found to result in significantly worse performance on C-SLAM problems as compared to Geodesic constraints that are used in iMESA [6]. Further, unlike iMESA, iDFGO does not handle online bookkeeping and instead assumes that agents have global knowledge of which variables are shared. Such global knowledge is impossible to obtain in incremental C-SLAM settings. Therefore, unlike iMESA, iDFGO’s properties make it difficult to apply to the incremental, communication limited C-SLAM problems that we address in this paper.

VI. EXPERIMENTS

In this section we evaluate the performance of iMESA on a variety of real and synthetic C-SLAM problems and demonstrate that it is able to achieve superior performance to state-of-the-art incremental, distributed C-SLAM back-ends.

A. Experimental Setup

1) *Prior Works & Baselines*: In all of our experiments we compare the performance of iMESA to the incremental distributed C-SLAM solvers DDF-SAM2 [49] and DLGBP [50]. DDF-SAM2 was implemented by the authors of this work and uses the Naive-Bayes approximation as suggested in the original paper. Implementation of DLGBP was provided by the original authors. We evaluate against standard DLGBP as well as Windowed DLGBP with window size $w = 30$.

In addition to these prior works we also validate iMESA against two baselines. The first is a centralized solver that aggregates all measurements into a single system and solves it incrementally using the iSAM2 optimizer [18]. The second is an independent solver in which all robots use iSAM2 to solve their local factor-graph without any inter-robot collaboration.




Fig. 2: Example ground-truth synthetic datasets: (a) 3D PGO dataset from Exp. VI-B, (b) 25 robot 2D dataset from Exp. VI-C, and (c) 5000 length 2D dataset from Exp. VI-D. Each color is the trajectory of a different robot.

2) *Synthetic Datasets:* We evaluate the above methods on synthetic and real-world datasets. Synthetic datasets are used to explore the effect of problem conditions on algorithm performance. Synthetic datasets are generated in both 2D and 3D by randomly sampling odometry from a categorical distribution with options of $1m$ forward motion as well as $\pm 90^\circ$ rotation around each available axis. Intra-robot loop closures are detected whenever a robot returns within $2m$ of a previous pose and added with probability $p = 0.4$. Intra-robot loop closures are always modeled as relative pose measurements. Inter-robot loop-closures are added every 5 timesteps to robots within $30m$ of each other. Inter-robot loop-closures can be modeled as relative pose measurements, bearing-range measurements, or range measurements. Visualizations of example synthetic datasets are provided in Fig. 2.

Remark 11 (Synthetic Data Noise Models): *All synthetic datasets use diagonal noise models with the following standard deviations. Priors: ($\sigma_r = 0.1^\circ, \sigma_t = 0.01m$), Odometry: ($\sigma_r = 1^\circ, \sigma_t = 0.05m$), Relative Pose Loop-Closures: ($\sigma_r = 1^\circ, \sigma_t = 0.1m$), Bearing+Range: ($\sigma_b = 1^\circ, \sigma_d = 0.1m$), Range: ($\sigma_d = 0.1m$). Where the standard deviations denote the measurement component to which they apply: r for rotation, t for translation, b for bearing, and d for distance/range.*

3) *Communication:* For all experiments we simulate communication connectivity between robots. Robots attempt to communicate at each timestep of the algorithm. Robots are only permitted a single pairwise communication per timestep to any robot within a communication range of $30m$. Additionally, we randomly drop 10% of communications to simulate instability in robots establishing a connection.

4) *Metrics:* We evaluate performance of all methods using Average Trajectory Error (ATE) to inspect the practical quality of results. However, we look to evaluate this performance not only for the final timestep, but, for intermediate timesteps during which a real robotic system will need accurate state estimates for downstream tasks. To measure this performance, we use an incremental variant of ATE [61, Eq. 3]. We compute iATE independently for each robot and sum all errors into two values, summarizing translational and rotational accuracy of the team. In all experiments iATE is reported in meters and radians for translational and rotational error respectively.

B. Measurement Models Experiment

In our first experiment we investigate how iMESA and prior works handle different problem formulations. Specifically we investigate 2D and 3D PGO, as well as 2D C-SLAM with bearing+range and range only inter-robot measurements. These different scenarios have different observability and

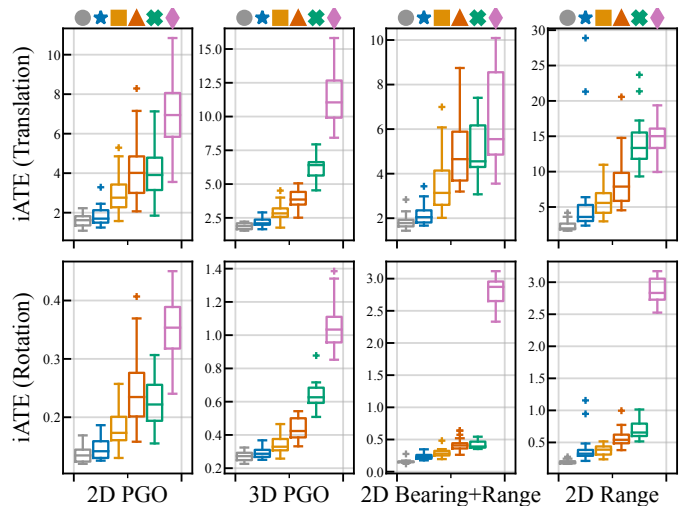


Fig. 3: iATE for our proposed algorithm iMESA (\star) along with centralized (\bullet) and independent (\diamond) baselines and prior works DDFSAM2 (\ast), DLGBP (\blacksquare), and Windowed DLGBP (\blacktriangle) on different C-SLAM formulations. Across problem formulations iMESA provides accurate results and outperforms state-of-the-art prior works.

nonlinearities, making them each an interesting optimization challenge. We run each method on 20 random datasets with 5 robots, each traversing a trajectory 1000 poses long.

We summarize the iATE achieved on each trial via a box and whisker plot in Fig. 3. There, we see that iMESA is able to outperform all prior works across problem scenarios and recover the closest solution to that of a centralized optimizer. Further, we can see Windowed DLGBP performs significantly worse than its non-windowed variant due to the presence of loop-closures in the datasets. Finally, it is notable that in the range scenario, iMESA, Windowed DLGBP, and DDF-SAM2 reported outlier performance on a small number of trials, likely due to the significant nonlinearity of range measurements.

C. Problem Scale Experiment

Next we evaluate how the methods perform as the size of the team grows. For each team size we generate 20 random 2D PGO datasets with robots traversing trajectories of 500 poses.

We summarize results from these trials in Fig. 4 in which we can observe that again, iMESA outperforms prior works. We can also see that the relative performance of all distributed methods is consistent across scales. We note that the error increases as team size increases as we sum errors from all robots (See. VI-A4). We can gain additional insight by looking at the trend in runtime across team size. In Fig. 5, we plot the total runtime averaged over all trials. Here, we can see that unlike the distributed methods, the centralized baseline computational requirements scale very poorly with team size. These results add evidence to the claim that centralized approaches are difficult to employ with large multi-robot teams [31]. The distributed methods, on the other hand, all scale very well with the size of the team.

D. Long-term Operation Experiment

In our next synthetic experiment we investigate how the methods handle long-term operation. For this experiment we




Fig. 4: iATE for our proposed method iMESA (★) along with centralized (●) and independent (◆) baselines and prior works DDFSAM2 (✱), DLGBP (■), and Windowed DLGBP (▲) as robot team size increases. Our proposed algorithm iMESA outperforms prior works and all methods provide consistent performance as team size grows.




Fig. 5: Total runtime for our proposed method iMESA (★) along with centralized (●) and independent (◆) baselines and prior works DDFSAM2 (✱), DLGBP (■), and Windowed DLGBP (▲) as team size increases. All distributed methods scale well as team size increases, however, the centralized solver sees significantly increased runtime as team size grows.

evaluate on 20 random 2D PGO datasets consisting of 5 robots traversing trajectories of increasing length.

Long-term operation performance is summarized in Fig. 6. In this figure we can see the continued trend that iMESA outperforms existing distributed C-SLAM back-ends and achieves results closest to a centralized solver. Interestingly, we can see in this experiment that DLGBP and its windowed variant perform worse as the trajectory length grows. We hypothesize that DLGBP struggles to properly incorporate very long term loop-closure information even without windowing. However, the exact cause of this degradation in performance was not studied in depth and deserves focus in future work.

We also use the results from this experiment to evaluate the feasibility of online performance for all methods. Using results from the 10k length trials, we compute the iteration run-time averaged across any parallel computation (i.e. across robots for all methods except the centralized approach). Iteration runtime includes all computation required by updates and during any communication performed at that iteration. However, it does not include network transmission time that would be required during actual communication. We further average these run-times across all 20 trials and plot the results in Fig. 7.

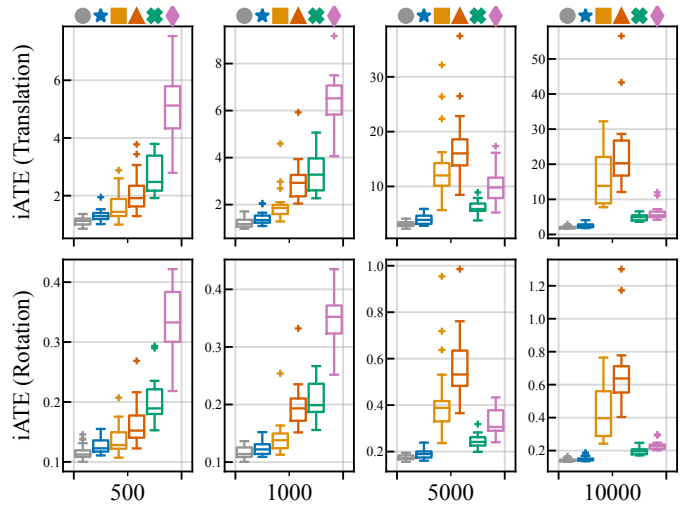


Fig. 6: iATE for our proposed method iMESA (★) along with centralized (●) and independent (◆) baselines and prior works DDFSAM2 (✱), DLGBP (■), and Windowed DLGBP (▲) as operational time increases. Our algorithm iMESA outperforms prior works across all lengths. Interestingly, DLGBP and its windowed variant see degraded performance on longer problems.




Fig. 7: Iteration runtime for our proposed method iMESA (★) along with centralized (●) and independent (◆) baselines and prior works DDFSAM2 (✱), DLGBP (■), and Windowed DLGBP (▲) as the factor-graph size increases. Our algorithm iMESA achieves very efficient performance requiring only marginally more computation time per iteration than locally running iSAM2 (independent). DLGBP iteration times grow significantly indicating that it may be difficult to apply to real-world applications. Note: the y-axis scale is linear up to 0.1 second and log scale above that threshold.

These iteration runtime results demonstrate that iMESA is very efficient. Even with 10k poses, the average iteration time remains less than 0.04 seconds requiring little more time than standard iSAM2 on problems of the same size (i.e. independent approach). This indicates that iMESA will be able to achieve real-time performance even in long-duration applications. To contextualize this result, we also run MESA on a single 10k length dataset. To optimize only the final iteration, MESA required over 800 pairwise communications and over 2 minutes of computation on each robot. This highlights the importance of incrementalization and the difficulties one would face if naively applying batch algorithms to incremental problems. These results also show that, as expected, Windowed DLGBP is the fastest. We note that its iteration runtime does increase despite being theoretically constant as larger data structures require more time to update. Further we can see that DLGBP is by far the most expensive, with final iteration times well over 1 second. While DLGBP may be feasible for small incremental C-SLAM problems, it will likely not achieve real-time performance for larger problems.




Fig. 8: iATE for our proposed method iMESA with different amounts of maximum communication delay [0, 4, 8, 12] along with centralized (●) and independent (◆) baselines. Even with increasing delays to some information, iMESA is able to perform consistently and provide accurate solutions.

E. Delayed Communication Experiment

In all experiments we simulate asynchronous communication between robots by limiting communication range and restricting robots to pairwise communication which causes robots to always be working with some outdated information. One aspect of communication networks that this does not explore is delays that may occur when communicating over networking infrastructure. In this experiment we explore the effect of delays on the performance of iMESA. We induce gaps in communication by restricting robots to communication only every other timestep and increase communication failure rate to 0.4. Further, robots communicate their oldest complete information up to a maximum delay of k timesteps. This experiment is run on 20 random PGO datasets consisting of 8 robots traversing trajectories 500 poses long.

We summarize the results that iMESA achieves under increasing maximum delay and summarize the iATE achieved in Fig 8. In this figure, we can see that iMESA maintains consistent performance across scales of maximum delay. This indicates that iMESA will be able to tolerate the networking delays we may encounter in real-world deployments. We also note that iMESA performance gap to the centralized solution is slightly larger than in other experiments due to the slower communication rate and increased communication dropout.

F. Real-Data Multi-Robot Experiment

In our final experiment we explore how incremental distributed C-SLAM methods perform on real-world data. We use the open-source Nebula Autonomy multi-robot datasets. These datasets consist of LiDAR data and 3D pose-graphs generated by LAMP 2.0 in subterranean tunnel environments [27]. We replay these pose-graphs incrementally for the algorithms investigated in this work, simulating communication using the method described in Sec. VI-A3 with a larger communication range of 100m. Additionally, due to the real-world nature of these datasets some loop-closures in the pose-graphs are outliers. In this work we do not focus on handling outliers and instead preprocess datasets to remove outlier measurements.

Accuracy results for all methods on the four datasets are summarized in Table I. In this table we can see that iMESA performs consistently with results comparable to the centralized solver. We can also see this qualitatively in iMESA’s solution to the *finals* dataset in Fig. 9. Interestingly, we found that DLGBP diverged on the *urban* and *finals* dataset. This was our only observation of this behavior, but indicates DLGBP is sensitive to problem conditions.




Fig. 9: Trajectory (left) and corresponding map (right) computed by iMESA on the *finals* dataset. Different colors denote different robots and lighter colors on trajectory plot depict groundtruth.

TABLE I: iATE (Translation, Rotation) on the Nebula Multi-Robot Datasets

Algorithm	Dataset							
	tunnel		urban		finals		KU	
Centralized	1.48	0.093	1.20	0.076	0.53	0.096	3.05	0.118
iMESA	1.46	0.091	1.20	0.071	0.65	0.094	3.14	0.117
DLGBP	1.91	0.093	–	–	–	–	3.37	0.120
WDLGBP	1.71	0.086	1.23	0.075	0.85	0.142	4.85	0.140
DDFSAM2	1.59	0.095	1.63	0.094	0.65	0.094	3.00	0.117
Independent	1.64	0.102	1.38	0.083	0.71	0.099	3.71	0.124

Finally, we want to note that some distributed methods appear to outperform the centralized solver in this experiment. This occurs as measurements are derived from real-data, and therefore the optimum of (2) does not necessarily correspond to the groundtruth trajectory. As such, slight deviations from this optimum that distributed methods exhibit can marginally improve the trajectory error with respect to the groundtruth.

VII. CONCLUSION AND FUTURE WORK

In this paper we present iMESA, a novel incremental distributed C-SLAM algorithm. iMESA is able to achieve real-time performance and accurate results with only sparse communication between robots. We demonstrate that iMESA: can handle generic C-SLAM problem formulations, enabling it to be used across application domains, can scale to large robot teams making it useful as future applications increase team sizes, can handle long-term operations and therefore will be applicable to field robotic applications, and outperforms existing incremental distributed C-SLAM back-ends.

However, iMESA requires additional work to see success in the real-world. iMESA, like other NLS based C-SLAM solvers, is sensitive to outlier measurements. Future work should investigate methods to make iMESA robust to both local and inter-robot outlier measurements. Additionally, to enforce invariants, iMESA requires that the two-stage communication process (Alg. 3) be fully completed if started. Future work should look to add algorithmic robustness for the case that communication drops out between robots while this process is occurring. Finally, while iMESA has empirically demonstrated quality performance across C-SLAM problem scenarios, future work should look to provide theoretical guarantees on its convergence properties and investigate iMESA’s ability to run on-board multi-robot teams operating in the field.

In addition to future work on the iMESA algorithm, we will also note the need for additional investigation into DLGBP based methods. These methods remain promising, but this

work found multiple instances where they struggled to converge or diverged entirely. Future work to explain and resolve this behavior would greatly benefit the C-SLAM community.

VIII. ACKNOWLEDGEMENTS

This work was partially supported by NASA award 80NSSC22PA952 and the NSF Graduate Research Fellowship Program. The authors would also like to thank Riku Murai for providing their implementation of DLGBP and their insightful discussion regarding Loopy Belief Propagation.

REFERENCES

- [1] P. Lajoie, B. Ramtoula, F. Wu, and G. Beltrame, "Towards collaborative simultaneous localization and mapping: a survey of the current research landscape," *Journal of Field Robotics*, vol. 2, no. 1, pp. 971–1000, 2022.
- [2] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. on Robotics (TRO)*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [3] D. Drew, "Multi-agent systems for search and rescue applications," *Current Robotics Reports*, vol. 2, pp. 189–200, 2021.
- [4] C. Yuan, Y. Zhang, and Z. Liu, "A survey on technologies for automatic forest fire monitoring, detection, and fighting using unmanned aerial vehicles and remote sensing techniques," *Canadian Journal of Forest Research*, vol. 45, no. 7, pp. 783–792, 2015.
- [5] NASA Jet Propulsion Laboratory, "Cooperative Autonomous Distributed Robotic Exploration (CADRE)." <https://www.jpl.nasa.gov/missions/cadre>.
- [6] D. McGann, K. Lassak, and M. Kaess, "Asynchronous distributed smoothing and mapping via on-manifold consensus ADMM," arXiv:2310.12320 [cs.RO], arXiv preprint, 2023.
- [7] Y. Tian, K. Khosoussi, and J. How, "Resource-aware algorithms for distributed loop closure detection with provable performance guarantees," in *Algorithmic Foundations of Robotics XIII: Proceedings of the 13th Workshop on the Algorithmic Foundations of Robotics*, pp. 422–438, Springer, 2020.
- [8] F. Dellaert and M. Kaess, "Factor graphs for robot perception," *Foundations and Trends in Robotics (FNT)*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [9] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Fast incremental smoothing and mapping with efficient data association," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, (Rome, Italy), pp. 1670–1677, Apr. 2007.
- [10] S. Choudhary, L. Carlone, C. Nieto, J. Rogers, H. Christensen, and F. Dellaert, "Distributed mapping with privacy and communication constraints: Lightweight algorithms and object-based models," *Intl. J. of Robotics Research (IJRR)*, vol. 36, no. 12, pp. 1286–1311, 2017.
- [11] E. Cristofalo, E. Montijano, and M. Schwager, "GeoD: Consensus-based geodesic distributed pose graph optimization," arXiv:2010.00156 [cs.RO], arXiv preprint, 2020.
- [12] Y. Tian, A. Koppel, A. Bedi, and J. How, "Asynchronous and parallel distributed pose graph optimization," *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 4, pp. 5819–5826, 2020.
- [13] Y. Tian, K. Khosoussi, D. Rosen, and J. How, "Distributed certifiably correct pose-graph optimization," *IEEE Trans. on Robotics (TRO)*, vol. 37, no. 6, pp. 2137–2156, 2021.
- [14] T. Fan and T. Murphey, "Majorization minimization methods for distributed pose graph optimization with convergence guarantees," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 5058–5065, 2020.
- [15] T. Fan and T. Murphey, "Majorization minimization methods for distributed pose graph optimization," *IEEE Trans. on Robotics (TRO)*, vol. 40, pp. 22–42, 2024.
- [16] L. Paull, S. Saeedi, M. Seto, and H. Li, "AUV navigation and localization: A review," *IEEE J. of Oceanic Engineering (JOE)*, vol. 39, no. 1, pp. 131–149, 2014.
- [17] E. Boroson, R. Hewitt, N. Ayanian, and J. de la Croix, "Inter-robot range measurements in pose graph optimization," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 4806–4813, 2020.
- [18] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *Intl. J. of Robotics Research (IJRR)*, vol. 31, no. 2, pp. 216–235, 2012.
- [19] Y. Zhang, M. Hsiao, J. Dong, J. Engel, and F. Dellaert, "MR-iSAM2: Incremental smoothing and mapping with multi-root bayes tree for multi-robot slam," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 8671–8678, 2021.
- [20] T. Bailey, M. Bryson, H. Mu, J. Vial, L. McCalman, and H. Durrant-Whyte, "Decentralised cooperative localisation for heterogeneous teams of mobile robots," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 2859–2865, 2011.
- [21] R. Dubé, A. Gawel, H. Sommer, J. Nieto, R. Siegwart, and C. Cadena, "An online multi-robot SLAM system for 3d LiDARs," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 1004–1011, 2017.
- [22] P. Schmuck and M. Chli, "Multi-UAV collaborative monocular SLAM," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 3863–3870, 2017.
- [23] M. Karrer, P. Schmuck, and M. Chli, "CVI-SLAM—collaborative visual-inertial SLAM," *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 4, pp. 2762–2769, 2018.
- [24] P. Schmuck and M. Chli, "CCM-SLAM: Robust and efficient centralized collaborative monocular simultaneous localization and mapping for robotic teams," *Journal of Field Robotics*, vol. 36, no. 4, pp. 763–781, 2019.
- [25] P. Schmuck, T. Ziegler, M. Karrer, J. Perraudin, and M. Chli, "COVINS: Visual-inertial SLAM for centralized collaboration," in *Proc. IEEE Intl. Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pp. 171–176, 2021.
- [26] Y. Zhang, M. Hsiao, Y. Zhao, J. Dong, and J. Engel, "Distributed client-server optimization for SLAM with limited on-device resources," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 5336–5342, 2021.
- [27] Y. Chang, K. Ebadi, C. Denniston, M. Ginting, A. Rosinol, A. Reinke, M. Palieri, J. Shi, A. Chatterjee, B. Morrell, A. Agha-mohammadi, and L. Carlone, "LAMP 2.0: A robust multi-robot SLAM system for operation in challenging large-scale underground environments," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 9175–9182, 2022.
- [28] T. Zhang, L. Zhang, Y. Chen, and Y. Zhou, "CVIDS: A collaborative localization and dense mapping framework for multi-agent based visual-inertial SLAM," *IEEE Transactions on Image Processing*, vol. 31, pp. 6562–6576, 2022.
- [29] Y. Chang, N. Hughes, A. Ray, and L. Carlone, "Hydra-multi: Collaborative online construction of 3d scene graphs with multi-robot teams," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 10995–11002, 2023.
- [30] A. Cramariuc, L. Bernreiter, F. Tschopp, M. Fehr, V. Reijngwart, J. Nieto, R. Siegwart, and C. Cadena, "maplab 2.0 – a modular and multi-modal mapping framework," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 2, pp. 520–527, 2023.
- [31] K. Ebadi, L. Bernreiter, H. Biggie, G. Catt, Y. Chang, A. Chatterjee, C. Denniston, S. Deschênes, K. Harlow, S. Khattak, L. Nogueira, M. Palieri, P. Petráček, M. Petrík, A. Reinke, V. Krátký, S. Zhao, A. Agha-mohammadi, K. Alexis, C. Heckman, K. Khosoussi, N. Kottege, B. Morrell, M. Hutter, F. Pauling, F. Pomerleau, M. Saska, S. Scherer, R. Siegwart, J. Williams, and L. Carlone, "Present and future of SLAM in extreme environments: The DARPA SubT challenge," *IEEE Trans. on Robotics (TRO)*, pp. 1–20, 2023.
- [32] L. Paull, M. Seto, and J. Leonard, "Decentralized cooperative trajectory estimation for autonomous underwater vehicles," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 184–191, 2014.
- [33] M. Schuster, C. Brand, H. Hirschmüller, M. Suppa, and M. Beetz, "Multi-robot 6d graph SLAM connecting decoupled local reference filters," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 5093–5100, 2015.
- [34] M. Schuster, K. Schmid, C. Brand, and M. Beetz, "Distributed stereo vision-based 6d localization and mapping for multi-robot teams," *Journal of Field Robotics*, vol. 36, no. 2, pp. 305–332, 2019.
- [35] R. Dubois, A. Eudes, J. Moras, and V. Frémont, "Dense decentralized multi-robot SLAM based on locally consistent TSDF submaps," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 4862–4869, 2020.
- [36] Y. Huang, T. Shan, F. Chen, and B. Englot, "DiSCo-SLAM: Distributed

- scan context-enabled multi-robot LiDAR SLAM with two-stage global-local graph optimization,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 2, pp. 1150–1157, 2022.
- [37] P. Lajoie and G. Beltrame, “Swarm-SLAM: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 9, no. 1, pp. 475–482, 2024.
- [38] S. Choudhary, L. Carlone, H. Christensen, and F. Dellaert, “Exactly sparse memory efficient SLAM using the multi-block alternating direction method of multipliers,” in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, (Hamburg, DE), pp. 1349–1356, Oct. 2015.
- [39] C. Chen, B. He, Y. Ye, and X. Yuan, “The direct extension of ADMM for multi-block convex minimization problems is not necessarily convergent,” *Mathematical Programming*, vol. 155, no. 1-2, pp. 57–79, 2016.
- [40] L. Carlone, R. Tron, K. Daniilidis, and F. Dellaert, “Initialization techniques for 3D SLAM: A survey on rotation estimation and its use in pose graph optimization,” in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, (Seattle, USA), pp. 4597–4604, 2015.
- [41] T. Cieslewski, S. Choudhary, and D. Scaramuzza, “Data-efficient decentralized visual SLAM,” in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 2466–2473, 2018.
- [42] P. Lajoie, B. Ramtoula, Y. Chang, L. Carlone, and G. Beltrame, “DOOR-SLAM: Distributed, online, and outlier resilient SLAM for robotic teams,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 5, no. 2, pp. 1656–1663, 2020.
- [43] S. Zhong, Y. Qi, Z. Chen, J. Wu, H. Chen, and M. Liu, “DCL-SLAM: A distributed collaborative LiDAR SLAM framework for a robotic swarm,” *IEEE Sensors Journal*, 2023.
- [44] C. Holmes, F. Dümbgen, and T. Barfoot, “On semidefinite relaxations for matrix-weighted state-estimation problems in robotics,” arXiv:2308.07275 [cs.RO], arXiv preprint, 2023.
- [45] Y. Chang, Y. Tian, J. How, and L. Carlone, “Kimera-multi: a system for distributed multi-robot metric-semantic simultaneous localization and mapping,” in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, (Xi’an, CN), pp. 11210–11218, 2021.
- [46] Y. Tian, Y. Chang, F. Herrera-Arias, C. Nieto-Granda, J. How, and L. Carlone, “Kimera-multi: Robust, distributed, dense metric-semantic SLAM for multi-robot systems,” *IEEE Trans. on Robotics (TRO)*, vol. 38, no. 4, pp. 2022–2038, 2022.
- [47] Y. Tian, Y. Chang, L. Quang, A. Schang, C. Nieto-Granda, J. How, and L. Carlone, “Resilient and distributed multi-robot visual SLAM: Datasets, experiments, and lessons learned,” in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 11027–11034, 2023.
- [48] A. Cunningham, M. Paluri, and F. Dellaert, “DDF-SAM: Fully distributed slam using constrained factor graphs,” in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, (Taipei, TW), pp. 3025–3030, Oct. 2010.
- [49] A. Cunningham, V. Indelman, and F. Dellaert, “DDF-SAM 2.0: Consistent distributed smoothing and mapping,” in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, (Karlsruhe, DE), pp. 5220–5227, May 2013.
- [50] R. Murai, J. Ortiz, S. Saeedi, P. Kelly, and A. J. Davison, “A robot web for distributed many-device localization,” *IEEE Trans. on Robotics (TRO)*, vol. 40, pp. 121–138, 2024.
- [51] K. Matsuka and S. Chung, “Localized and incremental probabilistic inference for large-scale networked dynamical systems,” *IEEE Trans. on Robotics (TRO)*, vol. 39, no. 5, pp. 3516–3535, 2023.
- [52] J. Solà, J. Derafay, and D. Atchuthan, “A micro lie theory for state estimation in robotics,” arXiv:1812.01537 [cs.RO], arXiv preprint, 2021.
- [53] M. Hong, Z. Luo, and M. Razaviyayn, “Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems,” *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 337–364, 2016.
- [54] Y. Wang, W. Yin, and J. Zeng, “Global convergence of ADMM in nonconvex nonsmooth optimization,” *Journal of Scientific Computing*, vol. 78, pp. 29–63, Jan. 2019.
- [55] J. Li, S. Ma, and T. Srivastava, “A riemannian ADMM,” arXiv:2211.02163 [math.OC], arXiv preprint, 2023.
- [56] K. Sun and X. Sun, “A two-level distributed algorithm for nonconvex constrained optimization,” *Computational Optimization and Applications*, vol. 84, pp. 609–649, Mar. 2023.
- [57] E. Wei and A. Ozdaglar, “On the $O(1/k)$ convergence of asynchronous distributed alternating direction method of multipliers,” in *Proc. IEEE Global Conference on Signal and Information Processing*, (Austin, USA), pp. 551–554, 2013.
- [58] O. Shorinwa, T. Halsted, and M. Schwager, “Scalable distributed optimization with separable variables in multi-agent networks,” in *Proc. American Control Conference (ACC)*, (Denver, USA), pp. 3619–3626, July 2020.
- [59] F. Dellaert, “Factor graphs and GTSAM: A hands-on introduction,” technical report, Georgia Institute of Technology, 2012.
- [60] K. Murphy, Y. Weiss, and M. Jordan, “Loopy belief propagation for approximate inference: An empirical study,” in *Proc. Fifteenth Conference on Uncertainty in Artificial Intelligence*, (Stockholm, SE), p. 467–475, 1999.
- [61] D. McGann, J. G. Rogers III, and M. Kaess, “Robust incremental smoothing and mapping (riSAM),” in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, (London, GB), pp. 4157–4163, May 2023.