

MIT Open Access Articles

Talk Resource-Efficiently to Me: Optimal Communication Planning for Distributed Loop Closure Detection

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Giamou, Matthew et al. "Talk Resource-Efficiently to Me: Optimal Communication Planning for Distributed Loop Closure Detection." IEEE International Conference on Robotics and Automation (ICRA), May 2018, Institute of Electrical and Electronics Engineers, September 2018. © 2018 IEEE

As Published: <http://dx.doi.org/10.1109/icra.2018.8460783>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/126635>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Talk Resource-Efficiently to Me: Optimal Communication Planning for Distributed Loop Closure Detection

Matthew Giamou[†] Kasra Khosoussi[†] Jonathan P. How

Abstract—Due to the distributed nature of cooperative simultaneous localization and mapping (CSLAM), detecting inter-robot loop closures necessitates sharing sensory data with other robots. A naïve approach to data sharing can easily lead to a waste of mission-critical resources. This paper investigates the logistical aspects of CSLAM. Particularly, we present a general resource-efficient communication planning framework that takes into account both the total amount of exchanged data and the induced division of labor between the participating robots. Compared to other state-of-the-art approaches, our framework is able to verify the same set of potential inter-robot loop closures while exchanging considerably less data and influencing the induced workloads. We develop a fast algorithm for finding globally optimal communication policies, and present theoretical analysis to characterize the necessary and sufficient conditions under which simpler strategies are optimal. The proposed framework is extensively evaluated with data from the KITTI odometry benchmark datasets.

I. INTRODUCTION

Multi-robot, or cooperative, simultaneous localization and mapping (CSLAM) is an active area of research with a wide spectrum of applications that span from robotic search and rescue in challenging environments to navigating fleets of autonomous cars; see [1]–[3] for recent surveys. Communication is a crucial aspect of the approach, without which CSLAM would simply reduce to decoupled copies of conventional SLAM. In applications without pre-existing infrastructure, ad-hoc wireless communication is subject to many shortcomings, including energy constraints, bandwidth, and range limitations; see, e.g., [2], [4]. Overlooking these challenges could lead to impractical solutions. Ensuring that agents are able to effectively and resource-efficiently communicate with one another is one of the most challenging problems facing distributed CSLAM architectures [2].

Communication is an essential prerequisite for establishing loop closures between different robots’ trajectories and maps. To search for inter-robot loop closures, robots need to compare and match the data acquired throughout each of their individual trajectories. However, each robot initially has access only to the data collected by its own onboard sensors. As a result, robots need to frequently share data among themselves. State-of-the-art techniques either employ a centralized architecture, or simply require each robot to broadcast a down-sampled history of its sensory readings; see, e.g., [5] and [6]–[9], respectively. A naïve approach to the data sharing problem can easily lead to a waste of

mission-critical resources including battery, wireless bandwidth, and CPU time.

We present a general communication planning framework for resource-efficient data exchange in the search for inter-robot loop closures in distributed CSLAM front-ends. Our framework has several appealing features:

- ◊ A guarantee to be *lossless* in the sense that, for any given set of candidate matches, the proposed framework allows for a *complete search* of all inter-robot loop closures that exist within that set.
- ◊ Efficient algorithms for finding optimal exchange policies with respect to the total amount of data transmission with minimal computational overhead.
- ◊ Providing a mechanism through which one can retain communication efficiency while influencing the final induced division of labor between the robots. This allows the team to balance the resulting induced workloads based on the distribution of computational resources among the robots.
- ◊ Applicability to systems that use measurements and maps composed of any data type, including dense 3D laser scans and local image features, e.g., BRIEF [10].

A. Related Work

In a series of papers, Indelman et al. and Dong et al. [6]–[8] develop a pose-graph CSLAM framework based on Expectation-Maximization. From the perspective of data exchange efficiency, in [6]–[8] robots broadcast a downsampled subset of their (informative) raw measurements (e.g., laser scans) with each other. Our work can be employed alongside these and similar systems to provide an alternative resource-efficient communication plan.

Cieslewski and Scaramuzza [11] investigate the scalability of decentralized visual place recognition—in terms of the amount of exchanged data per place recognition query—in large teams of robots. In particular, they propose a decentralized approach whose scalability is comparable to that of centralized architectures and significantly better than the existing decentralized approaches. In [11] it is empirically shown that their heuristic approximation only suffers a mild reduction in place recognition recall. The core idea in [11] is to send partial queries to every other robot, assess the returned image similarity scores, and send the full query only to the robot with the most likely candidate match. Unlike the online (frame-query) flavour of the problem addressed in [11], our work considers a batch formulation that arises in occasional, but larger, data exchanges. The batch setting is especially well suited to applications in which multiple

[†]Equal contribution.

Authors are with the Laboratory for Information and Decision Systems (LIDS), MIT, Cambridge, MA — {mgiamou, kasra, jhow}@mit.edu




Fig. 1: An overview of distributed sensory data exchange approaches in CSLAM. Figure 1a illustrates a monolog (unidirectional policy), in which one robot sends all of its scans to the other. In some of the state-of-the-art techniques, this process happens also in the opposite direction (both robots share all of their scans with each other); see, e.g., [6]. In addition to the sensory data, robots also need to transmit a smaller amount of information (“metadata” \mathcal{M}) to help identify potential loop closures (e.g., compact bag-of-words (BoW) vectors for visual place recognition or sparse trajectories). Figures 1b, 1c, and 1d illustrate the proposed approach. Contrary to Figure 1a, here robots engage in a dialog, and each shares a subset of its sensory data with the other robot. We demonstrate that this process can significantly reduce the total amount of exchanged data. In our approach, robots still need to exchange metadata. The broker (\rightleftarrows) then solves the optimal data exchange problem and sends the optimal exchange policy π^* for execution to the robots.

robots are distributed to cover a large space, and communication is only possible during rendezvous. In such settings, rendezvous are seen as short-lived valuable opportunities that can be leveraged to better achieve the mission objective. Furthermore, while [11] concerns fleet-wide communication efficiency in detecting potential inter-robot loop closures, here we focus on local (pairwise) efficiency in exchanging sensory data. Despite these differences, an extension of our framework to n -way data exchanges can use the idea behind [11] to improve the communication efficiency of its metadata exchange phase (see Section II).

Sharing compressed beliefs and graphs constitutes another type of information exchange that arises in CSLAM. State-of-the-art techniques often marginalize out unnecessary intermediate poses from the belief to reduce the amount of exchanged data; see, e.g., [4], [12]. The resulting information matrix, however, is generally dense. This has led to the study of approximate sparsification techniques to “compress” the reduced beliefs. Paull et al. [4] investigate CSLAM with acoustic communication in the context of autonomous underwater vehicles. They propose a consistent (conservative) sparsification scheme based on Kullback-Leibler divergence. Lazaro et al. [9] propose to transmit a reduced representation of robots’ graphs (“condensed graphs”), as well as the most recent laser scans. Sharing only the most recent laser scans comes at the cost of losing potential loop closures in the regions that robots had explored separately prior to the encounter (i.e., before establishing a communication link). Cunningham et al. [12] propose a fast RANSAC-based data association scheme for CSLAM. The communication module in [12] shares the reduced beliefs (“condensed maps”) with a bounded number of robots within communication range. In contrast to our work, [12] considers a feature-based formulation with purely geometric (point) features. In that setting, each landmark measurement consists of a pair of range and bearing values, which is typically too lightweight to necessitate a communication planning framework.

Forster et al. [5] propose a centralized framework, in which the base station aggregates all visual information and establishes inter-robot and intra-robot loop closures; see

[2] for more centralized CSLAM approaches. Centralized approaches have limited applications and, compared to our work, leave no room for communication efficiency.

Montijano et al. [13] and Leonardos et al. [14] explore elegant formulations and algorithms for *solving* the distributed data association problem with an emphasis on maintaining association consistency across the communication graph. Unlike [13], [14], our paper takes a step back and investigates the logistics of distributed data association through exchanging data between pairs of robots. Our approach is orthogonal to such techniques and can be employed alongside distributed solvers.

In summary, our framework neither tells the agents *what* to say to each other—a question that is partly a system-dependant design choice, partly addressed by belief compression methods, see e.g., [4], [9], [15], and partly addressed by measurement selection schemes, e.g., see [16], [17]—nor does it tell them *what to do* with the exchanged data—i.e., how to solve the data association or the resulting inference problem, which is addressed by works such as [7], [12]–[14] among others; it rather advises them on *how* to communicate more effectively and efficiently.

B. Contribution

This paper addresses the data exchange problem, a key prerequisite for realizing resource-efficient distributed inter-robot loop closure detection and place recognition. We formalize the problem, provide a theoretical analysis, and shed light on its connection to the weighed minimum bipartite vertex cover problem. These insights ultimately lead to a fast algorithm for finding globally optimal communication plans based on linear programming. Additionally, we experimentally validate the proposed framework based on real benchmark datasets.

Notation

Bold lower-case and upper-case letters are reserved for vectors and matrices, respectively. $\mathbf{1}$ and $\mathbf{0}$ denote, respectively, the column vectors of all ones and all zeros. Sets are shown by upper-case letters. $|A|$ denotes the cardinality of

set \mathcal{A} . The disjoint set union operator is denoted by \uplus such that $\mathcal{A} \uplus \mathcal{B} = \mathcal{A} \cup \mathcal{B}$ and implies that $\mathcal{A} \cap \mathcal{B} = \emptyset$. For any two vertices u and v in a given graph, $u \sim v$ means that there is an edge connecting u to v . Finally, for any set of vertices \mathcal{S} , $\mathcal{N}(\mathcal{S})$ is the neighbourhood of \mathcal{S} in the graph.

II. OPTIMAL DATA EXCHANGE

This section proposes a resource-efficient framework to facilitate the search for inter-robot loop closures in CSLAM via exchanging collected sensory data (collectively called “scans” in this paper). Each exchange operation is moderated by an *exchange broker*, which can be a trusted hardware/software component located at one of the two participating robots (Figure 1d), or a trusted third party (e.g., another robot or a base station)—see Figures 1b and 1c. The broker has the duty of initiating, planning, and executing the operation. Unlike the *servers* in centralized approaches, the exchange brokers are not meant to aggregate scans, but rather to advise robots on the “optimal” exchange policy. Although only one broker is needed per exchange process, the total number of (potential) brokers in a team may vary between 1 (a central broker) and the number of robots (each robot can act as a broker if necessary), as long as the broker is able to communicate with the two participating robots during an exchange.

A. Initiating an Exchange

The exchange process can be initiated between two robots when they are within communication range. First, the broker has to form the *exchange graph* \mathcal{G} .

Definition 1 (Exchange Graph). An exchange graph is an undirected bipartite graph $\mathcal{G} = (\mathcal{V}_1 \uplus \mathcal{V}_2, \mathcal{L})$ whose vertices correspond to the two robots’ poses involved in the data exchange problem, and $\mathcal{V}_1 \ni u \sim v \in \mathcal{V}_2$ iff there is a “potential” inter-robot loop closure between their corresponding poses.

Without loss of generality and for convenience, we assume the degree of each vertex in the exchange graph is at least one. \mathcal{L} is a set of plausible inter-robot loop closure candidates, determined based on geometry (e.g., trajectory estimates and sensor characteristics such as field of view and range) and/or appearance (e.g., visual place recognition systems). Visual place recognition techniques like the DBoW2 system of [18] can be used to form elements of \mathcal{L} using only information local to individual measurements. In the case of DBoW2, this information is vocabulary labels of BRIEF [10] features extracted from query images. In both cases, \mathcal{G} is populated without sharing the entirety of the robots’ measurement data. Instead, \mathcal{L} is formed using a compact representation of the sensory data (hereafter, “metadata”), e.g., a collection of bag-of-words (BoW) vectors. Robots cooperate with the broker (by, e.g., providing information about their beliefs over their trajectories or BoW vectors) to form \mathcal{L} . In practice, a considerable number of potential edges are not plausible given the available information. This often makes \mathcal{G} far less dense than the complete bipartite

Algorithm 1 Execute a DEP

```

1: for  $v \in \mathcal{V}$  do
2:   if  $\pi(v) = 1$  then
3:     Share  $\mathcal{S}_v$  (corresponding sensory data).
```

graph. The structure of \mathcal{G} ultimately depends on a variety of factors, including the particular sensors and perception models utilized, the level of uncertainty in the robots’ beliefs, and perceptual aliasing.

B. Optimal Data Exchange

The optimal data exchange problem is now formally defined.

Definition 2 (Data Exchange Policy). A data exchange policy (DEP) is a vertex labeling that specifies which “scans” should be exchanged between a pair of robots. Formally, we call $\pi : \mathcal{V} \rightarrow \{0,1\}$ a DEP over $\mathcal{V} \triangleq \mathcal{V}_1 \uplus \mathcal{V}_2$ in which $\pi(v) = 1$ (resp., $\pi(v) = 0$) indicates that the scan collected at vertex v should (resp., should *not*) be sent to the other robot.

Based on the above definition, π can be *executed* simply by scanning the labels and transmitting the scans marked with “1” (i.e., to be sent); see Algorithm 1.

Definition 3 (Admissible Policy). A DEP is called *admissible* iff it allows for a *complete* search; i.e., finding all possible loop closures in \mathcal{L} . This can be achieved iff, for each edge in the exchange graph, at least one robot shares its associated scan with the other robot. Formally, π is admissible iff for all $u \sim v$, $\pi(u) + \pi(v) \geq 1$.

Definition 4 (Monolog). Let $\mathcal{V}_{\text{source}} \in \{\mathcal{V}_1, \mathcal{V}_2\}$. An exchange policy $\pi : \mathcal{V} \rightarrow \{0,1\}$ is called a *monolog* if

$$\pi : v \mapsto \begin{cases} 1 & v \in \mathcal{V}_{\text{source}}, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Lemma 1. Every monolog is admissible.

The broker can guarantee the completeness of search by proposing an admissible policy—but which one of them? Two primitive objectives are considered in this work:

1) **Communication:** The first objective quantifies the communication cost incurred during the execution of an exchange policy—mainly due to bandwidth and energy consumption. The communication cost is measured by the total amount of exchanged data. From this perspective, π is preferred over π' iff it can conduct a complete search by exchanging less data between the two robots. More precisely, let $w_s : \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$ be a weight function defined over \mathcal{V} such that $w_s(v)$ quantifies the “size” of scan \mathcal{S}_v collected at the corresponding pose. Then, the communication cost incurred as a result of executing policy π can be modelled as

$$f_{\clubsuit}(\pi) \triangleq \sum_{v \in \mathcal{V}} w_s(v) \pi(v). \quad (2)$$

In the special case of uniform weights, f_{\clubsuit} reflects the number of exchanges made by π (up to a constant).




Fig. 2: A simple data exchange problem between two robots. Each vertex corresponds to a robot pose, and each edge represents a potential loop closure between the corresponding robot poses. There is a “scan” associated with each robot pose. To verify a potential inter-robot loop closure between two connected vertices, at least one robot needs to share its scan with the other robot. (a) A simple exchange graph. (b) An admissible exchange policy in which each robot shares the sensory data collected at its red vertex with the other robot. The orientation of each edge signifies the direction of exchange (i.e., vertex label). (c) The workload induced by π : \mathbb{R}_2 is responsible for searching for loop closures among the green candidates, \mathbb{R}_1 will search among the blue candidates. Note that the thick candidate edge will be screened by both robots.

2) Induced Division of Labor: Upon executing an exchange policy, each robot has to perform sensor registration on a subset of \mathcal{L} . The exchange policy implicitly determines the distribution of the workload between the robots. The second objective captures this induced workload. To quantify this workload, first note that any admissible policy π divides the initial candidate set into $\mathcal{L} = \mathcal{L}_1^\pi \cup \mathcal{L}_2^\pi$ in which \mathcal{L}_1^π (resp., \mathcal{L}_2^π) is the set of edges incident to \mathcal{V}_2 (resp., \mathcal{V}_1) at a vertex v such that $\pi(v) = 1$. These sets can be empty (monolog) and are not necessarily disjoint: $\mathcal{L}_{12}^\pi \triangleq \mathcal{L}_1^\pi \cap \mathcal{L}_2^\pi$ is the set of edges like $\{u, v\} \in \mathcal{L}$ such that $\pi(u) = \pi(v) = 1$ (see Figure 2b). $\mathcal{L}_1^\pi \setminus \mathcal{L}_2^\pi$ (resp., $\mathcal{L}_2^\pi \setminus \mathcal{L}_1^\pi$) can only be searched by the first (resp., second) robot. On the contrary, in principle both robots can screen the candidates in \mathcal{L}_{12}^π . We can either divide the burden of searching in \mathcal{L}_{12}^π between the robots, or simply let each robot screen it on its own. The latter is preferred due to the following advantages. First, from a robustness perspective, verifying \mathcal{L}_{12}^π separately on each robot creates a desirable redundancy in case robots are unable to exchange their newly discovered loop closures due to problems like communication failure. Furthermore, the cost of post-exchange communication will be slightly reduced since we do not need to exchange the loop closures found in \mathcal{L}_{12}^π (Section II-D). Finally, as we will see shortly, this choice leads to tractable optimization problems.

Suppose the computational cost of verifying candidate inter-robot loop closure $\{u, v\}$ is quantified by $c_{uv} \geq 0$. The total computational cost due to sensor registration induced by exchange policy π on robot $i \in \{1, 2\}$ is given by

$$\ell_i^\pi = \sum_{v \in \mathcal{V} \setminus \mathcal{V}_i} \sum_{u \sim v} c_{uv} \pi(v). \quad (3)$$

Note that under uniform $\{c_{uv}\}_{u \sim v}$, ℓ_i^π is proportional to the number of potential loop closures that must be verified by robot i as a result of exchange policy π . Let α_1 and α_2 be non-negative parameters that control the induced workload balance between the two robots, such that, e.g., increasing α_i will shift the balance in favor of robot i . For example, in a heterogeneous data exchange between a typical robot and a tactical supercomputer, one may seek to choose an admissible policy such that most of the induced workload is redirected toward the tactical supercomputer. This narrative

results in

$$f_{\mathbb{R}}(\pi; \alpha_1, \alpha_2) \triangleq \alpha_1 \ell_1^\pi + \alpha_2 \ell_2^\pi \quad (4)$$

$$= \sum_{v \in \mathcal{V}} w_\ell(v) \pi(v), \quad (5)$$

in which

$$w_\ell : v \mapsto \begin{cases} \alpha_2 \sum_{u \sim v} c_{uv} & v \in \mathcal{V}_1, \\ \alpha_1 \sum_{u \sim v} c_{uv} & v \in \mathcal{V}_2. \end{cases} \quad (6)$$

Problem 1 (Optimal Data Exchange Problems (ODEP)).

$$\begin{aligned} & \underset{\pi}{\text{minimize}} && f_{\mathbb{R}}(\pi; \alpha_1, \alpha_2) \\ & \text{subject to} && \pi \text{ is admissible.} \end{aligned} \quad (\text{P}_1)$$

$$\begin{aligned} & \underset{\pi}{\text{minimize}} && f_{\mathbb{R}}(\pi) \\ & \text{subject to} && \pi \text{ is admissible.} \end{aligned} \quad (\text{P}_2)$$

$$\begin{aligned} & \underset{\pi}{\text{minimize}} && f_{\bullet}(\pi; \alpha_1, \alpha_2, \omega) \\ & \text{subject to} && \pi \text{ is admissible.} \end{aligned} \quad (\text{P}_3)$$

$$f_{\bullet}(\pi; \alpha_1, \alpha_2, \omega) \triangleq f_{\mathbb{R}}(\pi) + \omega f_{\mathbb{R}}(\pi; \alpha_1, \alpha_2) \quad (7)$$

$$= \sum_{v \in \mathcal{V}} w_{\bullet}(v) \pi(v), \quad (8)$$

in which $w_{\bullet} : v \mapsto w_s(v) + \omega w_\ell(v)$.

C. Solving the Optimal Data Exchange Problem

It is easy to see that $\text{P}_{1:3}$ are all instances of the weighted minimum bipartite vertex cover problem.¹ To see this, first note that the admissibility constraint needed for guaranteeing the completeness of search is identical to the constraint in vertex cover. Translating an instance of one of these narratives to an equivalent instance of the other (i.e., mapping a lossless exchange policy to an equivalent vertex cover $\pi \mapsto \Pi$ and vice versa) is trivial: $\Pi \triangleq \{v \in \mathcal{V} : \pi(v) = 1\}$ and $\pi : v \mapsto \mathbb{1}_{\Pi}(v)$ where

$$\mathbb{1}_{\Pi}(v) \triangleq \begin{cases} 1 & \text{if } v \in \Pi, \\ 0 & \text{if } v \in \mathcal{V} \setminus \Pi. \end{cases} \quad (9)$$

¹Finding a subset of vertices in a vertex-weighted bipartite graph with the minimum sum of vertex weights such that it covers every edge.

Finally, note that the cost of π (in ODEP) is equal to that of Π in the weighted minimum bipartite vertex cover, and vice versa. Consequently $P_{1:3}$ can all be solved using the same machinery. Furthermore, this result characterizes the communication cost incurred in the search for inter-robot loop closures and the induced workload balance in terms of the graph topology and vertex/edge weights through a well-understood graph invariant.

Algorithm

Although the weighted minimum vertex cover problem is NP-hard in general, it can be solved efficiently in bipartite graphs; see, e.g., [19]. Therefore, by virtue of the above-mentioned observation, we can solve any ODEP efficiently by casting it as a weighted minimum bipartite vertex cover problem. Moreover, Algorithm 1 can be slightly restructured to execute the vertex cover translation of an optimal policy—see Algorithm 2. It remains to describe an algorithm based on linear programming (LP) for efficiently solving ODEP. Let $w \in \{w_\ell, w_s, w_\bullet\}$. The corresponding ODEP can then be formulated as the following integer linear program (ILP):

$$\begin{aligned} & \underset{\pi}{\text{minimize}} && \sum_{v \in \mathcal{V}} w(v) \pi(v) \\ & \text{subject to} && \pi(u) + \pi(v) \geq 1 && u \sim v, \\ & && \pi(u) \in \{0,1\} && u \in \mathcal{V}. \end{aligned} \quad (\text{P}_{\text{ILP}})$$

The admissibility constraint in P_{ILP} can be compactly written as $\mathbf{A}^\top \boldsymbol{\pi} \geq \mathbf{1}$, in which \mathbf{A} is the unoriented incidence matrix of the exchange graph, and $\boldsymbol{\pi}$ is the stacked vector of values $\pi(u)$ for $u \in \mathcal{V}$. Let \mathbf{w} be the stacked vector of vertex weights. P_{ILP} admits a natural LP relaxation by expanding its feasible set \mathcal{F}_{ILP} into $\mathcal{F}_{\text{LP}} \triangleq \{\boldsymbol{\pi} : \mathbf{A}^\top \boldsymbol{\pi} \geq \mathbf{1}, \boldsymbol{\pi} \geq \mathbf{0}\} \supset \mathcal{F}_{\text{ILP}}$:

$$\begin{aligned} & \underset{\boldsymbol{\pi}}{\text{minimize}} && \mathbf{w}^\top \boldsymbol{\pi} \\ & \text{subject to} && \boldsymbol{\pi} \in \mathcal{F}_{\text{LP}}. \end{aligned} \quad (\text{P}_{\text{LP}})$$

It is well known that \mathbf{A} is *totally unimodular*, and therefore \mathcal{F}_{LP} is integral; i.e., P_{LP} has an integral solution that can be found using the simplex algorithm (see, e.g., [19, Ch. 18]). Any integral solution corresponds to an optimal exchange policy for Problem 1. In the special case of uniform weights, we can construct the optimal policy directly from the maximum bipartite matching in \mathcal{G} ; see König's theorem [19].

Optimality Conditions for Monologs

ODEP is built on the presumption that exploiting bidirectional communication can lead to more resource-efficient strategies. While this is generally true, in some special cases, monologs may perform optimally as well. The following theorem offers the necessary and sufficient condition for the most general form of $P_{1:3}$ under which a monolog is optimal.

Theorem 1. Consider a vertex-weighted exchange graph \mathcal{G} with non-negative weights assigned by $w : \mathcal{V} \rightarrow \mathbb{R}_{\geq 0}$. Let $\mathcal{V}_\circ \in \{\mathcal{V}_1, \mathcal{V}_2\}$. The monolog π defined as

$$\pi_\circ : v \mapsto \begin{cases} 1 & v \in \mathcal{V}_\circ, \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

minimizes the cost function

$$f(\pi) \triangleq \sum_{v \in \mathcal{V}} w(v) \pi(v) \quad (11)$$

among all admissible policies if and only if \mathcal{G} satisfies what we call the *generalized Hall's condition* (GHC):

$$\forall \mathcal{S} \subseteq \mathcal{V}_\circ : \quad \sum_{v \in \mathcal{S}} w(v) \leq \sum_{v \in \mathcal{N}(\mathcal{S})} w(v). \quad (\text{GHC})$$

Proof. [\Rightarrow] We show the contrapositive. Suppose there exists a $\mathcal{S} \subseteq \mathcal{V}_\circ$ that violates GHC. Consider,

$$\pi^* : v \mapsto \begin{cases} 1 & v \in (\mathcal{V}_\circ \setminus \mathcal{S}) \uplus \mathcal{N}(\mathcal{S}) \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

π^* is admissible since the vertices in $\mathcal{V}_\circ \setminus \mathcal{S}$ cover the edges that are not incident to \mathcal{S} , while those in $\mathcal{N}(\mathcal{S})$ cover every edge incident to \mathcal{S} . Now since \mathcal{S} violates GHC we have,

$$f(\pi^*) = \sum_{v \in \mathcal{V}_\circ \setminus \mathcal{S}} w(v) + \sum_{v \in \mathcal{N}(\mathcal{S})} w(v) \quad (13)$$

$$< \sum_{v \in \mathcal{V}_\circ \setminus \mathcal{S}} w(v) + \sum_{v \in \mathcal{S}} w(v) \quad (14)$$

$$= \sum_{v \in \mathcal{V}_\circ} w(v) \quad (15)$$

$$= f(\pi_\circ). \quad (16)$$

[\Leftarrow] Now we show GHC is sufficient. Suppose GHC holds and let π^* be the optimal admissible policy. For simplicity and without loss of generality let us assume $\mathcal{V}_\circ = \mathcal{V}_1$. Define $\Pi^* \triangleq \{v \in \mathcal{V} : \pi^*(v) = 1\}$ and $\Pi_i^* \triangleq \Pi^* \cap \mathcal{V}_i$ ($i = 1, 2$). If Π_2^* is empty, $\pi^* = \pi_\circ$. Furthermore, based on GHC, Π_1^* cannot be empty unless \mathcal{V}_1 and \mathcal{V}_2 have equal costs, which also implies that $\pi^* = \pi_\circ$. Thus we can assume both are non-empty. Since π^* is admissible, there must be no edges between $\mathcal{V}_1 \setminus \Pi_1^*$ and $\mathcal{V}_2 \setminus \Pi_2^*$. Therefore, $\mathcal{N}(\mathcal{V}_1 \setminus \Pi_1^*) \subseteq \Pi_2^*$. From GHC and the fact that vertex weights are non-negative we have,

$$\sum_{v \in \mathcal{V}_1 \setminus \Pi_1^*} w(v) \leq \sum_{v \in \mathcal{N}(\mathcal{V}_1 \setminus \Pi_1^*)} w(v) \quad (17)$$

$$\leq \sum_{v \in \Pi_2^*} w(v). \quad (18)$$

Consequently,

$$f(\pi^*) = \sum_{v \in \Pi_1^*} w(v) + \sum_{v \in \Pi_2^*} w(v) \quad (19)$$

$$\geq \sum_{v \in \Pi_1^*} w(v) + \sum_{v \in \mathcal{V}_1 \setminus \Pi_1^*} w(v) \quad (20)$$

$$= \sum_{v \in \mathcal{V}_\circ} w(v) \quad (21)$$

$$= f(\pi_\circ). \quad (22)$$

This concludes the proof. \square

Remark 1. Theorem 1 states that π_\circ is optimal iff, for any subset of vertices in \mathcal{V}_\circ , the amount of data that needs to

Algorithm 2 Execute a DEP via Vertex Cover II

- 1: **for** $v \in \Pi$ **do**
 - 2: Send \mathcal{S}_v to the other robot.
-

be transmitted from \mathcal{V}_o to the other robot is not greater than the amount of data needs to be transmitted in the opposite direction. Although this result is intuitive, the fact that the GHC is both necessary and sufficient is non-trivial.

Corollary 1. Let $\mathcal{V}_{\max} \in \{\mathcal{V}_1, \mathcal{V}_2\}$ be the vertex set with the larger α_i . The monolog π_1 defined as

$$\pi_1 : v \mapsto \begin{cases} 1 & v \in \mathcal{V}_{\max}, \\ 0 & \text{otherwise.} \end{cases} \quad (23)$$

is optimal with respect to P_1 .

Corollary 1 implies that P_1 always has a trivial optimal monolog solution. Nonetheless note that P_3 still allows us influence the induced division of labor based while retaining communication efficiency. Moreover, Corollary 1 also implies that the two objective functions f_{\boxplus} and f_{\boxtimes} blended together in P_3 are competing with each other to shift the structure of the optimal policy towards monologs (ideal workload balance) and dialogs (communication efficiency), respectively.

Corollary 2. Let $\mathcal{V}_{\min} \in \{\mathcal{V}_1, \mathcal{V}_2\}$ be the vertex set with smaller cardinality. The monolog π_2 defined as

$$\pi_2 : v \mapsto \begin{cases} 1 & v \in \mathcal{V}_{\min}, \\ 0 & \text{otherwise,} \end{cases} \quad (24)$$

is optimal with respect to P_2 under uniform weights iff \mathcal{G} satisfies Hall’s condition (HC): $\forall \mathcal{S} \subseteq \mathcal{V}_{\min} : |\mathcal{S}| \leq |\mathcal{N}(\mathcal{S})|$.

Corollary 2 states the necessary and sufficient condition under which the monolog π_2 is optimal. This result also follows directly from Hall’s marriage theorem and König’s theorem [19]. As an example, consider the case of k -regular bipartite graphs.¹ A well-known application of Hall’s marriage theorem implies that k -regular bipartite graphs satisfy HC [19]. Similarly, it is easy to check that HC holds in the complete bipartite graph. Corollary 3 follows from this result and Corollary 2.

Corollary 3. The monolog π_2 is optimal with respect to P_2 under uniform weights in k -regular, and in complete bipartite graphs.

D. Post-Exchange Protocol

After executing the optimal policy π^* , each robot has to verify the potential loop closures in a subset of \mathcal{L} (\mathcal{L}_1^π and \mathcal{L}_2^π ; see Section II-B) via sensor registration. Examining the candidates will lead to a set of inter-robot loop closures $\mathcal{L}^\boxplus \subseteq \mathcal{L}$. Because of the admissibility constraint, we know that $\mathcal{L}_1^\boxplus \cup \mathcal{L}_2^\boxplus = \mathcal{L}^\boxplus$ in which \mathcal{L}_i^\boxplus is the set of loop closures discovered by robot $i \in \{1, 2\}$ after executing an admissible exchange policy (i.e., the search is guaranteed to

¹A graph is called k -regular if all of its vertices have degree $k \geq 1$.

Algorithm 3 Optimal Data Exchange

- 1: **Robots:** Send the essential metadata to the broker
 - 2: **Broker:** Form \mathcal{G} (w/ dynamic pricing)
 - 3: **Broker:** Form and solve ODEP via LP relaxation
 - 4: **Robots:** Execute π^* — exchange scans
 - 5: **Robots:** Search for loop closures in \mathcal{L}_1^π and \mathcal{L}_2^π
 - 6: **Robots:** Exchange the discovered loop closures: $\mathcal{L}_i^\boxplus \setminus \mathcal{L}_{12}^\boxplus$
-

be complete). At this point, each robot is aware of its own set of newly discovered inter-robot loop closures; these sets will have a non-empty overlap iff $\mathcal{L}_{12}^{\pi^*} \cap \mathcal{L}^\boxplus$ is non-empty. If the communication channel is still available, robots can immediately share their newly discovered positive matches with each other by transmitting $\mathcal{L}_i^\boxplus \setminus \mathcal{L}_{12}^\boxplus$ ($i = 1, 2$). The exchange process ends here. At this stage, robots are able to closely examine every potential candidate, perform geometric verification, solve the sensor registration and data association problems, and establish relative measurements; see, e.g., [8], [13], [14], [20].

E. Exchange Inertia and Dynamic Pricing

In Problem 1, vertex weights quantify quantities such as the size of a scan, computational cost of sensor registration for the corresponding potential loop closures, and the desired workload balance. From a broader perspective, the weights can be interpreted as the *exchange inertia*, such that a smaller weight signifies more desire to share the associated scan with other robots, and vice versa. This broader interpretation allows us to incorporate a wider spectrum of objectives and constraints using the same underlying framework. In particular, robots and/or the broker may utilize a dynamic pricing strategy driven by various internal/external incentives. For example, these dynamic pricing schemes may depend on the specific role of a robot in the team, its capabilities, clearance level, privacy restrictions, and the available mission-critical resources.

III. EXPERIMENTS

Algorithm 3 summarises the entire ODEP process. This section presents results obtained using the KITTI dataset [21] to formulate realistic ODEP instances. KITTI was chosen for its long, data-rich trajectories, and accurate ground truth. ODEP instances are solved with the Gurobi LP solver.² Solving ODEP takes about 0.41 seconds in one of the largest exchange graphs encountered in our datasets (with more than 2×10^3 vertices and 96×10^3 edges) on an Intel Core i7-6820HQ CPU operating at 2.70 GHz. The runtime in realistic settings and using DBoW2 with $\alpha = 0.3$ (see Section III-B) is about 0.03 seconds. Due to space limitation, in this section we focus mainly on P_2 .

A. Trajectory Geometry Experiments

In order to create instances of ODEP with the KITTI dataset, we chose sequences of the odometry benchmark that contained considerable amounts of self-intersection and re-tracing in their ground truth trajectory. Each sequence is

²<http://www.gurobi.com>

divided into two parts corresponding to two distinct robots. For each pose in the trajectory, Oriented FAST and Rotated BRIEF (ORB) features [22] exceeding a variable FAST detection threshold are extracted from the associated color camera image. Since this set of features can be used to detect and compute loop closures between poses as part of a SLAM system [20], the number of extracted features determines the vertex weight $w_s(v)$ for the pose at vertex v . In regions with greater environmental detail, a greater number of ORB features are extracted. The KITTI dataset’s odometry ground truth is then used to form edges between nearby poses associated with each robot. This process results in an exchange graph \mathcal{G} with weights $w_s(v)$ that depends on a number of parameters:

- 1) FAST threshold k_F used to detect ORB features,
- 2) Data rate or measurement frequency f (KITTI data is provided at 10 Hz),
- 3) Maximum distance d_{\max} between poses that are candidate matches (i.e. $u \sim v$),
- 4) Minimum fraction η of range limited camera field of view (FOV) between poses that are candidate matches.

Varying these parameters leads to different structures in \mathcal{G} and variable communication savings when using ODEP. In practice, different sensors and varying confidence in robot trajectory estimates would permit empirical modelling of exchange graph formation. In this paper, we analyze ranges of the above parameters to capture a variety of problem instances. For example, large values of d_{\max} correspond to scenarios where each robot’s trajectory estimate is highly uncertain and, therefore, a greater range of nearby poses need to be considered loop-closure candidates. Figure 3 displays edges of \mathcal{L} in green for a particular set of parameters on KITTI odometry sequence 0 and sequence 6. Figure 4a and 4d display the communication savings of the optimal policy relative to monolog policies for sequences 0 and 6 when \mathcal{L} is formed between poses within a variable d_{\max} . Figures 4b and 4e report similar results when \mathcal{L} is formed with a variable minimum FOV overlap η . The abrupt jumps in cost seen in Figures 4d and 4e are caused by sequence 6’s particular trajectory. Figure 3b displays the simple elongated loop that sequence 6 follows, along with some candidate edges formed by the field of view threshold of $\eta = 0.4$. These settings lead to candidate edges across the thin loop which vanish for shorter values of d_{\max} and higher values of η , reducing the required communication cost.

Figure 4 demonstrates that solving ODEP enables the robots to reduce the amount of data to be exchanged by up to 5 MB over some monologs. Note that in some of the state-of-the-art systems, full bidirectional communication of measurements is utilized by default, resulting in at least the sum of the communication costs of both monologs (red and green curves) in Figures 4a-4f [6]. For a typical 11 Mb/s ad hoc WiFi network tested in our laboratory, 5 MB corresponds to approximately 5 seconds of transmission time. Thus, in addition to reducing use of network bandwidth and battery usage, communication reduction could potentially help to

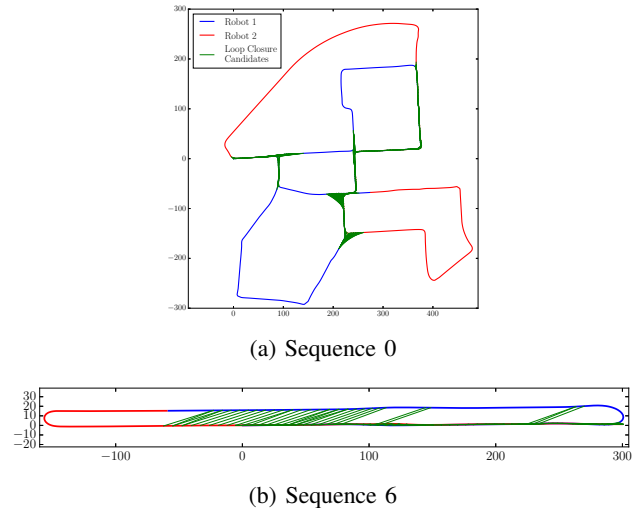


Fig. 3: Ground truth for KITTI odometry dataset sequences 0 and 6 with parameters $f = 2$ Hz, $\eta = 0.4$, $d_{\max} = 30$ m between robot 1 (blue) and robot 2’s (red) for edges (green). The edges and weights formed with ORB feature counts produce the exchange graph \mathcal{G} .

significantly shorten robot rendezvous periods in time-critical missions.

B. Place Recognition Experiments

An appearance-based place recognition system like DBoW2 [18] can also be used for inter-robot loop closure detection and to generate the candidate edge set \mathcal{L} . In a situation where robots do not have an accurate estimate of the transformation between their trajectories’ frames of reference, place recognition must be leveraged instead of viewpoint proximity to find potential loop closures. To facilitate place recognition, DBoW2 only needs the “word” in the bag-of-words vocabulary describing each ORB feature [18]. This word can typically be described in 3 or fewer bytes, which is less than one tenth of the size the standard 32 byte BRIEF descriptor used in ORB. Thus, an inexpensive exchange of vocabulary vectors (i.e., metadata) allows robots to search for promising candidates and form an exchange graph \mathcal{G} for ODEP.

In our experiments, we trained a DBoW2 vocabulary with parameters $k_w = 10$ and $L_w = 6$ on ORB features from 5 KITTI odometry benchmark sequences. The two candidate edges with the highest normalized score exceeding threshold α [18] were used to form candidate edges in \mathcal{L} . Communication savings from ODEP instances produced with KITTI odometry sequences 0 and 6 are displayed in Figures 4c and 4f. Although the structure of exchange graphs resulting from appearance based methods were very different from the geometric methods of the previous section, the cost of communication can still be significantly reduced using our method. For low α thresholds, the optimal policy affords significant savings of around 5-10 MB over the monolog policies in sequence 0 (Figure 4c). In the smaller sequence 6, the net communication savings are smaller because there are fewer candidates, but the optimal policy is still almost 10% more efficient than the best monolog policies at α thresholds in the range of 0.2 to 0.4. It is also worth noting

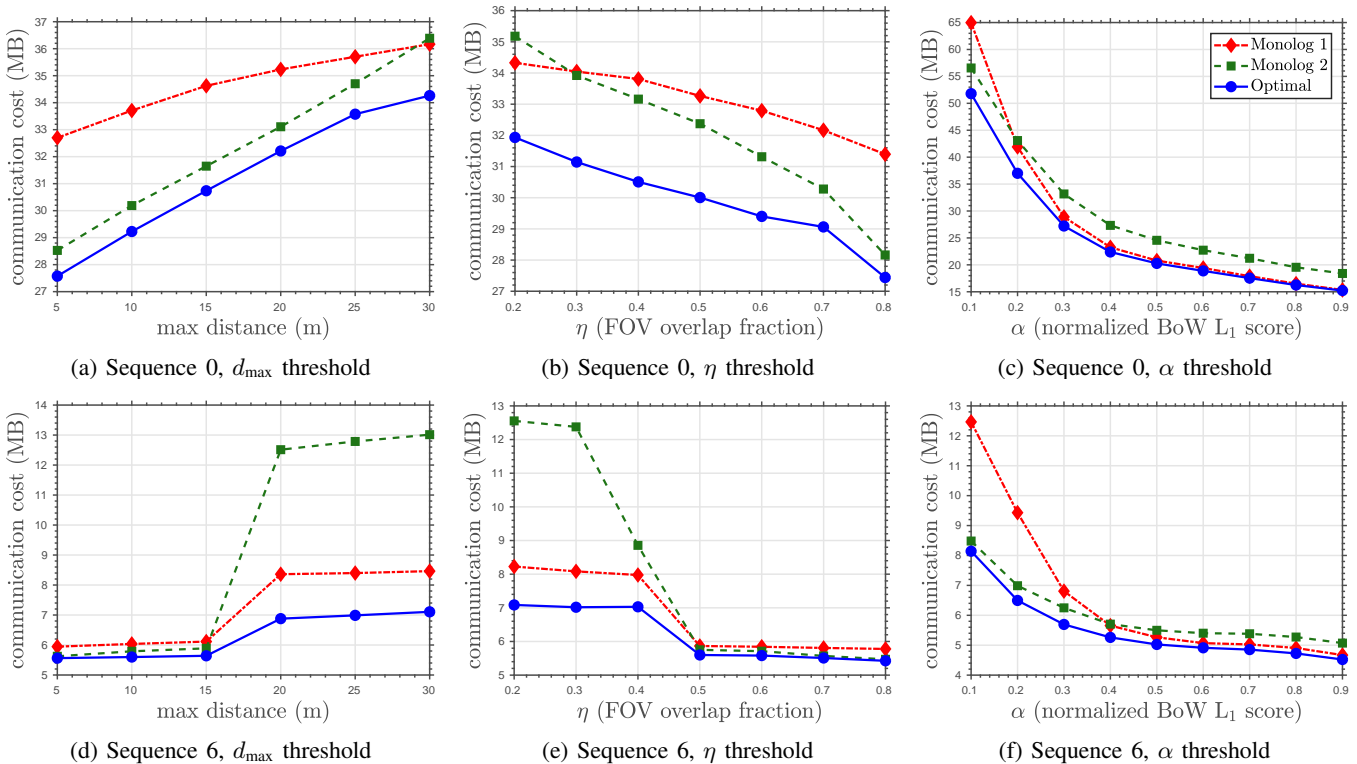


Fig. 4: Communication cost for KITTI odometry sequences 0 and 6 with $k_F = 100$, $f = 10$ Hz, and varying d_{\max} , η , or α (x-axes). Figures 4a and 4d form the candidate edge graph \mathcal{L} using maximum Euclidean distance d_{\max} between poses from KITTI groundtruth data, whereas Figures 4b and 4e use the fraction of overlapping fields of view η to determine candidates. Figures 4c and 4f form the candidate graph using the candidates with the 2 best DBoW2 scores greater than threshold α .

that relying on a single communication direction (i.e. using only one fixed monolog policy) throughout a mission is a poor communication heuristic that could produce arbitrarily bad and inconsistent results. This is illustrated in Figures 4f and 4b where depending on the value of α or η , the better choice of monolog and its performance penalty relative to the optimal policy changes drastically.

IV. CONCLUSION AND FUTURE WORK

Given the limitations of onboard resources, it is essential that robots communicate wisely. State-of-the-art techniques often have to sacrifice content by down-sampling the exchanged data, e.g., [6]–[8], [11]. This comes at the risk of losing potential valuable inter-robot loop closures which are the essence of cooperative localization and mapping. This paper addressed this challenge by investigating the logistical aspect of sensory data sharing in distributed CSLAM frontends. First, we formalized the optimal data exchange problem that encompasses a wide range of sensing modalities (e.g., vision, 2D and 3D lasers). This led to a resource-efficient and provably lossless (i.e., ensuring a complete search) communication planning framework. The proposed framework takes into account both the quantity of exchanged data, and the resulting division of labor induced by the executed exchange policy. This allows us to design efficient communication plans while distributing the induced workload based on, for example, the distribution of computational resources among robots. Additionally, ODEP can seamlessly

incorporate privacy and security constraints through the concept of exchange inertia and dynamic pricing schemes. Our approach benefits greatly from several fundamental results in graph theory and combinatorial optimization. In particular, these results lead to a fast and provably tight LP relaxation scheme to find the globally optimal exchange policies. In addition, our theoretical analysis characterized the necessary and sufficient conditions under which simpler unidirectional exchange policies are optimal. Finally, we experimentally validated geometric and appearance-based realizations of the proposed framework using the KITTI odometry benchmark datasets.

In retrospect, several crucial insights played major roles in the success of our approach. First and foremost, identifying plausible inter-robot loop-closure candidates *before* transmitting the bulk of sensory data is what makes communication planning possible. Forming the exchange graph and exploiting its unique structure (topology and the vertex/edge weights) allowed us to identify more efficient, yet lossless, exchange policies—often emerging as natural dialogs. Although this requires exchanging “metadata”, the incurred cost is often not comparable to the that of the actual data exchange. For example, visual place recognition systems like DBoW2 form loop closure candidates with sparse feature vectors that use an order of magnitude less data than the full descriptors used for subsequent loop closure verification. In our experiments where robots found candidate edges by exchanging pose graphs, poses are described by $SE(2)$ or

SE(3) objects that are much smaller than hundreds of visual descriptors. Furthermore, we exploited the sparsity pattern of the graph in our implementation to solve the resulting LP even faster.

This paper provides a solid foundation for optimal communication planning in distributed CSLAM front-ends. Our approach is able to find the optimal exchange policy between a pair of robots during pairwise encounters. n -way ($n > 2$) scan exchange problems naturally arise in robotic networks with denser communication graphs. Although the proposed approach can still be used in these cases, it may not necessarily lead to the optimal strategy. Addressing the sensory data exchange between more than two robots requires exploring new mechanisms such as data caching and routing. The optimal n -way data exchange problem is our next challenge.

ACKNOWLEDGEMENT

This work was supported in part by the NASA Convergent Aeronautics Solutions project Design Environment for Novel Vertical Lift Vehicles (DELIVER) and by the Northrop Grumman Corporation. The authors would like to thank Sergii Iglin for the Graph Theory Toolbox¹, the team behind the KITTI dataset [21], and Noam Buckman of LIDS for assistance with experiments.

REFERENCES

- [1] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [2] S. Saeedi, M. Trentini, M. Seto, and H. Li, "Multiple-robot simultaneous localization and mapping: A review," *Journal of Field Robotics*, vol. 33, no. 1, pp. 3–46, 2016.
- [3] V. Indelman, "Distributed perception and estimation: a short survey," *Principles of Multi-Robot Systems, workshop in conjunction with Robotics Science and Systems*, 2015.
- [4] L. Paull, G. Huang, M. Seto, and J. J. Leonard, "Communication-constrained multi-robot cooperative SLAM," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 509–516.
- [5] C. Forster, S. Lynen, L. Kneip, and D. Scaramuzza, "Collaborative monocular SLAM with multiple micro aerial vehicles," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 3962–3970.
- [6] J. Dong, E. Nelson, V. Indelman, N. Michael, and F. Dellaert, "Distributed real-time cooperative localization and mapping using an uncertainty-aware expectation maximization approach," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 5807–5814.
- [7] V. Indelman, E. Nelson, J. Dong, N. Michael, and F. Dellaert, "Incremental distributed inference from arbitrary poses and unknown data association: Using collaborating robots to establish a common reference," *IEEE Control Systems*, vol. 36, no. 2, pp. 41–74, 2016.
- [8] V. Indelman, E. Nelson, N. Michael, and F. Dellaert, "Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 593–600.
- [9] M. T. Lazaro, L. M. Paz, P. Pinies, J. A. Castellanos, and G. Grisetti, "Multi-robot SLAM using condensed measurements," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 1069–1076.
- [10] M. Calonder, V. Lepetit, C. Strecha, and P. Fua, "Brief: Binary robust independent elementary features," *Computer Vision—ECCV 2010*, pp. 778–792, 2010.
- [11] T. Cieslewski and D. Scaramuzza, "Efficient decentralized visual place recognition using a distributed inverted index," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 640–647, 2017.
- [12] A. Cunningham, K. M. Wurm, W. Burgard, and F. Dellaert, "Fully distributed scalable smoothing and mapping with robust multi-robot data association," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 1093–1100.
- [13] E. Montijano, R. Aragues, and C. Sagues, "Distributed data association in robotic networks with cameras and limited communications," *IEEE Transactions on Robotics*, vol. 29, no. 6, pp. 1408–1423, 2013.
- [14] S. Leonardos, X. Zhou, and K. Daniilidis, "Distributed consistent data association via permutation synchronization," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 2645–2652.
- [15] L. Paull, G. Huang, and J. J. Leonard, "A unified resource-constrained framework for graph SLAM," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016, pp. 1346–1353.
- [16] K. Khosoussi, G. S. Sukhatme, S. Huang, and G. Dissanayake, "Designing sparse reliable pose-graph SLAM: A graph-theoretic approach," *International Workshop on the Algorithmic Foundations of Robotics*, 2016.
- [17] V. Ila, J. M. Porta, and J. Andrade-Cetto, "Information-based compact pose SLAM," *Robotics, IEEE Transactions on*, vol. 26, no. 1, pp. 78–93, 2010.
- [18] D. Gálvez-López and J. D. Tardós, "Bags of binary words for fast place recognition in image sequences," *IEEE Transactions on Robotics*, vol. 28, no. 5, pp. 1188–1197, October 2012.
- [19] A. Schrijver, *Combinatorial optimization: polyhedra and efficiency*. Springer Science & Business Media, 2003, vol. 24.
- [20] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM: a versatile and accurate monocular SLAM system," *IEEE Transactions on Robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [21] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *International Journal of Robotics Research (IJRR)*, 2013.
- [22] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *Computer Vision (ICCV), 2011 IEEE international conference on*. IEEE, 2011, pp. 2564–2571.

¹<https://goo.gl/XJdWUf>