

Distributed Simultaneous Localisation and Auto-Calibration using Gaussian Belief Propagation

Riku Murai, Ignacio Alzugaray, Paul H.J. Kelly and Andrew J. Davison

Abstract—We present a novel scalable, fully distributed, and online method for simultaneous localisation and extrinsic calibration for multi-robot setups. Individual a priori unknown robot poses are probabilistically inferred as robots sense each other while simultaneously calibrating their sensors and markers extrinsic using Gaussian Belief Propagation. In the presented experiments, we show how our method not only yields accurate robot localisation and auto-calibration but also is able to perform under challenging circumstances such as highly noisy measurements, significant communication failures or limited communication range.

Index Terms—Distributed Robot Systems, Localization, Calibration and Identification

I. INTRODUCTION

AS robots become increasingly ubiquitous, more of them are expected to jointly operate and coordinate in shared environments to perform complex and collaborative tasks. For many of the tasks, it is critical to estimate the robot's pose accurately and precisely, also known as robot localisation. While many localisation systems have been proposed in the literature, most of the works focus on single-robot setups. However, relying solely on a single robot's proprioceptive and exteroceptive sensors for localisation is challenging and often limiting, especially in a multi-robot scenario, where relative localisation is crucial to ensure that robots interact safely and effectively.

This motivates *multi-robot collaborative localisation*, where robots utilise each other's observations to enhance their own localisation accuracy and therefore, the overall accuracy of the multi-robot system. The accuracy of co-localisation, however, is heavily reliant on the quality of extrinsic calibration of the sensors (e.g. visual camera rigs, rangefinders) and the markers they can detect on other robots (e.g. AprilTags, reflective markers). While most works often take such extrinsic calibration for granted, in practice, the default in-factory calibration can only be precise to a certain degree. This is particularly important in multi-robot setups, where manual calibration becomes impractical and highly accurate in-factory, per-robot calibration incurs high operational costs.

In this paper, we envision a system (see Fig. 1) in which multiple robots co-localise themselves as they move and sense

Manuscript received: July, 14, 2023; Revised November, 20, 2023; Accepted December, 18, 2023.

This paper was recommended for publication by Editor Sven Behnke upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by EPSRC (EP/K008730/1 and EP/P010040/1).

Authors are with Department of Computing, Imperial College London, UK {riku.murai15, i.alzugaray, p.kelly, a.davison}@imperial.ac.uk

Digital Object Identifier (DOI): see top of this page.

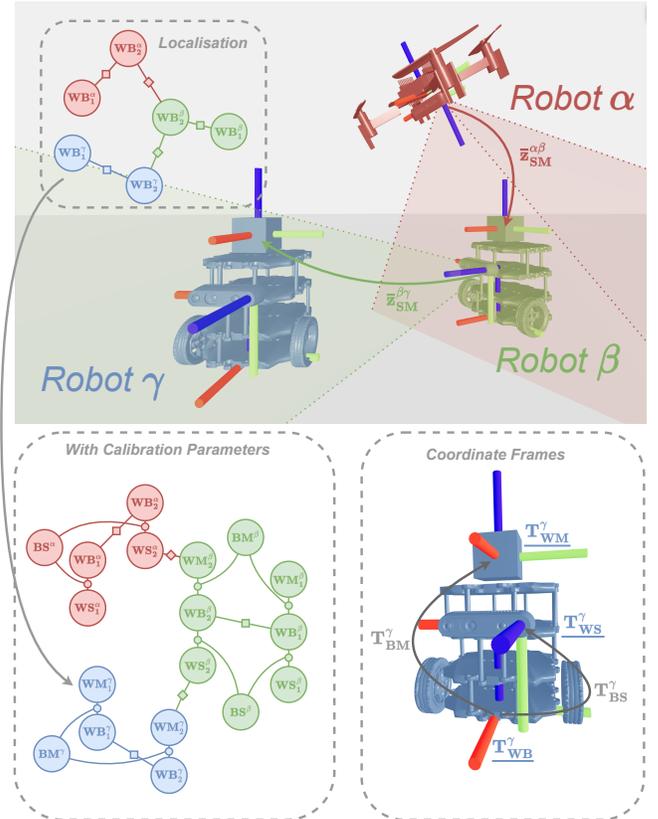


Fig. 1: Overview of the proposed auto-calibrating localisation system for three heterogeneous robots (top). Each robot observes the markers M placed on its peers to establish measurement \bar{z}_{SM} using sensor S mounted on top of a moving base B . Using the proposed methodology, the robots' relative positions and their calibration parameters can be retrieved in a distributed and asynchronous fashion performing probabilistic inference on a factor graph. We refer to \mathbf{T}_{WB} as \mathbf{WB} for clarity.

each other while simultaneously estimating and refining the extrinsic calibration of their sensor and their onboard marker on-the-fly. In summary, the contributions of our work are:

- A novel method for distributed multi-robot localisation and extrinsic calibration of both the sensor and the observed marker on the robots. Our approach builds on top of our previous work *Robot Web* [1], originally limited to pose variables in $\mathbf{SE}(2)$ given range-bearing observations. We extend the framework by simultaneously estimating the $\mathbf{SE}(3)$ pose of the robots and their extrinsic calibrations using Gaussian Belief Propagation

(GBP) to further improve the accuracy of multi-robot localisation.

- We present a formulation of the inter-robot factor that avoids the sharing of the calibration variables amongst multiple robots, sparing communication effort between robots and thus enhancing the scalability of the system.
- We provide an extensive evaluation of our approach in comparison with other state-of-the-art alternatives and measure the performance of the method under extreme conditions such as a large number of communication failures, a large proportion of outlying measurements, and a limited communication range.

II. BACKGROUND

A. Multi-Device Calibration

Calibration is vital for robotic operations and, as such, the body of literature on the subject is vast. Due to space limitations, our literature review focuses on calibration processes that involve multiple cameras or multiple robots.

Accurate extrinsic calibration is often critical in multi-device systems. Different robots have different base frames, and within a robot, the exact position of the sensor and observable onboard marker may not be available. A common instance of this is hand-eye calibration. From a set of known relative transformations, the calibration process seeks to establish the undetermined relationship, often the relative transformation between the robot base frame and sensor frame [2]. These methods can be extended to support multiple robots using iterative methods [3] or probabilistic approaches [4]. However, these methodologies primarily focus on offline settings where calibration precedes operational activities. In multi-robot setups, [5] proposes a method to perform online calibration of infra-red sensors while estimating the parameters of the underlying physical sensor model.

Multi-camera rigs are becoming increasingly popular as they can significantly extend the surrounding perceptive field for any robot and even directly yield stereo-depth capabilities provided there is view overlap. However, accurate calibration of such these rigs is often challenging, leading to several works on automatic offline calibration [6], [7], [8]. The method in [9] carries out continuous self-calibration using an extended Kalman filter in a stereo setup. Beyond two cameras, self-calibration of multiple cameras extrinsic is achieved on an aerial vehicle in [10], whereas an information-theoretic approach described in [11] is able to operate on a rig of eight cameras. Notably, while these methods are online, they predominantly address setups with a single robot with multiple onboard sensors, rather than a truly distributed, multi-robot system.

In the field of sensor networks, CaliBree [12] performs fully distributed sensor calibration by measuring disagreement between uncalibrated and calibrated sensors upon rendezvous event. This method, however, is only limited to calibration and does not address the localisation of the devices. In [13], GBP is used in a distributed fashion for intrinsic calibration and refinement of the camera poses. The method solves structure-from-motion, where multiple cameras are stationary; hence, it

is not applicable to online robotic applications with a moving onboard sensor. Non-parametric belief propagation is used in [14] to perform calibration and localisation of sensors. The method is sampling-based; hence less efficient than GBP and assumes that the sensors are stationary. LaSLAT [15] performs localisation and calibration of the sensors together with tracking of a target. In LaSLAT, the sensor poses are assumed to be static and are not suitable for localising multiple moving robots, which is the problem we address in this work.

B. Multi-Robot Localisation

Many recent advancements in multi-robot localisation leverage the advancements in distributed pose-graph optimisation (PGO). For example, [16] uses chordal-relaxation to make the underlying PGO problem linear, and solves them using a Gauss-Seidel solver. Semidefinite Programming (SDP) relaxation together with Riemannian block coordinate descent is used in [17], [18] which enables verification of the correctness of the estimates, and is decentralised and asynchronous. A distributed SLAM system [19] is built upon these approaches, demonstrating their practicality. However, the above approaches require a full relative transformation between the robots and can only handle isotropic covariance. This could be limiting, for example, if the inter-robot observations are all range-bearing, as we will explore in this paper. More recently, methods use range measurements [20] or bearing measurements [21] and show that it is possible to obtain certifiably optimal solutions again via SDP relaxation. While all PGO-based methods achieve good localisation accuracy, the formulation is often tailored and these methods do not generalise to other problem instances.

More general methods for multi-robot localisation such as DDF-SAM [22], and DDF-SAM2 [23] operate on factor graphs. They rely on Gaussian elimination and require robots to exchange Gaussian marginals about shared variables. The communication; however, increases quadratically with the number of shared variables. Alternating Direction Method of Multipliers (ADMM) has been employed for distributed SLAM [24] or to efficiently share map points for distributed bundle adjustment [25]. However, unlike elimination-based approaches only the point-estimates are recovered.

Our recent work, Robot Web [1] uses GBP to perform localisation. The proposed approach operates on general factor graphs; the underlying GBP algorithm is asynchronous and decentralised, thus the system is scalable to many robots.

III. PRELIMINARIES

A. Factor Graphs

A factor graph is a type of bipartite graph $G = (X, F, E)$ consisting of variable nodes $X = \{\mathbf{x}_i\}_{i=1:N_v}$ connected by edges E to factor nodes $F = \{f_s\}_{s=1:N_f}$. A factor graph represents the factorisation of the joint distribution: $p(X) \propto \prod_{s=1}^{N_f} f_s(X_s)$, where $X_s = n(f_s)$. Here, $n(v)$ is the set of neighbouring nodes connected via edge to the node v .

B. Belief Propagation

Belief Propagation (BP) [26] is an algorithm used to infer a marginal distribution of $p(\mathbf{x}_i)$ for each variable $\mathbf{x}_i \in X$ of a joint distribution $p(X)$ by passing messages among neighbouring nodes. As such operation is local, the algorithm is highly parallelisable and suitable for distributed inference. BP has promising applications in robotics, including multi-hypothesis and non-Gaussian inference [27].

In BP, at each iteration t , factor f_s sends a message $m_{f_s \rightarrow \mathbf{x}_i}^t$ to all variables $\mathbf{x}_i \in n(f_s)$, and variable \mathbf{x}_i sends back a message $m_{\mathbf{x}_i \rightarrow f_s}^t$ to all factors $f_s \in n(\mathbf{x}_i)$. At a variable \mathbf{x}_i , the product of all the incoming messages forms the belief $b^t(\mathbf{x}_i)$:

$$b^t(\mathbf{x}_i) = \prod_{f_s \in F_i} m_{f_s \rightarrow \mathbf{x}_i}^{t-1}(\mathbf{x}_i), \quad (1)$$

where $F_i = n(\mathbf{x}_i)$. The outgoing message to f_s from \mathbf{x}_i is simply the product of all other incoming messages from F_i :

$$m_{\mathbf{x}_i \rightarrow f_s}^t(\mathbf{x}_j) = \frac{b^t(\mathbf{x}_i)}{m_{f_s \rightarrow \mathbf{x}_i}^{t-1}(\mathbf{x}_i)}. \quad (2)$$

The message from a factor to a variable message is:

$$m_{f_s \rightarrow \mathbf{x}_i}^t(\mathbf{x}_i) = \sum_{X_s \setminus \mathbf{x}_i} f_s(X_s) \prod_{\mathbf{x}_j \in X_s \setminus \mathbf{x}_i} m_{\mathbf{x}_j \rightarrow f_s}^t(\mathbf{x}_j), \quad (3)$$

where $X_s = n(f_s)$. A factor takes a product of all the incoming messages together with its potential and marginalises out all the other adjacent variables $X_s \setminus \mathbf{x}_i$ to compute the outgoing message to \mathbf{x}_i .

C. Gaussian Belief Propagation

GBP [28] is a subset of BP where all the factors and, hence, the joint posterior distribution are Gaussians. Note that GBP has no theoretical convergence guarantees when applied to graphs with cycles and yet, it has shown robust performance across many different tasks [29], [30], [31].

We use canonical representation of Gaussian: $\mathcal{N}^{-1}(\mathbf{x}; \boldsymbol{\eta}, \boldsymbol{\Lambda}) \propto \exp(-\frac{1}{2}\mathbf{x}^\top \boldsymbol{\Lambda} \mathbf{x} + \boldsymbol{\eta}^\top \mathbf{x})$, where we define information vector $\boldsymbol{\eta} = \boldsymbol{\Lambda} \boldsymbol{\mu}$ and information matrix $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^{-1}$. Measurements $Z = \{\bar{\mathbf{z}}_m\}_{m=1:N_z}$ are modelled as $\bar{\mathbf{z}}_m = h_m(X_m) + \epsilon$, where $h_m(\cdot)$ is a measurement prediction function, and ϵ is a zero-mean Gaussian $\epsilon \sim \mathcal{N}(0, \boldsymbol{\Sigma}_m)$. Assuming independence of the observations, the likelihood of the observations is:

$$l(X; Z) = \prod_m^{N_z} l_m(X_m; \bar{\mathbf{z}}_m) \propto \prod_m^{N_z} \exp\left(-\frac{1}{2} \|\bar{\mathbf{z}}_m - h_m(X_m)\|_{\boldsymbol{\Sigma}_m}^2\right), \quad (4)$$

where the likelihood $l(X; Z) \propto p(Z|X)$. This notation is used to clarify that likelihood is a function of X , with Z as its parameter [32]. We are interested in estimating the configuration which maximises the Maximum A Posteriori (MAP) estimate: $X^{MAP} = \arg \max_X l(X; Z)P(X)$. GBP achieves this via marginal inference by computing the marginal posterior $P(\mathbf{x}_i|Z) = \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ where $\boldsymbol{\mu}_i = \mathbf{x}_i^{MAP}$ for all $\mathbf{x}_i \in X$. In

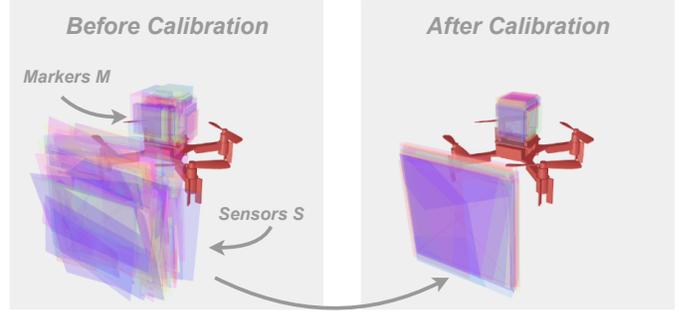


Fig. 2: Example of calibration of the extrinsic of the sensors' pose and markers' position using the proposed method, where we artificially set the ground-truth extrinsics to be the same for visual clarity. We overlay the calibration estimates of 64 robots from randomly initialised states (left), and visualise the estimated extrinsics after the calibration (right).

the factor graph representation of the posterior $P(X|Z)$, each \mathbf{x}_i is a variable node, and each l_m is a factor node connected via edge to nodes X_m . The reader is referred to [33], [34] for more details on GBP.

D. Gaussian Belief Propagation with Lie Groups

Given a non-linear measurement prediction function h_m it's first-order Taylor approximation is: $h_m(\mathbf{x}) \approx h_m(\mathbf{x}_0) + \mathbf{J}_m h_m(\mathbf{x} - \mathbf{x}_0)$, $\mathbf{J}_m = \frac{dh_m}{d\mathbf{x}}|_{\mathbf{x}=\mathbf{x}_0}$. We retrieve the likelihood of the descent direction $\Delta \mathbf{x} = \mathbf{x} - \mathbf{x}_0$ that minimises the local residual as in [1]:

$$l_m(\Delta X_m; \bar{\mathbf{z}}_m) = \mathcal{N}^{-1}(\Delta X_m; \mathbf{J}_m^T \boldsymbol{\Lambda}_m (\bar{\mathbf{z}}_m - h_m(\mathbf{x}_0)), \mathbf{J}_m^T \boldsymbol{\Lambda}_m \mathbf{J}_m). \quad (5)$$

For clarity, the Δ is dropped from the notation of the likelihood for the remaining of the paper. Energy of the factor is defined as: $E_m = (\bar{\mathbf{z}}_m - h_m(x_0))^\top \boldsymbol{\Lambda}_m (\bar{\mathbf{z}}_m - h_m(x_0))$.

We further apply GBP to variables on the Lie Group as in [1], which rely on the following expressions:

$$\mathcal{Y} = \mathcal{X} \oplus \mathcal{X}_\tau \triangleq \mathcal{X} \circ \text{Exp}(\mathcal{X}_\tau) \in \mathcal{M}, \quad (6)$$

$$\mathcal{X}_\tau = \mathcal{Y} \ominus \mathcal{X} \triangleq \text{Log}(\mathcal{X}^{-1} \circ \mathcal{Y}) \in T_{\mathcal{X}}\mathcal{M}, \quad (7)$$

where \mathcal{X}, \mathcal{Y} are points on the Lie Group \mathcal{M} , and \mathcal{X}_τ is a tangent vector in Lie algebra $T_{\mathcal{X}}\mathcal{M}$ defined locally around a point \mathcal{X} . Functions Exp, Log are the exponential, logarithmic mapping from $\mathcal{M} \rightarrow T_{\mathcal{X}}\mathcal{M}$ and $T_{\mathcal{X}}\mathcal{M} \rightarrow \mathcal{M}$ respectively, allowing us to move back and forth between Lie Group and Lie algebra. Composition operation \circ uses the group property and imposes that the composition of the elements remains in the Lie Group. To simplify the notation, we will drop \circ when composing two transformations.

IV. DISTRIBUTED LOCALISATION AND EXTRINSIC CALIBRATION

A. Gaussian Belief Propagation for Distributed Inference

Let G be the global factor graph which we want to perform inference over. In distributed GBP, robots own a factor graph G^ω each, and their union is $G = \bigcup_{\omega \in \Omega} G^\omega$. Each robot has

ownership of its pose variables and the factors corresponding to the observations it made, i.e. the nodes that their local graph G^ω consists of. This is important as the global factor graph G is partitioned amongst the robots Ω ; hence, the marginal estimates obtained by solving the distributed problem using GBP are *exactly* the same as the marginal estimates of the global problem obtained via centralised GBP under an assumption of perfect communication.

Distributed inference is achieved by each robot $\alpha \in \Omega$ performing GBP message passing on their local graph G^α . Along the edges of a factor $f^{\alpha\beta}$ owned by α , a factor-to-variable message $m_{f^{\alpha\beta} \rightarrow x^\beta}(x^\beta)$ is sent from α to β via inter-robot communication. Similarly, β sends back to α variable-to-factor message, $m_{x^\beta \rightarrow f^{\alpha\beta}}(x^\beta)$. A factor-to-variable message is a Gaussian distribution $\mathcal{N}(\mu_{f \rightarrow x}, \Lambda_{f \rightarrow x}^{-1})$, and a variable-to-factor message is a linearisation point \bar{x} and a Gaussian distribution $\mathcal{N}(\mu_{x \rightarrow f}, \Lambda_{x \rightarrow f}^{-1})$.

B. Problem Formulation

This section details how multi-robot localisation and the extrinsic calibration of their onboard sensors and observable markers are simultaneously and distributedly performed. In the considered setting, each robot is equipped with a range-bearing sensor S that observes the other robot's onboard marker M . While the addressed setup is representative of realistic constraints of many robotic applications, here we describe the formulation in general terms, so that it can easily be extended to support more information-rich measurements such as direct relative transformations (e.g. using visual sensors and fiducial markers).

The relative transformation from the base of the robot B to its sensor S is $\mathbf{T}_{BS} \in \mathbf{SE}(3)$, where the notation \mathbf{T}_{BS} represents the pose of S in the coordinate frame of B . Similarly, $\mathbf{t}_{BM} \in \mathbb{R}^3$ represents the marker position M relative to B . Since only range-bearing sensors are used in the current setup, the orientation of the marker is not observable and thus not included in this specific problem definition. When the sensor S^α in robot α observes the marker M^β in robot β , a relative measurement $\bar{\mathbf{z}}_{SM}^{\alpha\beta}$ is generated. The initial estimates of \mathbf{T}_{BS} and \mathbf{t}_{BM} are expected to be noisy due to inaccurate calibration. Our work optimises over the extrinsic calibration using the observations $\bar{\mathbf{z}}_{SM}^{\alpha\beta}$ robots accumulate over time as depicted in Fig. 2.

Let $\Omega = \{\alpha, \beta, \gamma, \dots\}$ be the set of robots, T be the number of considered time-steps, and $\mathbf{T}_{W\omega t}^\omega$ denote the pose of the base of the robot ω at time t in the world coordinate. To perform marginal inference over all, $\mathbf{T}_{BS}^\omega, \mathbf{t}_{BM}^\omega, \mathbf{T}_{W\omega t}^\omega, \forall \omega \in \Omega, \forall t \in \{1, \dots, T\}$, we consider the following factors.

1) *Range Bearing Sensor*: In our setup, robots can observe the other robots using range-bearing sensors. We use spherical coordinate (r, θ, ϕ) , i.e. radial distance, azimuthal angle, and elevation angle respectively. All angles are parameterised using $\mathbf{SO}(2)$, hence; the range bearing measurement is $\bar{\mathbf{z}}_{SM}^{\alpha\beta} \in \langle \mathbb{R}, \mathbf{SO}(2), \mathbf{SO}(2) \rangle$, a composite manifold.

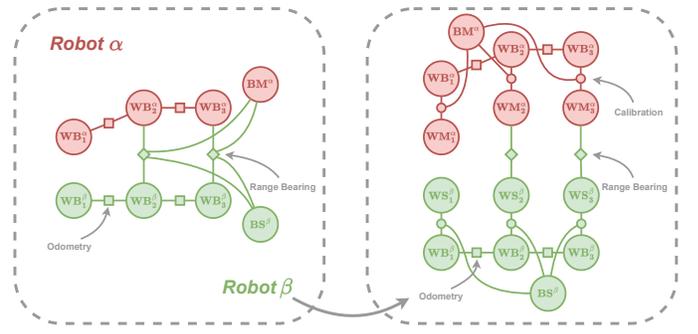


Fig. 3: Reducing the inter-robot communication by restructuring the factor graph. We refer to \mathbf{T}_{WB} as **WB** for clarity. **Left**: The inter-robot factor (range bearing) depends on four variables: the poses of the robots, the marker M , and the sensor S pose with respect to the robot base B . **Right**: We introduce the marker and sensor variable in the world coordinate frame using Eq. (10). While the total number of variables increases, the inter-robot factors depend on fewer variables, thus reducing the communication requirements.

The range bearing factor relating the sensor \mathbf{T}_{WSt}^α and the marker \mathbf{t}_{WMt}^β at time t is:

$$l_s(\mathbf{T}_{WSt}^\alpha, \mathbf{t}_{WMt}^\beta; \bar{\mathbf{z}}_{SM}^{\alpha\beta}) \propto \exp\left(-\frac{1}{2} \|\bar{\mathbf{z}}_{SM}^{\alpha\beta} \diamond h_s(\mathbf{T}_{WSt}^\alpha, \mathbf{t}_{WMt}^\beta)\|_{\Sigma_s}^2\right), \quad (8)$$

where we use the notation \diamond from [35], an operation on the composite manifold (\ominus operation is applied to each block of composites separately), and h_s is the function that predicts range bearing measurement between \mathbf{T}_{WS}^α and \mathbf{t}_{WM}^β .

2) *Robot Odometry*: We assume that odometry measurements $\bar{\mathbf{T}}_{B_{t-1}B_t} \in \mathbf{SE}(3)$ (e.g. IMU/wheel odometry) are made available to each robot. An odometry factor penalises the deviation between observation $\bar{\mathbf{T}}_{B_{t-1}B_t}$ and the two estimated consecutive poses $\mathbf{T}_{WB_{t-1}}, \mathbf{T}_{WB_t} \in \mathbf{SE}(3)$:

$$l_o(\mathbf{T}_{WB_{t-1}}, \mathbf{T}_{WB_t}; \bar{\mathbf{T}}_{B_{t-1}B_t}) \propto \exp\left(-\frac{1}{2} \|\bar{\mathbf{T}}_{B_{t-1}B_t} \ominus (\mathbf{T}_{WB_{t-1}}^{-1} \mathbf{T}_{WB_t})\|_{\Sigma_o}^2\right). \quad (9)$$

This assumes that the odometry measurement is measured in the base frame B . This property can be enforced by choosing a suitable base frame given prior information about the wheel/IMU position. However, if the odometry is provided via the sensor S (i.e. visual odometry), we can replace the transformations in the base frame B with a transformation in the sensor from S .

3) *Calibration Factor*: In Eq. (8), we have used $\mathbf{T}_{WSt}, \mathbf{t}_{WMt}$ position of sensor S and marker M in the world coordinate frame W at time t . A simple solution to obtain the position in the world coordinate is to use all $\mathbf{T}_{WBt}^\alpha, \mathbf{T}_{BS}^\alpha, \mathbf{T}_{WBt}^\beta, \mathbf{t}_{BM}^\beta$ inside the likelihood function, as shown in the left row of Fig. 3. However, this has a clear disadvantage: robots must communicate both the calibration estimate and the pose estimate with each other. This not only doubles the inter-robot communication effort but also exposes internal states (i.e. sensor calibration) that do not need to be revealed to other robots. Furthermore, it creates

small cycles which often leads to overconfidence [28]. Hence, this motivates the redesign of a factor to only share the pose estimate of the sensor and the marker between robots.

The objective of the calibration is to find a transformation $\mathbf{T}_{BS} \in \mathbf{SE}(3)$ such that: $\mathbf{T}_{WS} = \mathbf{T}_{WB}\mathbf{T}_{BS}$. This relationship as a likelihood is defined as:

$$l_c(\mathbf{T}_{WS}, \mathbf{T}_{WB}, \mathbf{T}_{BS}) \propto \exp\left(-\frac{1}{2}\|\text{Log}(\mathbf{T}_{WS}^{-1}\mathbf{T}_{WB}\mathbf{T}_{BS})\|_{\Sigma_c}^2\right). \quad (10)$$

The likelihood for calibration of the marker can be derived in a similar way. This allows us to create a factor graph as illustrated in the right row of Fig. 3, where only \mathbf{T}_{WSt} , \mathbf{t}_{WMt} is connected between the robots. While this formulation increases the total number of variables in the factor graph, fewer variables are connected to the inter-robot factor, thus reducing the data transfer between the robots.

C. Adaptive Regulariser on the Factor

Due to the nature of $\mathbf{SE}(3)$, the objective function which we are minimising is non-linear and non-convex; challenging for any iterative optimisers, but especially for local ones such as GBP with no access to the global objective function. In our case, the Lie group extension of GBP [1] was insufficient to consistently reach convergence. Hence, here, we introduce an adaptive regularisation term in GBP to assist convergence.

For each of the factors f_m , with a likelihood as defined in Eq. (5), we add a zero-mean prior $\mathcal{N}^{-1}(0, \lambda_m \mathbf{I})$. The term λ_m^t is local to the factor and is updated adaptively based on the difference between the current local factor energy E_m^t and last iterations E_m^{t-1} :

$$\lambda_m^t = \begin{cases} \lambda_m^{t-1} \cdot \lambda_{\uparrow} & E_m^t - E_m^{t-1} > \epsilon_{\lambda} \\ \lambda_m^{t-1} / \lambda_{\downarrow} & \text{otherwise} \end{cases}, \quad (11)$$

where a threshold ϵ_{λ} is required to avoid the weighting from increasing when the factors' energy stops changing significantly near convergence, and $\lambda_{\uparrow}, \lambda_{\downarrow}$ are the increase, decrease factor respectively. Intuitively, at the beginning when far from optima, the adaptive regulariser encourages small descent steps. As the factors become more confident about its approximation of the curvature (i.e. made multiple successive descents), larger descent steps are performed. While the principle of this approach is the same as Levenberg-Marquardt, fundamental this weighting scheme is computed and applied purely locally, and the step is always taken even if the local energy increases. This way, no synchronisation or communication is required when applied distributedly.

Assuming that the objective function is strictly convex, the addition of the adaptive regularisation term will not change the optimal solution. As the GBP converges, $\lim_{\lambda_m \rightarrow 0} l_s(X_s; \bar{\mathbf{z}}_s) \mathcal{N}^{-1}(0, \lambda_m \mathbf{I}) = l_s(X_s; \bar{\mathbf{z}}_s)$, and $\lambda_m \rightarrow 0$ as the energy decreases or reaches local convergence.

V. EVALUATION

We mainly evaluate our approach in a simulated environment with a vast number of robots, as obtaining the ground-truth extrinsic calibration and robot poses in the real-world

TABLE I: Accuracy of the proposed method ('Ours') and the global, centralised NLLS LM solver ('LM') at convergence as a function of the number of robots N and the enabling of autocalibration. Results include the RMSE ATE and ARE of the robot poses of their bases in the world frame \mathbf{T}_{WB} , the extrinsic calibration of their sensor \mathbf{T}_{BS} and marker \mathbf{t}_{BM} (only translation) where applicable.

N	T	Initial		LM w/ Calib.		Ours w/ Calib.		Ours w/o Calib.	
		[m]	[deg]	[m]	[deg]	[m]	[deg]	[m]	[deg]
16	\mathbf{T}_{WB}	0.432	7.422	0.065	1.858	0.084	1.970	0.093	2.313
	\mathbf{T}_{BS}	0.080	8.852	0.023	1.156	0.027	1.268	-	-
	\mathbf{t}_{BM}	0.085	-	0.020	-	0.022	-	-	-
32	\mathbf{T}_{WB}	0.434	7.471	0.051	1.742	0.062	1.811	0.075	2.138
	\mathbf{T}_{BS}	0.082	8.856	0.021	1.035	0.025	1.251	-	-
	\mathbf{t}_{BM}	0.087	-	0.019	-	0.022	-	-	-
64	\mathbf{T}_{WB}	0.436	7.402	0.043	1.684	0.054	1.761	0.066	2.082
	\mathbf{T}_{BS}	0.083	8.810	0.020	0.969	0.025	1.214	-	-
	\mathbf{t}_{BM}	0.088	-	0.018	-	0.020	-	-	-
128	\mathbf{T}_{WB}	0.434	7.385	0.039	1.646	0.049	1.732	0.060	2.041
	\mathbf{T}_{BS}	0.085	8.740	0.018	0.969	0.022	1.202	-	-
	\mathbf{t}_{BM}	0.087	-	0.017	-	0.020	-	-	-

for such experiments would be extremely challenging. As a verification of the applicability of our method to the real-world, we evaluate using UTIAS MR.CLAM dataset [36].

To simulate sensor noise, observations are corrupted by applying zero-mean Gaussian noise. Odometry measurements are corrupted with noise with σ_B^t , a standard deviation of 0.01 meter per meter travelled for the translation, and with σ_B^R , a standard deviation of 1 degree per 90 degrees rotated for the rotation. Inter-robot measurements are corrupted in their range and bearing readings with a standard deviation of σ_s : (0.05m, 5°). 3D range bearing measurements (r, θ, ϕ) are restricted to the three closest observable robots, to imitate realistic and limited inter-robot observability. We further restrict the range-sensing to be limited to $|\theta| < 60^\circ$ and $|\phi| < 60^\circ$, to simulate the field of view limitations of, for instance, visual sensors. Robots are randomly initialised in translation and orientation within a $20m \times 20m \times 20m$ space. We assume that the initial pose of the robots is known to a certain degree, within a noisy initial guess with standard deviation 0.01m, 1° respectively for the translation and the rotation. We simulate the robots' motion by drawing random samples from a uniform distribution, $\mathcal{U}(0, 1)m$ for translational motion and $\mathcal{U}(-\pi, \pi)^\circ$ for rotational motion, across all three dimensions. Finally, the initial calibration of sensor S and marker M also deviate from the ground-truth with a standard deviation of the translational part of extrinsic of sensor σ_S^t and marker σ_M^t set to 0.05m, and the standard deviation of the rotation part of the sensor frame $\sigma_S^R = 5^\circ$.

To enhance the stability of GBP, for the adaptive regularisation, we use the default parameter of $\lambda_m = 10, \lambda_{\downarrow} = 9, \lambda_{\uparrow} = 11, \epsilon_{\lambda} = 10^{-4}$. While not sensitive to the choice of parameters, we found GBP to diverge in many cases without adaptive regularisation. Additionally, 30% of both internal and external GBP messages are randomly dropped, as an empirical heuristic to improve the convergence of the system [37]. Unless specified otherwise, for robustness against outlying measurements, we dynamically scale the information matrix of the range-bearing sensor factor using a DCS robust kernel [38]: $s_m = \min\left(1, \frac{2\Phi}{\Phi + E_m}\right)$, with $\Phi = 10$. For each range-bearing

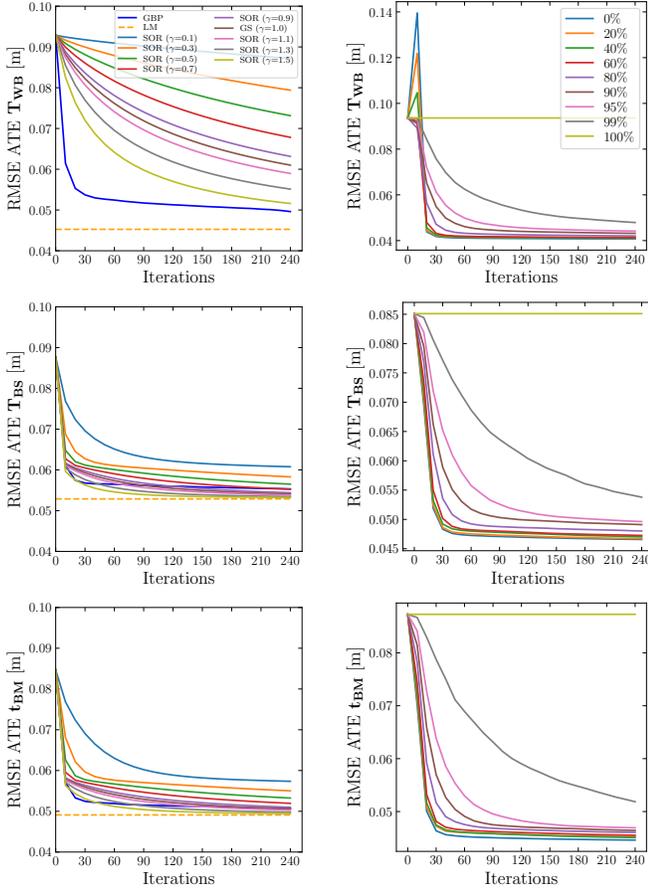


Fig. 4: From top to bottom: RMSE ATE for T_{WB} , T_{BS} , t_{BM} . RMSE ARE omitted as it follows the same trend. **Left:** Comparison of different distributed alternatives (Final RMSE ATE of global, non-distributed LM shown for reference). **Right:** Analysis of robustness regarding communication failures by randomly dropping a percentage of the inter-robot GBP messages in each iteration. 100% indicates that all inter-robot messages are dropped, preventing co-localisation.

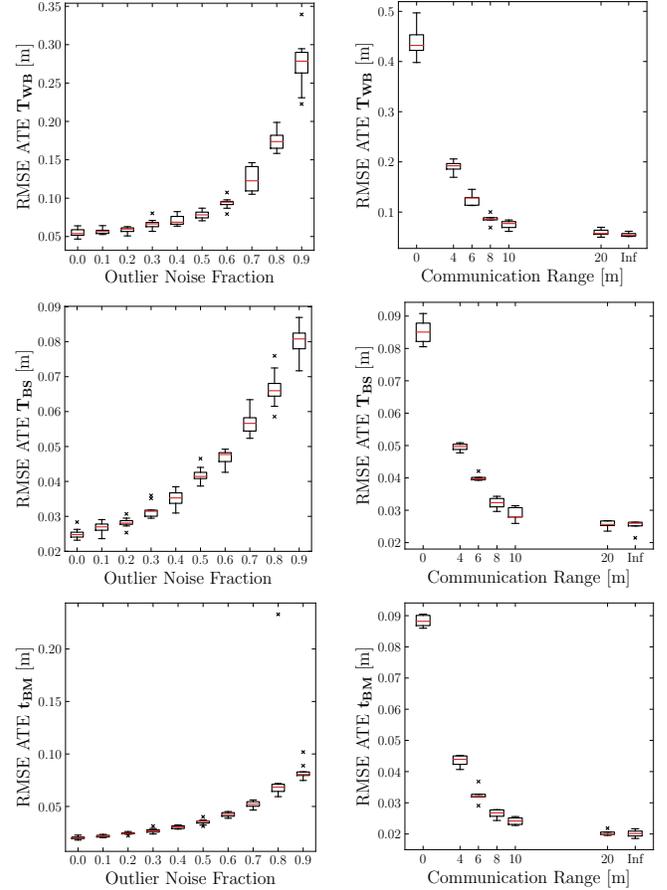


Fig. 5: From top to bottom: RMSE ATE for T_{WB} , T_{BS} , t_{BM} . RMSE ARE omitted as it follows the same trend. **Left:** Effect of increasing the fraction of outlier noise. Non-Gaussian noise is added to the inter-robot sensor measurement to simulate outliers. **Right:** Impact on the overall accuracy when robots are limited to only communicating with peers within the specified range.

measurement, the information matrix Λ_m is scaled by s_m^2 .

In all the presented experiments, unless specified otherwise, we consider $N = 64$ robots that randomly execute 50 motions, incrementally growing the underlying factor graph (see Section IV-B) and performing 30 GBP message-passing iterations after each of these motions. The experimental results aggregate information from a total of 10 randomised runs, where we often report the average Root Mean Squared Error (RMSE) of Absolute Trajectory Error (ATE) and Absolute Rotation Error (ARE) to measure the accuracy of the system as described in [16] for multi-robot setup.

A. Comparison with Centralised Factor-Graph Solvers

Here we compare the proposed incremental GBP-based approach with a global Non-Linear Least Squares (NLLS) Levenberg-Marquardt (LM) solver (implemented in Theus [39]) that processes the full graph as a whole batch. We evaluate how the accuracy of the overall system varies as a function of the number of robots N and the effect resulting

from enabling or disabling auto-calibration for the proposed method, i.e. whether the initial noisy calibration is optimisable or remains fixed, respectively. The robust kernel is disabled for this experiment to simplify the comparison. The accuracy of the different alternatives is compared in Table I. Despite the proposed method being distributed and without any global, second-order perspective of the whole problem, experimental results show no significant differences with respect to the global LM solver at convergence. As expected, the larger the number of robots N , the higher the accuracy of all methods as the underlying factor graph becomes denser and thus, more information-rich. Observe that the proposed GBP-based approach is still able to profit from a denser graph despite including more cycles. We additionally report the results of our method while considering that the noisy initial calibration is correct. This yields obviously worse results than when we optimise the graph which considers the calibration parameters.

B. Comparison with Distributed Factor-Graph Solvers

We compare our method against other distributed solvers: block Gauss-Seidel (GS) and its relaxations block Successive Over-Relaxation (SOR) [40]. In GBP messages are exchanged in parallel and thus do not require coordinated updates. We favour GS and SOR by counting each iteration as an ordered sweep, where robots sequentially exchange their updated state in a specific order. In this comparison, all the methods are provided with the whole graph to be optimised from the beginning instead of incrementally growing and solving the problem, with 16 robots making 10 random motions. We use the same relaxation parameter as reported in [16]. The robust kernel is disabled for this experiment to simplify the comparison.

Results are presented in the left column of Fig. 4. GBP shows a faster convergence rate than GS and SOR in the number of iterations (the aforementioned global and centralised LM method is also shown for reference). Presented results match prior comparisons between GBP and Successive Over-Relaxation in [28]. Note that GS and SOR produce marginally better extrinsic at by trading off a significantly worse body frame localisation.

C. Robustness Analysis

1) *Communication Failure and Asynchronicity*: We analyse the robustness of the system regarding potential communication failures, modelled by randomly dropping a percentage of the inter-robot GBP messages in each iteration, and present the results in the right column of Fig. 4. In this experiment, the whole graph is available from the beginning of the proposed algorithm with 16 robots and 50 random motions to clearly identify the convergence trends. The behaviour of the system remains largely unaffected by communication failure up to around 80%, communication failure. ATE of \mathbf{T}_{WB} initially increases as poses are initially uncertain and are down-weighted by the robust kernel. However, within a few iterations, the poses are correctly optimised, reducing the ATE. Even at an extremely high communication failure rate, the RMSE ATE still gradually decreases as we perform more iterations and thus, more rounds of communication. The experiment further demonstrates the asynchronicity of our approach, where the message order does not significantly impact the overall performance. This is a crucial property required for real-world deployment, where the communication channel is potentially unreliable, especially at scale.

2) *Robustness to Outlier Measurements*: As real-world inter-robot sensing is often challenging (e.g. misidentification, sensor failure), we investigate the robustness of the system to extreme, non-Gaussian outlier measurements following a uniform noise. As presented in the left row of Fig. 5, results indicate that, while performance is reasonably impacted as the fraction of outliers increases, the system remains stable even for an extremely high percentage of non-Gaussian outliers. Even at 40%, we observe a relatively small increase in the error compared to no outlying noise, demonstrating the robustness of our approach.

TABLE II: Evaluation of our method on real-world data. Results include the RMSE ATE of our system with and without the autocalibration enabled, on UTIAS MR.CLAM dataset 1-4. The noise column indicates whether a noise was artificially added to the sensor calibration to simulate an uncalibrated system.

Noise	Auto Calib.	1	2	3	4	Avg.
✗	✗	0.102	0.0976	0.0690	0.0712	0.0851
✗	✓	0.102	0.0967	0.0706	0.0694	0.0847
✓	✗	0.122	0.121	0.0974	0.0855	0.107
✓	✓	0.111	0.120	0.0825	0.0809	0.0984

3) *Communication Range*: To mimic realistic, real-world conditions, we further limit the communication radius of the robots to 4, 6, 8, 10 and 20m and report the results on the right row of Fig. 5, including also no communication (0m) and infinite communication range ('Inf') for completeness. While a longer communication range proves to be indeed beneficial, our method is able to optimise all the parameters effectively even with a severely limited communication radius. For reference, only 24% of the robots are within communication range at a 10m radius whereas the percentage increases to 85% at a 20m radius and yet such a drastic increase only yields insignificant returns in terms of accuracy.

The imposed limit on the communication range also tests the system's asynchronicity. Since an agent cannot communicate at the time of observation if the other agent is too far away, the agent needs to wait for a rendezvous event in order to exchange information. Our result hence further highlights that our system is capable of handling asynchronous events.

4) *Real-world Experiments*: To verify the applicability of our approach to real-world data, we have performed localisation and auto-calibration using UTIAS MR.CLAM dataset [36]. As the robots are ground vehicles, we model them with $\mathbf{SE}(2)$ poses and 2D range-bearing observations. We use datasets 1-4, where the landmarks and the robots are randomly scattered. We set $\sigma_B^t = (0.05\text{m}, 0.01\text{m})$, $\sigma_B^r = 5^\circ$, $\sigma_s = (0.08\text{m}, 2^\circ)$. The dataset is subsampled at 1s intervals, and we use sliding window-based GBP [1] with a window size of 30.

As presented in Table II, as the robots are calibrated, auto-calibration does not yield better RMSE ATE. We simulate an uncalibrated system by artificially adding noise to the sensor calibration with a standard deviation of 0.05m and 10° for translation and rotation respectively. While this manipulation of calibration is artificial, this data still contains challenging real-world sensor and odometry noise. In such a case, auto-calibration reduces the ATE, indicating that our method has successfully filtered out the biases even with real-sensor data.

VI. CONCLUSION

In this paper, we presented a method for online, simultaneous localisation and automatic extrinsic calibration of sensors and observable markers, by building on our previous work Robot Web [1]. Our work performs distributed and asynchronous inference on the factor graph using GBP, and we have demonstrated its robustness against large amounts of communication failure, outlying measurements, and restricted communication ranges.

Distributed and asynchronous properties of GBP offer attractive features for multi-robot systems. Automatic calibration ensures that the robots require as little maintenance as possible and the accurate localisation provides the basis required for multi-robot interaction.

ACKNOWLEDGMENT

The authors would like to thank Eric Dexheimer for insightful discussions.

REFERENCES

- [1] R. Murai, J. Ortiz, S. Saedi, P. H. Kelly, and A. J. Davison, "A robot web for distributed many-device localisation," *IEEE Transactions on Robotics*, 2023. **1, 2, 3, 5, 7**
- [2] Y. Shiu and S. Ahmad, "Calibration of wrist-mounted robotic sensors by solving homogeneous transform equations of the form $AX=XB$," *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 16–29, 1989. **2**
- [3] J. Wang, L. Wu, M. Q.-H. Meng, and H. Ren, "Towards simultaneous coordinate calibrations for cooperative multiple robots," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2014. **2**
- [4] Q. Ma, Z. Goh, S. Ruan, and G. S. Chirikjian, "Probabilistic approaches to the $AXB= YCZ AXB= YCZ$ calibration problem in multi-robot systems," *Autonomous Robots*, vol. 42, pp. 1497–1520, 2018. **2**
- [5] S. Goyal, A. Prorok, and A. Martinoli, "Two-phase online calibration for infrared-based inter-robot positioning modules," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2011. **2**
- [6] S. Esquivel, F. Woelk, and R. Koch, "Calibration of a Multi-camera Rig from Non-overlapping Views," in *Pattern Recognition. DAGM*, 2007, pp. 82–91. **2**
- [7] G. Carrera, A. Angeli, and A. J. Davison, "SLAM-based automatic extrinsic calibration of a multi-camera rig," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2011. **2**
- [8] Y. Lin, V. Larsson, M. Geppert, Z. Kukulova, M. Pollefeys, and T. Sattler, "Infrastructure-based multi-camera calibration using radial projections," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. **2**
- [9] T. Dang, C. Hoffmann, and C. Stiller, "Continuous stereo self-calibration by camera parameter tracking," *IEEE Transactions on Image Processing*, vol. 18, no. 7, pp. 1536–1550, 2009. **2**
- [10] L. Heng, G. H. Lee, and M. Pollefeys, "Self-calibration and visual SLAM with a multi-camera system on a micro aerial vehicle," *Autonomous Robots*, vol. 39, pp. 259–277, 2015. **2**
- [11] E. Dexheimer, P. Peluse, J. Chen, J. Pritts, and M. Kaess, "Information-theoretic online multi-camera extrinsic calibration," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4757–4764, 2022. **2**
- [12] E. Miluzzo, N. D. Lane, A. T. Campbell, and R. Olfati-Saber, "CaliBree: A self-calibration system for mobile sensor networks," in *Distributed Computing in Sensor Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 314–331. **2**
- [13] D. Devarajan, Z. Cheng, and R. J. Radke, "Calibrating distributed camera networks," *Proceedings of the IEEE*, vol. 96, no. 10, pp. 1625–1639, 2008. **2**
- [14] A. Ihler, J. Fisher, R. Moses, and A. Willsky, "Nonparametric belief propagation for self-calibration in sensor networks," in *Third International Symposium on Information Processing in Sensor Networks*, 2004, pp. 225–233. **2**
- [15] C. Taylor, A. Rahimi, J. Bachrach, H. Shrobe, and A. Grue, "Simultaneous localization, calibration, and tracking in an ad hoc sensor network," in *2006 5th International Conference on Information Processing in Sensor Networks*, 2006. **2**
- [16] 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," *International Journal of Robotics Research (IJRR)*, vol. 36, no. 12, pp. 1286–1311, 2017. **2, 6, 7**
- [17] Y. Tian, A. Koppel, A. S. Bedi, and J. P. How, "Asynchronous and parallel distributed pose graph optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5819–5826, 2020. **2**
- [18] Y. Tian, K. Khosoussi, D. M. Rosen, and J. P. How, "Distributed certifiably correct pose-graph optimization," *IEEE Transactions on Robotics (T-RO)*, vol. 37, no. 6, pp. 2137–2156, 2021. **2**
- [19] Y. Tian, Y. Chang, F. H. Arias, C. Nieto-Granda, J. P. How, and L. Carlone, "Kimera-multi: Robust, distributed, dense metric-semantic SLAM for multi-robot systems," *IEEE Transactions on Robotics (T-RO)*, vol. 38, no. 4, 2022. **2**
- [20] A. Papalía, A. Fishberg, B. W. O'Neill, J. P. How, D. M. Rosen, and J. J. Leonard, "Certifiably correct range-aided SLAM," *arXiv preprint arXiv:2302.11614*, 2023. **2**
- [21] Y. Wang, X. Wen, L. Yin, C. Xu, Y. Cao, and F. Gao, "Certifiably optimal mutual localization with anonymous bearing measurements," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9374–9381, 2022. **2**
- [22] A. Cunningham, M. Paluri, and F. Dellaert, "DDF-SAM: Fully distributed SLAM using constrained factor graphs," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2010. **2**
- [23] A. Cunningham, V. Indelman, and F. Dellaert, "DDF-SAM 2.0: Consistent distributed smoothing and mapping," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2013. **2**
- [24] D. McGann, K. Lassak, and M. Kaess, "Asynchronous distributed smoothing and mapping via on-manifold consensus ADMM," *arXiv preprint arXiv:2310.12320*, 2023. **2**
- [25] P. Bänninger, I. Alzugaray, M. Karrer, and M. Chli, "Cross-agent relocalization for decentralized collaborative SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2023. **2**
- [26] J. Pearl, "Reverend Bayes on inference engines: A distributed hierarchical approach," in *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 1982. **3**
- [27] D. Fourie, J. Leonard, and M. Kaess, "A nonparametric belief solution to the Bayes tree," in *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2016. **3**
- [28] Y. Weiss and W. T. Freeman, "Correctness of belief propagation in Gaussian graphical models of arbitrary topology," in *Neural Information Processing Systems (NIPS)*, 2000. **3, 5, 7**
- [29] A. Patwardhan, R. Murai, and A. J. Davison, "Distributing collaborative multi-robot planning with Gaussian belief propagation," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 552–559, 2023. **3**
- [30] J. Ortiz, M. Pupilli, S. Leutenegger, and A. J. Davison, "Bundle adjustment on a graph processor," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **3**
- [31] R. Scona, H. Matsuki, and A. Davison, "From scene flow to visual odometry through local and global regularisation in Markov random fields," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4299–4306, 2022. **3**
- [32] F. Dellaert and M. Kaess, "Factor Graphs for Robot Perception," *Foundations and Trends in Robotics*, vol. 6, no. 1–2, pp. 1–139, 2017. **3**
- [33] J. Ortiz, T. Evans, and A. J. Davison, "A visual introduction to Gaussian Belief Propagation," *arXiv preprint arXiv:2107.02308*, 2021. **3**
- [34] A. J. Davison and J. Ortiz, "FutureMapping 2: Gaussian Belief Propagation for Spatial AI," *arXiv preprint arXiv:1910.14139*, 2019. **3**
- [35] J. Solà, J. Deray, and D. Atchuthan, "A micro Lie theory for state estimation in robotics," *arXiv:1812.01537*, 2018. **4**
- [36] K. Y. Leung, Y. Halpern, T. D. Barfoot, and H. H. Liu, "The UTIAS multi-robot cooperative localization and mapping dataset," *International Journal of Robotics Research (IJRR)*, vol. 30, no. 8, pp. 969–974, 2011. **5, 7**
- [37] J. Ortiz, T. Evans, E. Sucar, and A. J. Davison, "Incremental Abstraction in Distributed Probabilistic SLAM Graphs," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2022. **5**
- [38] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard, "Robust map optimization using dynamic covariance scaling," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2012. **5**
- [39] L. Pineda, T. Fan, M. Monge, S. Venkataraman, P. Sodhi, R. T. Chen, J. Ortiz, D. DeTone, A. Wang, S. Anderson, J. Dong, B. Amos, and M. Mukadam, "Theseus: A Library for Differentiable Nonlinear Optimization," in *Neural Information Processing Systems (NIPS)*, 2022. **6**
- [40] D. Bertsekas and J. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Athena Scientific, 2015. **7**