

Radio-based Multi-Robot Odometry and Relative Localization

Andrés Martínez-Silva¹, David Alejo², Luis Merino¹ and Fernando Caballero¹

Abstract—Radio-based methods such as Ultra-Wideband (UWB) and RADAR Detection And Ranging (radar), which have traditionally seen limited adoption in robotics, are experiencing a boost in popularity thanks to their robustness to harsh environmental conditions and cluttered environments. This work proposes a multi-robot UGV-UAV localization system that leverages the two technologies with inexpensive and readily-available sensors, such as Inertial Measurement Units (IMUs) and wheel encoders, to estimate the relative position of an aerial robot with respect to a ground robot. The first stage of the system pipeline includes a nonlinear optimization framework to trilaterate the location of the aerial platform based on UWB range data, and a radar pre-processing module with loosely coupled ego-motion estimation which has been adapted for a multi-robot scenario. Then, the pre-processed radar data as well as the relative transformation are fed to a pose-graph optimization framework with odometry and inter-robot constraints. The system, implemented for the Robotic Operating System (ROS 2) with the Ceres optimizer, has been validated in Software-in-the-Loop (SITL) simulations and in a real-world dataset. The proposed relative localization module outperforms state-of-the-art closed-form methods which are less robust to noise. Our SITL environment includes a custom Gazebo plugin for generating realistic UWB measurements modeled after real data. Conveniently, the proposed factor graph formulation makes the system readily extensible to full Simultaneous Localization And Mapping (SLAM). Finally, all the code and experimental data is publicly available to support reproducibility and to serve as a common open dataset for benchmarking.

I. INTRODUCTION

Heterogeneous multi-robot systems have gained increasing popularity in the last few years. Deploying a team of different robots that work cooperatively - i.e. an Unmanned Ground Vehicle (UGV) and an Unmanned Aerial Vehicle (UAV) - can bring many benefits: one robot may be allowed to access locations that are blocked to the other, enabling the team to cover more terrain in less time. One key challenge is the localization of the robots in the same frame. This paper addresses how information sharing between robots via range-based radio sensors in GPS-denied environments potentially enables more robust multi-robot localization.

*This work was supported by the grants PICRAH 4.0 0 (PLEC2023-010353); funded by the Spanish Ministry of Science and Innovation and the Spanish Research Agency (MCIN/AEI/10.13039/501100011033); and COBUILD (PID2024-161069OB-C31), funded by the Spanish Ministry of Science, Innovation and Universities, the Spanish Research Agency (MICIU/AEI/10.13039/501100011033) and the European Regional Development Fund (FEDER, UE).

¹Andrés Martínez-Silva, Fernando Caballero and Luis Merino are with the Service Robotics Laboratory. University Pablo de Olavide, CP 41013, Seville, Spain.

²David Alejo is with the Department of System Engineering and Automation, University of Seville, Spain.

In a multi-robot team, each robot should be able to determine its own position and the relative position of the other robot. Although most of the research focuses on vision- or LiDAR-based methods, the unstructured layout of some areas, together with unfavorable environmental conditions such as fog, dust, or poor lighting critically impact the performance of these methods [1], [2]. Radio-based localization, in contrast, is gaining traction in the literature, especially Ultra-Wideband (UWB) and RADAR Detection And Ranging (radar) methods. UWB sensors provide high-frequency range measurements between anchors and tags with centimetric precision. They are also lightweight, cost-effective, and robust to harsh environmental conditions. Standard UWB-based localization systems require the anchors to be installed in the environment for effective trilateration, but this is often impractical. Deploying a team of robots, on the other hand, makes it possible to place the transceivers onboard to obtain inter-robot distances that can be fused with odometry measurements to determine the positions of both robots in real time. On the other hand, radars are able to estimate the position (in the form of range, azimuth and elevation) and radial velocity of detected objects by leveraging the Doppler effect. However, radar data brings a particular set of challenges regarding sparsity, lower angular and spatial resolution, and multimodal noise coming from reflections or signal scattering, among others [3].

While there is extensive research on relative transformation estimation based on range measurements, there is very little work on UWB-based heterogeneous multi-robot localization systems that: 1) leverage multiple on-board transceivers for redundancy and enhanced state observability, 2) fuse UWB with radar odometry to localize both robots with respect to a fixed frame, 3) involve robots operating in 2.5D and 3D spaces and 4) can run in real time.

The main contribution of this paper is a radio-based multi-robot UGV-UAV localization system that comprises three fundamental modules:

- A Nonlinear Least Squares (NLS) optimization framework to compute the relative transformation between the odometry frames of a UAV and a UGV based on multiple anchor-tag distance measurements, in real time.
- A pose-graph optimization framework that simultaneously optimizes the poses of both platforms by fusing all available sensor data, including the inter-robot transformation, the pre-processed radar point-clouds and other readily-available sensors such as IMUs or wheel encoders, in real time.
- A simulator plugin for Gazebo Harmonic that replicates UWB range measurements based on a measurement

model that has been validated against real sensor data. This plugin is compatible with multi-robot setups in the open-source PX4 Software-In-The-Loop (SITL) framework, narrowing the gap between simulation and real experimentation.

The system was validated in SITL simulations, and in a real-world experiment involving a UAV and a UGV. All the code and the dataset are publicly available ¹

The remainder of the paper is structured as follows: Section II reviews previous work on radio-based localization in multi-robot systems. Section III formalizes our approach by describing the proposed localization system, focusing on our three main contributions. Section IV validates our implementation in simulation and a real dataset. Finally, Section V presents some final remarks.

II. RELATED WORK

The problem of estimating inter-robot relative transformations based on distance measurements has been extensively studied in [4], [5], where the authors propose an algebraic method to compute a 6 Degrees of Freedom (DOF) rigid transformation between robot frames. A closed-form solution can be found from ten distance measurements, and a discrete set of 40 solutions from as few as six. In [6], the authors particularize the case to UWB sensors and achieve a 4-DOF closed-form solution from just six measurements. However, these methods are highly sensitive to outliers and noise and typically require further refinement with Nonlinear Least Squares (NLS). More recently, the authors of [7] perform an exhaustive analysis on 4-DOF relative frame transformation (RFT) estimation using UWB and visual-inertial odometry, proposing a global estimation approach to improve NLS initialization. However, these methods assume only one sensor per platform and do not exploit redundancy and frequency of multiple onboard measurements. Other research has considered up to four onboard UWB nodes [8] to compute instantaneous body frame transformations, though it relies on the availability of sufficient simultaneous measurements. In heterogeneous teams of robots, this assumption is often challenged due to limitations in sensor placement, communication quality and noise.

Radar-based localization and SLAM have been primarily explored in the context of autonomous driving, due to its advantages over other sensor modalities in adverse weather conditions. In [9], the authors present a radar-only motion estimation system using a Frequency Modulated Continuous Wave (FMCW) radar without Doppler information. In [10], a real-time pose-graph SLAM system that leverages Iterative Closest Point (ICP)-based scan-matching to obtain relative transformations between consecutive radar detections is introduced. To localize the vehicle, the odometry component implements an Unscented Kalman Filter (UKF) that fuses wheel speed, yaw rate, steering wheel angle and velocity of the radar sensor. The latter is obtained from radial (Doppler) velocities using a radar ego-motion

estimation algorithm [11]. In robotics, radar-based systems have only recently seen significant development. Notably, the authors of [12] introduced an open-source Robot Operating System (ROS) package that implements real-time 6-DOF SLAM for 4D radars. More recently, the authors of [3] proposed a graph-based radar-inertial odometry system with a loosely coupled ego-motion estimation algorithm for holonomic and non-holonomic ground robots. However, while these methods were designed with individual robot missions in mind, this paper applies the formulation to multi-robot missions by leveraging relative range data coming from UWB transceivers placed onboard two different platforms.

Finally, cooperative localization involving a team of robots has been widely covered in the past [13]–[15]. Recently, Brunacci et al. [16] used a model of the Dilution of Precision (DoP) of the UWB sensors to propose an active perception method to increase the accuracy of the relative transform estimation. Shalihan et al. [17] propose the MR-FLOUR method that uses UWB inter-robot measurements fused with wheel odometry and LiDAR inter-robot detections to enhance the precision of the localization system. However, it is difficult to find relevant work regarding teams of different robots that localize themselves with radio sensors, especially combining UWB and radar in 3-dimensional settings, as most of the methods in the literature focus on teams of UGVs.

To the best of our knowledge, this is the first work that integrates relative transformation estimation as an inter-robot factor in a multi-robot pose-graph optimization framework, together with radar factors in 3-dimensional settings.

III. PROBLEM FORMULATION

In this Section, we define the challenge to be tackled and introduce the main components of our multi-robot localization system. In Section III-A, we present our method to estimate the relative transformation between the odometry frames of the UAV and the UGV based on distance measurements. Both robots are equipped with UWB modules, which are placed in fixed locations relative to their body frame of reference. In Section III-B we introduce the radar odometry module, and finally, in Section III-C, we fuse all the available information in a pose graph optimization that computes the positions of both robots with respect to a common frame of reference.

A. Relative transform estimation

Suppose a heterogeneous team of two robots equipped with UWB transceivers moving in proximity, each of them describing a trajectory in their own local reference frame. The UAV (robot 1) is equipped with two tags, and the UGV (robot 2) is equipped with four anchors - see Figure 1.

Let ${}^{B_1}\tilde{\mathbf{p}}_{ti}, i \in \{1, 2\}$ be the known homogeneous coordinates of the UWB tag i in the body frame of robot 1, and ${}^{B_2}\tilde{\mathbf{p}}_{aj}, j \in \{1 : 4\}$ the known homogeneous coordinates of the UWB anchor j in the body frame of robot 2, and let ${}^{O_1}\mathbf{p}_1$ and ${}^{O_2}\mathbf{p}_2$ be the positions of robots 1 and 2 in their own local odometry frames, respectively. Consequently, the coordinates of tag i in the odometry frame of robot 1, ${}^{O_1}\tilde{\mathbf{p}}_{ti}$,

¹<https://github.com/robotics-upo/mr-radio-localization>




Fig. 1: Each robot describes a trajectory in its own local odometry frame while collecting range measurements (dashed lines) between two tags and four anchors. Measurement from different tags are represented in different colors.

and the coordinates of anchor j in the odometry frame of robot 2, ${}^{O_2}\tilde{\mathbf{p}}_{aj}$, are known.

Let ${}^{O_1}T_{O_2}$ be the rigid transformation between the odometry frames of the robots. The predicted distance $\hat{d}$ between tag i and anchor j is defined as follows:

$$\hat{d}(t_i, a_j) = \|{}^{O_1}\tilde{\mathbf{p}}_{ti} - {}^{O_1}T_{O_2} {}^{O_2}\tilde{\mathbf{p}}_{aj}\| + \epsilon \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, \sigma^2)$ represents independent, identically distributed Gaussian noise.

The objective is to find the inter-robot transformation ${}^{O_1}T_{O_2}$ between the odometry frames of both platforms, which allows us to map the local positions of robot 2 to the local reference frame of robot 1, according to the anchor-tag range measurements d . We solve this problem in terms of nonlinear least squares, where the cost function is defined as:

$$J = \min_{{}^{O_1}T_{O_2}} \sum_W \left(\sum_{i,j} \left(\frac{\hat{d}(t_i, a_j) - d}{\sigma} \right)^2 + \|e({}^{O_1}T_{O_2}, T_p)\|_{\Sigma}^2 \right) \quad (2)$$

The first term is a scalar that, for each anchor-tag pair, minimizes the difference between the predicted and the actual measurement, divided by its standard deviation σ , and the second term anchors the estimated transform to a prior T_p .

Finally, $\|e\|_{\Sigma}$ is the squared Mahalanobis distance $\mathbf{e}^T \Sigma^{-1} \mathbf{e}$ with covariance matrix Σ , and W is a distance-based sliding window comprising from 2 to 10 meters of each robot trajectory. Considering that both the UWB system and the odometry system provide estimates at a rate of $\approx 10\text{Hz}$ and there are eight available anchor-tag pairs reporting data simultaneously, each optimization problem produces in the order of ≈ 600 scalar UWB residuals when the robots are moving, although the number depends on the speed of the robots and the quality of the communications link.

In order to simplify the problem, we reduce it to 4-DOF by making the assumption that the roll and pitch of each robot's local odometry frame are obtained by means of the onboard IMUs, as they are able to estimate roll and pitch states very accurately even with low-cost IMUs [18] thanks to the integration of gyroscopes and accelerometers. Thus, only the

translation components and the yaw angle are expected to exhibit significant drift. In this way, we define the state vector $\mathbf{x} = [x, y, z, \theta]^T$, comprised of the three relative translation components and an elemental relative rotation along the vertical axis between the odometric frames. Our formulation uses SE(3) transformation matrices that we parameterize by means of the Lie Algebra formulation from [19]. Particularly, to reduce the problem to 4-DOF, we introduce the following error function $\mathbf{e}(\cdot)$:

$$\mathbf{e}(T_1, T_2) = \xi(T_1^{-1}T_2)_{[1:3,6]} \quad (3)$$

where $\xi(\cdot)_{[1:3,6]}$ takes the logarithmic map of its input, and retrieves the three translational components and the last angular component.

B. Radar Odometry

While LiDAR and vision-based methods are extensively used for localization via scan-matching or image correspondences, they are particularly sensitive to challenging environmental conditions. 4D radars, on the other hand, are robust to harsh weather and lighting conditions and can be used for localization using similar techniques. They provide range, azimuth, elevation and radial or Doppler velocity - i.e. the component of the object velocity along the sensor Line of Sight (LoS) - with each detection. This information is converted into a sparse point-cloud with per-point velocity estimates. However, the raw radar data must be first pre-processed to obtain intermediate representations that can then be used to estimate the robot positions (and the map, in the case of SLAM) in a factor graph optimization framework. With this purpose, we incorporate the existing work presented in [3] and extend it to a multi-robot scenario. In this method, the pre-processing consists of a filter to remove false positives and dynamic objects, as well as ego-motion estimation (this is, estimating the linear velocity of the sensor platform) by fusing radial velocity and IMU data in a loosely coupled formulation that accounts for holonomic ground vehicles.

Particularly, each Doppler velocity v_D is expressed as the projection of the robot velocity vector $\mathbf{v}$ onto the unit vector $\mathbf{h}$ pointing towards the detection:

$$v_D = \mathbf{h} \cdot \mathbf{v} = h_x v_x + h_y v_y + h_z v_z. \quad (4)$$

For a radar scan with N detections, this yields a linear system solved by least squares:

$$\begin{bmatrix} v_{D,1} \\ \vdots \\ v_{D,N} \end{bmatrix} = \begin{bmatrix} h_{1x} & h_{1y} & h_{1z} \\ \vdots & \vdots & \vdots \\ h_{Nx} & h_{Ny} & h_{Nz} \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix}. \quad (5)$$

For holonomic ground robots, roll (ϕ) and pitch (θ) data from the IMU is used to remove one DOF, leaving the following expression:

$$\mathbf{v} = \begin{bmatrix} \cos(\theta)v_x \\ \cos(\phi)v_y \\ \sin(\theta)v_x + \sin(\phi)v_y \end{bmatrix}. \quad (6)$$

This formulation improves robustness by accounting for feasible velocities and reduces noise in the vertical component, where radar data is particularly noisy due to low resolution and reduced FoV. However, because multirotors can exhibit significant velocity changes in all three directions, all three linear velocity components need to be estimated independently.

Finally, the estimated linear velocity is used to initialize registration-based scan-matching of the filtered pointclouds with Generalized ICP (GICP) [20]. The concatenation of estimated relative transforms can be used as a standalone odometry system, or fused with additional measurements in a larger localization problem.

C. Pose Graph Optimization

Pose graphs formulations constitute a widely documented approach to localization and SLAM problems [21]. A pose graph is composed of robot positions, which are represented as nodes, and observations coming from sensor measurements, which are represented as edges or factors. Edges are often relative pose measurements that serve as constraints between two consecutive nodes, in the case of odometry, or two arbitrary nodes, in the case of loop closing. Solving a pose graph consists of maximizing the posterior density $p(x|z)$ of states $\mathbf{x}$ given the relative measurements coming from observations $\mathbf{z}$. If noise models are assumed to be Gaussian, solving this problem is equivalent to solving nonlinear least squares of the form:

$$X^* = \arg \min_{\mathbf{x}} \sum_{i,j} \|\mathbf{z}_{ij} - \hat{\mathbf{z}}_{ij}(\mathbf{x}_i, \mathbf{x}_j)\|_{\Sigma_{ij}}^2 \quad (7)$$

where $\mathbf{z}_{ij}$ is a relative measurement between two nodes and $\hat{\mathbf{z}}_{ij}(\mathbf{x}_i, \mathbf{x}_j)$ is the predicted measurement for a configuration of those two nodes, given that the data association is providing a consistent estimate.

Pose Graphs can be extended to multi-robot systems with a few considerations. In this paper, we adopt the concept of encounters and anchor nodes introduced in [13]. Encounters are edges that connect the poses of two robots, and anchor nodes define the transformation between the local odometry frame of each robot and a common frame of reference. In this way, robot poses are expressed in their own local frames, and encounter constraints involve the anchor nodes to express the inter-robot measurement in a common reference frame. To reduce the problem to 4-DOF, each node is a 4D state $\mathbf{x} = [x, y, z, \theta]^T$ and we can make use of the relative error formulation presented in (3). This yields the following expression:

$$X^* = \arg \min_{\mathbf{x}} \sum_W \left[\sum_{r=1}^R (\phi^r + \sum_{i=1}^{N_r-1} \|\mathbf{e}(Z_{i,i-1}^r, \hat{Z}_{i,i-1}^r)\|_{\Sigma_{i,i-1}}^2) + \sum_{i,j \in C} \|\mathbf{e}(Z_{ij}, \hat{Z}_{ij})\|_{\Sigma_{ij}}^2 + \phi^1 \right] \quad (8)$$

The first term ϕ^r is a prior that anchors each one of the R robots local trajectory to a starting location, which is set to

the origin of its respective odometry frame. The second term groups the odometry constraints for each robot r (observed relative transformation Z^r and predicted relative transformation $\hat{Z}^r$) between the consecutive poses of each robot along their respective trajectories, for a total of N_r relevant poses or keyframes. The third term refers to encounter constraints, represented by the C set. In this term, $\hat{Z}_{ij}$ is the relative transformation between pose T_j of robot r , and pose T_i of robot r' , which is compared to the observation Z_{ij} . Finally, ϕ^0 is a prior for one of the anchor nodes and addresses the gauge freedom in the global trajectory, which is set to align with the odometry frame of robot 1, without loss of generality.

The optimization is performed over a spatial window W of N_r keyframes, which are added every 0.5 meters of trajectory for each of the robots.

In our setup, odometry constraints come either from radar, the wheel encoders in the case of the ground platform, or IMU integration regarding the UAV. As for inter-robot constraints, we make use of the relative estimation from UWB measurements introduced in III-A, which can be used to infer a virtual inter-robot observation Z_{ij} based on the odometry frame alignment provided by the optimization.

IV. EXPERIMENTS

The experiments were carried out with a holonomic UGV with a maximum payload of over 100 kilograms, and a fully autonomous DJI M210 UAV. Figure 2 highlights the location of the sensors onboard both of the robots, including four Eliko UWB anchors on the UGV and two Eliko UWB tags on the UAV platform which have been replicated in simulation, as well as two ARS548RDI radars. Notice that the radar in the aerial platform is oriented backwards and tilted to point below the robot, while the radar on the UGV points forward in the same direction as the robot. Finally, both robots are equipped with OS1-32 3D LiDAR sensors, which are used to obtain the ground truth trajectories in the real setting by applying the Direct LiDAR Localization [22] (DLL) method.

A. Ultra-Wideband characterization

A custom simulator plugin has been implemented to produce realistic range measurements on simulation. This plugin incorporates zero-mean Gaussian noise, affecting all measurements, and a small probability of missing samples, simulating communication issues.

We analyzed the range measurements from two of the anchors to both tags in the real system, with the objective of adjusting the virtual measurement to model real-life behavior (see Figure 3). The Mean Bias Error (MBE), the standard deviation σ and Root Mean Squared Error (RMSE) of the real data were computed to serve as a baseline, as seen in Table I. It was found that the average standard deviation across all pairs was under 10 centimeters in a segment when the robots were static, and slightly below 12 centimeters when they were in motion. Therefore, the standard deviation adopted for our measurement model in (1) is $\sigma = 12$ cm. Additionally, as the captured data shows a close to zero




Fig. 2: (a) and (c): Top and lateral view of the UGV used in the experiments. (b) and (d): Top and lateral view of the Matrice M210 used in the experiments. UWB sensors are highlighted in purple, and the radar sensors are highlighted in red.

Anchor / Tag Pair	MBE [cm]	σ [cm]	RMSE [cm]
Anchor 1 / Tag 1	8.69	12.55	15.08
Anchor 1 / Tag 2	-9.51	11.63	14.94
Anchor 2 / Tag 1	-0.04	9.00	8.84
Anchor 2 / Tag 2	-6.83	10.93	12.78

TABLE I: MBE, standard deviation, and RMSE of UWB ranges characterized on the real system.




Fig. 3: Measurements collected during the starting segment of the dataset trajectory for each pair, as compared with ground truth ranges.

average bias across pairs under LOS conditions, we decided to model the measurements with zero bias.

Finally, the geometric distribution of tags and anchors was configured to be identical to the real-world experiments.

B. Validation framework

The evaluation of our system has been carried out with the implementation shown in Fig. 4 and includes two variants: the SITL simulation environment (see Fig. 4a) and the system


developed for real-world experiments (see Fig. 4b). Both variants share interfaces with the relative transformation estimation node and the multi-robot pose-graph node, which are the main components of our system. The simulated version includes a custom simulator plugin to generate realistic virtual UWB measurements and an intermediate module that interfaces with PX4, parsing odometry into ROS 2 standard messages and commanding the robots to follow pre-generated trajectories. First, we evaluated the system in simulation without radar odometry, which is then added in the real-world setup as an additional constraint to the pose-graph framework.

C. Simulation


Two cooperative exploration missions were designed. In both cases, the robots start their trajectories close together (≈ 1 meter separation). In Mission 1, the robots describe a parametric trajectory in coordinated motion, keeping a close distance between one another. In Mission 2, the robots perform two different trajectories of varying lengths and speeds, with the aerial vehicle describing a rectangular motion and the UGV describing a circumference. This experiment aims to test the system when the robots are far apart from each other (≈ 10 meters) and then get close together again. In addition to this, we modify the odometry output of both platforms, as read from the SITL, to include a 2% drift that degrades the estimation as the robots advance.

To evaluate the 4-DOF relative transformation optimization described in Section III-A, we computed the RMSE of the angular and translational components of the transformation error and compared the results with the baseline linear method from the authors of [6], using the same standard deviation of 2 centimeters as in the original work. We selected Mission 1 as a reference, a use case where the robots stay close together. Figure 5 compares the solution obtained by the original method, as well as our method with and without the linear solution used as prior, represented as T_p in (2). Our standalone method without an initial guess achieves a mean translational RMSE of 0.57 m and rotational RMSE of 22.9° , compared to 0.71 m and 35.5° of the linear method. Furthermore, our method provides more consistent estimates with higher convergence rates when compared to the linear method. Regarding the contribution of the linear prior to our method, there is a small benefit from the linear prior in the initial 25 seconds. However, in some cases it can undermine the performance of our method (from 115 to 125 seconds). Consequently, our standalone method is able to provide a high rate, robust estimation of the relative transform.

Next, we evaluate the performance of our system in Missions 1 and 2 using our modeled dispersion of range measurements ($\sigma = 12$ cm). Figure 6 shows the reference trajectories as well as the final optimized poses for both simulated missions, with the UAV poses already transformed by the anchor nodes. In order to evaluate the algorithms, we compute the Absolute Trajectory Error (ATE) in translation, and the RMSE is used to compute the errors in the yaw angle. Table II shows the results of the final pose estimation



(a) SITL simulation architecture.



(b) System architecture for real experiments.

Fig. 4: Overall architecture of the system.




Fig. 5: Comparison of the RMSE of the estimated relative transform in the three evaluated methods with $\sigma = 2$ cm.

for both platforms in both missions. It shows that our system is capable of estimating the global pose of both platforms even without knowledge of the initial transformation between robots. As expected, the errors on Mission 2 are higher due to DoP induced by the growing distance between the robots at the central segment of the trajectory.

D. Dataset

The dataset experiment includes a UGV trajectory that spans 27 m in the horizontal plane and includes frequent fast heading changes, and a UAV trajectory which is 118 meters long, moving in 3D space with overall smoother motions.

Figure 7 indicates the rotational and translational components of the RMSE of the estimated inter-robot transfor-




Fig. 6: Optimized poses as compared to ground truth data in Mission 1 (a) and Mission 2 (b). UGV trajectories are highlighted in blue, and UAV trajectories in red. The baseline is represented as a continuous line, and the dots represent the optimized poses.

TABLE II: Translational ATE in meters and RMSE of the angular error of the poses of the UAV and UGV after pose-graph optimization.

Experiment	Robot	ATE (m)	RMSE (°)
Mission 1	UAV	0.98	25.9
	UGV	0.15	8.59
Mission 2	UAV	1.51	27.51
	UGV	0.02	8.21
Dataset	UAV	1.44	5.74
	UGV	0.94	11.4

mation. Approximately after 40 seconds, when both of the robots have advanced enough, the error quickly converges





Fig. 7: RMSE of the estimated relative transform with respect to ground truth from DLL, showing rotational and translational components.

to ≈ 1.1 m in translation and $\approx 5^\circ$ in rotation. Considering measurement error, which can be on the order of tens of centimeters at these operating ranges, and odometry error (in this case, only wheel and IMU odometry are being fed to the optimizer, both of which tend to accumulate significant drift), the results obtained for real robots are favorable and consistent with what was obtained in simulation, which can also be confirmed by analyzing the errors of the global pose estimation (see Table II).


Figure 8 represents the linear velocities estimated by the radar ego-motion algorithm, smoothed with a 10 sample moving average filter. Ground truth linear velocities are estimated by computing the derivative of the DLL pose estimation, and similarly smoothing the result. In this experiment, despite its holonomic configuration, the robot moves predominantly in the forward direction, where the radar is able to capture significant changes in velocity (≈ 0.5 m/s). In the UAV case, which proves more challenging as there is individual motion in the three directions, the method is also able to provide a consistent estimate.

Figure 9 shows the final optimized trajectories of the UAV and the UGV in 3D. As in the simulation scenario, the UGV trajectory is anchored to the map while the UAV trajectory is not, which means the relative localization error is manifested on the latter. Specifically, the bias in the Y direction of the UAV trajectory matches the errors found in the relative transformation (≈ 1.1 m). Additionally, it was found that radar scan-matching frequently produced inconsistent translations in the vertical direction that contaminated the estimation, which was compensated by setting a higher variance in this direction to the relative transformation constraint computed by GICP. This is a well-known issue in radar-based systems due to the low vertical resolution of the sensor [3], and is most notably seen in the UAV trajectory.

Finally, performance metrics for all optimization experiments are provided in Table III. All experiments were performed in real time from a stream of data stored in a ROS 2 bag file, and run on a Gigabyte AORUS 5 KE laptop with a 12th Gen Intel® Core™ i7-12700H CPU and 32GB of RAM.



(a) UGV



(b) UAV

Fig. 8: Radar-estimated velocities (blue) as compared with DLL velocities (red) in the UGV (a) and the UAV (b).

The temporal metric is provided by Ceres and corresponds to the total time spent in the solver. The relative transformation estimation is overall less demanding due to mostly scalar residuals. In the dataset, it takes under 200 ms with a variance in the order of 100 ms. Conversely, pose-graph optimization shows more variance across experiments, which can be attributed to the variable number of 4D residuals being included in the optimization problem depending on the availability of sensor measurements. The average time per solve in the case of the dataset implies that in real-world operation the system would be able to run at rates above 5 Hz, as both processes execute concurrently, making it suitable for real-time localization.

V. CONCLUSIONS

This work showed the potential of radio-based localization methods in GPS-denied environments where LiDAR or vision may fail - e.g., in poorly lit scenes or in places with




Fig. 9: 3D representation of the global optimized poses with anchor transformation.

TABLE III: Average and standard deviation of the execution time of the optimization processes for the RTE and Pose-Graph optimization problems. The number of runs is indicated in parentheses.

Experiment	RTE (s)	Pose-Graph (s)
Mission 1	0.31 ± 0.18 (352)	0.77 ± 0.74 (119)
Mission 2	0.31 ± 0.15 (168)	0.44 ± 0.48 (94)
Dataset	0.140 ± 0.141 (50)	0.118 ± 0.238 (226)

harsh weather conditions. A cooperative framework combining UWB and radar was proposed to localize a ground and an aerial platform. UWB anchor-tag measurements enabled trajectory alignment over short distances, reliably estimating translation and yaw even without an initial guess. However, performance degraded at longer ranges due to limited transceiver geometry and DoP. Alternatively, radar odometry estimates were obtained via ego-motion estimation and scan-matching, with Doppler velocity allowing consistent linear velocity estimation and aiding scan-matching initialization. These estimates were fused in a multi-robot pose-graph optimization with proprioceptive odometry, helping reduce drift and enforcing global consistency through anchor nodes and encounter constraints. Simulation and dataset results showed robust performance, with relative localization errors of approximately 1.2 % on long, complex UAV trajectories.

However, although our SITL framework models the dispersion of range measurements from real data, a more exhaustive analysis of the bias, including behavior under Non Line Of Sight (NLOS) conditions, is left for future work. Additionally, this system does not implement explicit place recognition, which leads to growing uncertainty over time. A promising direction is extending this approach to a fully-fledged multi-robot, radio-based SLAM system, which our factor graph formulation naturally supports. However, radar place recognition remains challenging due to narrow FoV and low vertical resolution, which could be overcome by

introducing more involved feature extraction methods such as learning-based embeddings [23].

REFERENCES

- [1] G. Brooker, R. Hennessy, C. Lobsey, M. Bishop, and E. Widzyk-Capehart, "Seeing through dust and water vapor: Millimeter wave radar sensors for mining applications," *Journal of Field Robotics*, vol. 24, p. 527–557, 2007.
- [2] J. Starr and B. Lattimer, "Evaluation of navigation sensors in fire smoke environments," *Fire Technology*, vol. 50, p. 1459–1481, 2014.
- [3] L. C. Elena, F. Caballero, and L. Merino, "Loosely coupled 4D-Radar-Inertial odometry for ground robots," *arXiv preprint arXiv:2411.17289*, 2024.
- [4] N. Trawny, X. S. Zhou, K. Zhou, and S. I. Roumeliotis, "Interrobot transformations in 3-D," *IEEE Transactions on Robotics*, vol. 26, no. 2, pp. 226–243, 2010.
- [5] N. Trawny, X. S. Zhou, and S. I. Roumeliotis, "3d relative pose estimation from six distances," in *Robotics: Science and systems*, 2009.
- [6] F. Molina Martel, J. Sidorenko, C. Bodensteiner, M. Arens, and U. Hugentobler, "Unique 4-dof relative pose estimation with six distances for UWB/v-slam-based devices," *Sensors*, vol. 19, no. 20, p. 4366, 2019.
- [7] T. H. Nguyen and L. Xie, "Relative transformation estimation based on fusion of odometry and UWB ranging data," *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2861–2877, 2023.
- [8] Z. Cao, R. Liu, C. Yuen, A. Athukorala, B. K. K. Ng, M. Mathanraj, and U.-X. Tan, "Relative localization of mobile robots with multiple ultra-wideband ranging measurements," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5857–5863.
- [9] S. H. Cen and P. Newman, "Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6045–6052.
- [10] M. Holder, S. Hellwig, and H. Winner, "Real-time pose graph SLAM based on radar," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 1145–1151.
- [11] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer, "Instantaneous ego-motion estimation using doppler radar," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 869–874.
- [12] J. Zhang, H. Zhuge, Z. Wu, G. Peng, M. Wen, Y. Liu, and D. Wang, "4DRadarSLAM: A 4D imaging radar SLAM system for large-scale environments based on pose graph optimization," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 8333–8340.
- [13] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller, "Multiple relative pose graphs for robust cooperative mapping," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 3185–3192.
- [14] I. Deutsch, M. Liu, and R. Siegwart, "A framework for multi-robot pose graph SLAM," in *2016 IEEE International Conference on Real-time Computing and Robotics (RCAR)*. IEEE, 2016, pp. 567–572.
- [15] 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 *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 593–600.
- [16] V. Brunacci, A. Dionigi, A. De Angelis, and G. Costante, "Infrastructure-less UWB-based active relative localization," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 14 079–14 086.
- [17] M. Shalihan, Z. Cao, B. P. L. Lau, R. Liu, C. Yuen, and U.-X. Tan, "MR-FLOUR: Multi-robot relative localization based on the fusion of LiDAR, odometry, and UWB ranging," *Mechatronics*, vol. 112, p. 103410, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957415825001199>
- [18] M. Kamal Mazhar, M. J. Khan, A. I. Bhatti, and N. Naseer, "A novel roll and pitch estimation approach for a ground vehicle stability improvement using a low cost IMU," *Sensors*, vol. 20, no. 2, 2020. [Online]. Available: <https://www.mdpi.com/1424-8220/20/2/340>
- [19] J. Sola, J. Deray, and D. Atchuthan, "A micro lie theory for state estimation in robotics," *arXiv preprint arXiv:1812.01537*, 2018.

- [20] A. Segal, D. Haehnel, and S. Thrun, "Generalized-ICP," in *Robotics: science and systems*, vol. 2, no. 4. Seattle, WA, 2009, p. 435.
- [21] G. Grisetti, R. Kümmerle, C. Stachniss, and W. Burgard, "A tutorial on graph-based SLAM," *IEEE Intelligent Transportation Systems Magazine*, vol. 2, no. 4, pp. 31–43, 2010.
- [22] F. Caballero and L. Merino, "DLL: Direct LiDAR Localization. A map-based localization approach for aerial robots," in *2021 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2021, pp. 5491–5498.
- [23] G. Peng, H. Li, Y. Zhao, J. Zhang, Z. Wu, P. Zheng, and D. Wang, "TransLoc4D: Transformer-based 4D Radar Place Recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 17 595–17 605.