

Lawrence Berkeley National Laboratory

LBL Publications

Title

MURP: Multi-Agent Ultra-Wideband Relative Pose Estimation With Constrained Communications in 3D Environments

Permalink

<https://escholarship.org/uc/item/2gr034mv>

Journal

IEEE Robotics and Automation Letters, 9(11)

ISSN

2377-3766

Authors

Fishberg, Andrew

Quiter, Brian J

How, Jonathan P

Publication Date

2024

DOI

10.1109/lra.2024.3451393

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

MURP: Multi-Agent Ultra-Wideband Relative Pose Estimation With Constrained Communications in 3D Environments

Andrew Fishberg^{1b}, Brian J. Quiter^{1b}, and Jonathan P. How^{1b}, *Fellow, IEEE*

Abstract—Inter-agent relative localization is critical for many multi-robot systems operating in the absence of external positioning infrastructure or prior environmental knowledge. We propose a novel inter-agent relative 3D pose estimation system where each participating agent is equipped with several ultra-wideband (UWB) ranging tags. Prior work typically supplements noisy UWB range measurements with additional *continuously* transmitted data (e.g., odometry) leading to potential scaling issues with increased team size and/or decreased communication network capability. By equipping each agent with multiple UWB antennas, our approach addresses these concerns by using only *locally* collected UWB range measurements, *a priori* state constraints, and *event-based* detections of when said constraints are violated. The addition of our learned mean ranging bias correction improves our approach by an additional 19% positional error, and gives us an overall experimental mean absolute position and heading errors of 0.24 m and 9.5° respectively. When compared to other state-of-the-art approaches, our work demonstrates improved performance over similar systems, while remaining competitive with methods that have significantly higher communication costs.

Index Terms—Distributed robot systems, multi-robot systems, range sensing, robotics in under-resourced settings, swarm robotics.

I. INTRODUCTION

WITHIN the last decade, ultra-wideband (UWB) has matured into a reliable, inexpensive,¹ and commercially available RF solution for data transmission, relative ranging, and localization. For robotics, UWB has several properties of note: precision of 10 cm, ranges up to 100 m, resilience to multipath, operates in non-line of sight (NLOS) conditions, low

Received 31 March 2024; accepted 26 July 2024. Date of publication 28 August 2024; date of current version 16 October 2024. This article was recommended for publication by Associate Editor X. Yu and Editor M. Ani Hsieh upon evaluation of the reviewers' comments. This work was supported by the US Department of Energy, National Nuclear Security Administration, Office of Defense Nuclear Nonproliferation Research and Development (DNN R&D). MIT's contributions were supported by DOE, NNSA, and ALB under Award DE-NA0003921 and Award AWD-000372-G2. LBNL's contributions were supported by the US Department of Energy under Contract DE-AC02-05CH11231. (Corresponding author: Andrew Fishberg.)

Andrew Fishberg and Jonathan P. How are with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: fishberg@mit.edu; jhow@mit.edu).

Brian J. Quiter is with the Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA (e-mail: bjquiter@lbl.gov).

Our datasets and videos: <https://github.com/mit-acl/murp-datasets>
Digital Object Identifier 10.1109/LRA.2024.3451393

¹At time of writing, UWB sensors sell for ~ \$10 as an integrated circuit and between \$20-\$100 as a development board or plug-and-play sensor.




Fig. 1. Diagram of proposed system. Here three agents fly at different altitudes while performing real-time 3D relative pose estimation. Each agent is equipped with six ultra-wideband (UWB) antennas, each capable of performing pairwise relative ranging between all other agents' individual antennas. By using trilateration, an improved sensor model, and *a priori* state constraints about altitude/roll/pitch, agents can perform *instantaneous* estimation entirely with *locally* collected UWB measurements (i.e., without the need to *continuously* transmit other measurements, such as odometry). Additionally, each agent *locally* monitors its *a priori* constraints via downward facing LiDAR and IMU, enabling an *event-based* communication model that only transmits if said assumptions change or are violated. The pictured drone is used in the experiments outlined in Section V-A.

power consumption, and 100 Mbit/s communication speeds [1]. Recent devices even extend the recommended and operational ranges to 300 m and 500 m respectively [2]. Nevertheless, UWB measurements are not immune from ranging errors or noise (see Section III-A), the modeling and correction of which is an active area of research [3], [4], [5], [6].

A common approach in UWB relative localization work fuses noisy UWB ranging measurements with additional *continuously* transmitted data, such as odometry [7] and visual inter-agent tracks [8], [9]. While these approaches achieve low absolute position error (APE) and absolute heading error (AHE), there are two prevalent shortcomings: (1) They often use a simplistic UWB measurement noise model (i.e., zero mean Gaussian), which then requires the use of supplementary measurements to compensate. (2) They rely on these supplementary




Fig. 2. Plots demonstrating the UWB noise and geometry characteristics outlined in Section III. (a) Histogram of our entire set of UWB measurements binned by range error. Demonstrates a non-zero mean and long tail (i.e., violates the zero mean Gaussian assumption that is typically used). (b) Same data plotted as error with respect to relative elevation showing that the measurement error’s mean and variance change significantly with relative elevation. The dotted blue line represents a learned 6-degree polynomial fit of measurement bias. (c) Simple example demonstrating dilution of precision (DOP) in a 2D ranging scenario. (d) Demonstrates how three (or more) ranging antennas within a single plane produce a pair of ambiguous solutions (i.e., if all antennas are in the base’s plane $z = 0$, while the target’s x and y coordinates are fully observable, the target’s altitude has an ambiguity between $\pm z$).

measurements (often not *locally*² collected, e.g., odometry), which mandates their *continuous* transmission between agents and impacts system scalability with increased swarm size or decreased communication throughput.

Our previous work [10] used UWB to demonstrate an *instantaneous*³ multi-tag approach to relative 2D pose estimation that achieved superior mean position accuracy and competitive performance on other metrics to Cao et al. [7] (the most comparable state-of-the-art work) using only *local* UWB measurements. These results were achieved by a trilateration nonlinear least squares (NLLS) optimization problem that, leveraging an improved UWB sensor model, accounted for various sources of measurement error.

This letter extends our prior 2D work [10] to an *instantaneous* multi-tag relative pose estimate in 3D environments (Fig. 1). Similar to our previous work, a key objective is to develop an approach that minimizes the communication load (i.e., we are willing to trade absolute accuracy for reduced communication). In 3D environments, each agent’s altitude/roll/pitch are *locally* measurable with respect to a common level-ground world frame (e.g., via altimeter or downward facing LiDAR and IMU). By constraining each agent to an altitude/roll/pitch envelope, this information only needs to be transmitted once *a priori*. These assumptions can then be *locally* monitored by each agent, with *event-based* communication only occurring when an agent changes or violates its constraint envelope (Fig. 3(b)). This formulation allows us to simplify the 6-DoF optimization (i.e., 3D position and orientation) to a 3-DoF optimization (i.e., x , y , yaw with known z , roll, pitch) that, with fewer free variables, more reliably produces 3D localization solutions. Thus our approach provides a 3D relative pose using only *locally* collected UWB range measurements, *a priori* state constraints, and detections of constraint violations.

This paper’s contributions are: (1) An in-depth analysis and modeling of the observed noise characteristics of 3D UWB ranging measurements (Section III). (2) A 3D *instantaneous* solution for UWB-based relative localization with minimal communication (i.e., only requires *event-based* communication when a *locally* monitored assumption is violated, which should occur infrequently or never). Our solution is formulated as a robust nonlinear least squares optimization with a learned measurement

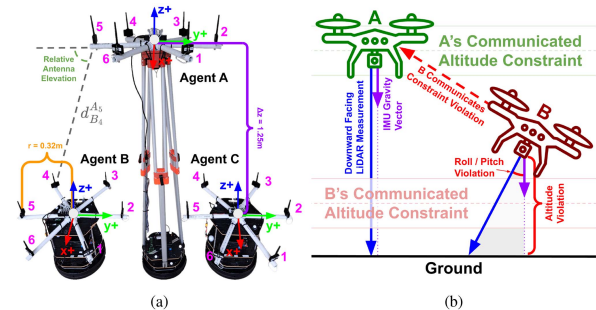


Fig. 3. (a) Annotated diagram of three UGV agents used in Section V-B. UGV agents are designed with the same baseline as the UAV in Fig. 1, making them comparable surrogates to the UAV and its experiments in Section V-A. (b) Diagram demonstrating agent *B* *locally* detecting violations in its *a priori* constraints, triggering an *event-based* communication. Other agents will exclude *B* from their relative pose estimation until *B* notifies the swarm of restored constraint status or provides a new constraint envelope.

bias correction as a function of relative elevation, improving our mean APE by 19% (Section IV). (3) Hardware experimental results that demonstrate the merits of our proposed solution, **a mean APE and AHE of 0.24 m and 9.5° respectively without continuously transmitting measurements** (Section V). (4) Our datasets, over 200 hours of pairwise UWB range measurements with ground truth.⁴

II. RELATED WORKS

UWB technology can be leveraged by mobile robotics in several orthogonal ways. This work uses UWB for infrastructure-free inter-agent relative measurements – thus, works involving GPS position, static UWB anchors, UWB radar, or measurements done by direct waveform analysis are out of scope of this letter [17], [18], [19] (see survey paper [20] or our previous work [10] for a broader review of UWB in robotics). Table I highlights the unique capability gap this work fills with respect to the recent related literature. While all works in Table I use UWB for inter-agent ranging, it is challenging to compare numerical results given the vastly different operating assumptions and priorities – specifically access to additional sensors (e.g., cameras, LiDAR, etc.), the size of agents with respect to their environment (see Section III-B), and overall communication model (e.g., unlimited vs constrained communication) all significantly change performance and scalability. To contextualize this work, we discuss [13] and [9] since they are the most complete

²“*locally*” refers to a measurement collected onboard, and thus can be used to estimate other agents’ pose without needing to be transmitted.

³“*instantaneous*” implies the target’s position is fully determined (i.e., observable) using only current measurements. As such, this approach does not require a dynamics model or measurements over a period of time. Furthermore, since estimation does not rely on integrating velocity or acceleration, the estimation should not drift/diverge over time.

⁴<https://github.com/mit-acl/murp-datasets>

TABLE I
COMPARISON OF RELEVANT RELATED WORKS TO HIGHLIGHT THE CAPABILITY GAP THIS LETTER FILLS

2D / 3D	Related Work	# UWB Per Agent	UWB Noise Model	Data Exchanged
2D	Guo et al. (2020) [11]	Single	Bounded Error Assumption	UWB + Velocity
	Cao et al. (2021) [7]	Many (4)	NLLS	UWB + Odometry
	Zheng et al. (2022) [12]	Many (2 or 4)	ZM Gaussian + NLOS Rejection	UWB + VIO
	Zhang et al. (2023) [13]	Many (2 or 4)	ZM Gaussian + Outlier Rejection	UWB + Point Clouds + Odom + Keyframes
	[Our Prev. Work] (2022) [10]	Many (4)	NLLS + Data Informed Sensor Model	UWB
3D	Xu et al. (2020) [8]	Single	ZM Gaussian	UWB + VIO + Visual Tracks
	Qi et al. (2020) [14]	Single	ZM Gaussian	UWB + IMU + GPS Heading
	Xianjia et al. (2021) [15]	Many (2 or 4)	NLLS	UWB
	Xu et al. (2022) [9]	Single	ZM Gaussian + Occlusion Rejection	UWB + VIO + Visual Tracks
	Xun et al. (2023) [16]	Single	ZM Gaussian	UWB + IMU + IR Visual Tracks
	[This Work] (2024)	Many (4 or 6)	NLLS + Data Informed Sensor Model	UWB + Event-Based Assumption Violations

Color emphasizes how a work fits in with our goals. Specifically, *green* indicates our needs are met, *yellow* indicates our needs are partially met, and *red* indicates our needs are not met.

and comprehensive 2D and 3D UWB systems respectively, [16] since it has a unique compact hardware solution, and [15] since it is the most directly comparable to this work.

Zhang et al. [13] presents a centralized 2D range-aided cooperative localization and consistent reconstruction system, merging a tightly coupled visual odometry and multi-tag ranging front-end with a pose graph optimization (PGO) back-end. This work differs from ours by being 2D and relying on *continuous* transmission of point cloud, odometry, and keyframe information. Omni-swarm [9] is a decentralized 3D swarm estimation system. This work differs from ours by relying primarily on visual tracks of neighboring agents from each agent’s omnidirectional camera – omitting UWB measurements only degrades their estimation by approximately 0.01 m. Additionally, each agent has only a single UWB tag and requires *continuously* sharing measurements between all pairs of agents, scaling poorly with larger swarms. CREPES [16] presents a custom compact hardware module, which tightly couples an IMU, UWB, IR LEDs, and IR fish-eye camera. The measurements are fused into 6-DoF relative pose via a centralized error-state Kalman filter (ESKF) and PGO, while also having the ability to provide an *instantaneous* estimate from a single frame of sensor measurements. This work differs from ours by relying heavily on visual tracks over UWB (similar to [9]). Additionally, the system relies on *continuously* transmitted measurements. Thus the state of the art approaches in [9], [13], [16] achieve APE and AHE on the order of 0.10 m and 1° respectively, but require additional sensors that must *continuously* transmit measurements, making them scale poorly to scenarios with larger swarm sizes or reduced communication throughput. Furthermore, these required additional sensors (e.g., LiDAR, specialized cameras, etc.) can cost thousands of dollars per agent, whereas UWB sensors (even multiple per robot) can cost $< \$100$ per agent.

Xianjia et al. [15] is the most similar to our work (i.e., 3D environment, multiple UWB tags per agent, and using only UWB measurements). When their agents have similar antenna baselines to our work, they achieve a mean *xy*-positional error of approximately 0.40 m, 0.65 m, and 0.85 m for simulated $8\text{ m} \times 8\text{ m}$ flights at various fixed altitudes. Furthermore, for real flights the mean *xy*-positional error becomes approximately

1 m with a variable altitude and a (beneficial) larger baseline. By comparison, although our experimental environment differs, we achieve a mean APE and AHE of 0.24 m and 9.5° respectively with real experiments operating in a $10\text{ m} \times 10\text{ m}$ mocap space. Ref. [15] differs from our approach in two important ways: 1) the authors model UWB ranging error as a zero mean Gaussian with a 0.10 m standard deviation, which does not reflect real data (see Section III-A); and 2) we define an explicit communication protocol that allows us to address several capability gaps without *continuously* transmitting measurements. Specifically, by constraining altitude/roll/pitch and improving our UWB noise model, we can better address UWB noise (see Section III-A) and an observability ambiguity (see Section III-B and Fig. 2(d)). See Section III-C for details.

III. EXTENSIONS TO 3D ENVIRONMENTS

The following subsections motivate our hardware and algorithmic choices for operating in 3D (Section III-C). We first characterize the UWB ranging error of the Nooploop LinkTrack P-B [2] from a set of initial experiments⁵ (Section III-A), and then provide the observability and error properties of trilateration-based localization (Section III-B).

A. Characterizing UWB Noise

Long Tail Distributions: Contrary to common noise assumptions in Table I, individual UWB ranging errors appear neither zero mean nor Gaussian. Instead, error distributions have long positive tails (see Fig. 2(a)). A trend towards positive bias can be attributed to the many ways positive ranging error can be introduced to a UWB measurement (e.g., multipath or a change of propagation medium). Furthermore, installing SMA cables between RF devices and antennas adds consistent positive bias to all measurements.

Sensitivity to Obstruction: Antenna obstruction (i.e., NLOS conditions) introduces positive bias and increased variance to collected ranging measurements (see Fig. 4(b)). While the UWB protocol is resilient to multipath and NLOS, UWB ranging is not completely absolved of these concerns. Furthermore, auxiliary metrics, like RSSI, do not appear to meaningfully indicate an obstructed measurement. Refs. [9], [10], [12], [13] address this by rejecting suspect measurements, whether detected via statistical tests, robust loss functions, or hardcoded rejection criteria.

Dependence on Relative Pose: Measurement noise is dependent on relative pose between antennas. The interplay of antenna attenuation patterns is a core concern of RF designers, but gets understandably overlooked by many end users. The LinkTrack P series comes equipped with a standard dipole antenna which we would expect to produce the behavior observed in Fig. 2(b) when aligned upright [2] – specifically, approximately uniform performance within an *xy*-plane (i.e., varying azimuth) and degraded performance outside the ground plane $z = 0$ (i.e., non-zero elevation). We note that both the mean bias and variance change with elevation. A relationship to elevation is noted in [9], although it is attributed to NLOS conditions similar to

⁵While these effects should generalize to other UWB ranging devices, the “black box” nature of commercial sensors means the exact manifestations could differ between products. Most works in Table I use a Nooploop LinkTrack P model [2], except for [11] (PulsON 440) and [15] (unspecified). Thus, we expect similar noise characteristics among those related works.

Fig. 4(b), but [9] chooses to address this by simply omitting range measurements with more than 37° relative elevation (feasible only because of their reliance on other measurements).

B. Geometry of Ranging-Based Pose Estimation

Observability of Position: In 3D trilateration, a minimum of 4 non-coplanar antennas are required to uniquely determine a target’s position [21]. That being said, 3 or more planar antennas produce only a pair of solutions that are a reflection across the antenna plane (see Fig. 2(d)). This ambiguity can be resolved with a simple z measurement (i.e., altitude).

Observability of Pose: Common in camera-robot registration, Horn’s Method is a closed-form solution for finding the pose between two Cartesian coordinate systems from a set of corresponding point pairs [22]. Specifically, given exactly 3 planar points in two frames with known correspondences, Horn’s Method fully specifies the 3D translation and orientation. Thus, by using trilateration to measure 3 points in a known configuration on our target, we can recover full pose.

Dilution of Precision: Most commonly associated with GPS, dilution of precision (DOP) is a trilateration sensitivity analysis for quantifying how 1D ranging errors propagate to 2D/3D point estimation error [23]. Specifically, given trilateration’s inherent nonlinearity, even identically noisy 1D range measurements can produce significantly different 2D/3D positional uncertainties – this is entirely determined by the given base station and target geometry (Fig. 2(c)). In general we note that increasing the “baseline” distance between base stations (i.e., the antenna separation on a single vehicle), or having a target closer to the base stations (i.e., decreasing distance between two vehicles), reduces trilateration’s 2D/3D uncertainty. DOP also accounts for the improved accuracy seen by [15] for configurations with larger baselines. Thus, the accuracy of trilateration systems with vastly different baselines or operating ranges can be misleading. See [12], [13] for a detailed discussion of DOP as it pertains to 2D UWB localization.

C. System Design

Based on the noise characteristics (Section III-A) and underlying geometry (Section III-B), we propose the system design shown in Fig. 1 for our multirotor drones. Specifically, each agent is equipped with 6 ranging antennas (attached 0.31 m from the center at the end of each propeller arm) to maximize the baselines and improve DOP geometry. Since there are 6 coplanar antennas, this is an over-constrained problem that still has a $\pm z$ ambiguity across the $z = 0$ plane. Although an additional nonplanar antenna could be added to each agent, it is challenging to get a large z -baseline without impacting the flight characteristics (i.e., heavy/awkward configuration). Instead, we note that altitude/roll/pitch can be measured within the flat-ground world frame *locally* via altimeter or downward facing LiDAR and onboard IMU.⁶ In many drone applications, it is common for multirotors to maintain roll/pitch near 0° and a constant altitude. Thus, we can specify a minimalist communication protocol that shares these intended constraints *a priori* and then monitors them *locally* to ensure they are satisfied. *Event-based* communication

⁶Altitude measurements are direct and should not drift. Substantial drift in roll/pitch w.r.t. gravity would impede basic flight operations before meaningfully impacting our estimation, making it negligible for our system.

TABLE II
THEORETICAL STANDARD DEVIATION OF ABSOLUTE POSITION ERROR BETWEEN TWO OF OUR UAVS (ILLUSTRATED IN FIG. 1 AND SECTION V-A)

StdDev of Abs Pos Error [m]			Evaluated at rel. position [5 m, 0 m, Δz] where Δz equals:					
name	local only	alt. assumpt.	0.0 m	1.0 m	2.5 m	5.0 m	10.0 m	25.0 m
no_assumpt	✓	n/a	∞	9.98	2.05	1.35	2.81	14.38
extr_conserv	✓	± 2.00 m	4.02	2.89	1.43	1.21	2.66	13.73
very_conserv	✓	± 1.00 m	1.08	1.00	0.84	1.00	2.38	12.48
conserv	✓	± 0.20 m	0.39	0.41	0.49	0.79	1.96	10.21
proposed	✓	± 0.10 m	0.39	0.41	0.49	0.78	1.96	10.17
cont_shared		± 0.04 m	0.39	0.41	0.49	0.78	1.96	10.16

Computation uses a UWB ranging uncertainty of $\sigma = 0.24$ m (see Figure 2a). Altitude assumptions are the standard deviation of the target’s assumed altitude when *local only*, and ± 0.04 m (the advertised precision of the downward LiDAR) when *continuously* sharing measurements. Since the system is nonlinear, table columns show evaluation at several relative altitudes. Table rows show how different relative altitude assumptions affect uncertainty. Results demonstrate the advantage of relative altitude assumptions and shows how said assumptions enable comparable uncertainty with only *locally* collected measurements vs systems that *continuously* share altitude measurements. See Section III-C for additional discussion.

would then only occur if a constraint is *locally* detected to have been violated/changed (Fig. 3(b)).

Table II summarizes the theoretical uncertainty of the proposed system and demonstrates the value of adding relative altitude assumptions to a communication constrained system.⁷ Although assumed altitude bounds vary with the platform, we see even *overly* conservative assumptions (e.g., ± 2.0 m) still greatly improves performance over no assumption (no_assumpt vs extr_conserv or very_conserv). This is especially true for level or near-level flights (i.e., unobservable at $\Delta z = 0$, high sensitivity at $\Delta z \approx 0$), a common formation for UAV systems. By comparing the positional uncertainty of our proposed assumption (i.e., proposed, experimentally selected using Fig. 4(a)), a twice as lenient assumption (i.e., conserv), or *continuously* communicated altitude measurements (i.e., cont_shared, using the advertised downward LiDAR accuracy of ± 0.04 m) we see our minimal communication protocol does not substantially degrade uncertainty while using only *locally* collected measurements.

IV. TECHNICAL APPROACH

A. Pose Parameterization

Given a set of 3D reference frames \mathcal{F} , consider any pair of frames $A \in \mathcal{F}$ and $B \in \mathcal{F}$. Let $\mathbf{T}_B^A \in SE(3)$ be the relative pose from frame A to frame B . We note \mathbf{T}_B^A is a 6-DoF value that can be equivalently parameterized as the 2-tuple $(\mathbf{R}_B^A, \mathbf{t}_B^A)$ or 6-tuple $\langle x_B^A, y_B^A, z_B^A, \alpha_B^A, \beta_B^A, \gamma_B^A \rangle$, where $\mathbf{R}_B^A \in SO(3)$ is the relative rotation between A and B parameterized by relative roll α_B^A , pitch β_B^A , and yaw γ_B^A and $\mathbf{t}_B^A \in \mathbb{R}^3$ is a relative translation vector between A and B parameterized by relative x_B^A , y_B^A , and z_B^A .

B. Local Robot Definitions

Let there be some 3D world frame W defined with respect to the level ground and gravity vector (i.e., $z = 0$ is the “floor” and gravity faces down). Consider \mathcal{R} , a set of $N_{\mathcal{R}}$ robots operating in W . Each robot $R \in \mathcal{R}$, has N_R relative ranging antennas rigidly affixed to R ’s body, where R_k denotes R ’s k th antenna. Additionally, for each R let there be a *static a priori* 3-tuple of known information $\mathcal{I}_R = \langle \hat{\mathbf{r}}_R, \hat{\mathbf{r}}_R, \mathcal{P}_R \rangle$, where:

⁷See explanation and code in the GitHub repository for more information on how the theoretical values in Table II were computed.

Algorithm 1: Procedure for robot $A \in \mathcal{R}$, a member of robot swarm \mathcal{R} , to generate instantaneous relative poses w.r.t. all other agents in swarm $B \in (\mathcal{R} \setminus \{A\})$. A communicates violated assumptions with $\text{transmit}(A)$. A locally tracks which agents have communicated violated assumptions, and checks B 's last received status with $\text{invalid}(B)$.

Input: \mathcal{R} , $A \in \mathcal{R}$, $\{\mathcal{I}_R \forall R \in \mathcal{R}\}$, ℓ , $\bar{d}_{B_j}^{A_i}$
Output: $\{\mathbf{T}_B^A \forall B \in (\mathcal{R} \setminus \{A\}) \text{ s.t. } \neg \text{invalid}(B)\}$
1: **function** LOCALPROCEDURE
2: $\hat{\kappa}_A \leftarrow$ local measurements of A w.r.t. W
3: **if** $\neg(|\hat{\kappa}_A - \tilde{\kappa}_A| \leq \tilde{\kappa}_A)$ **then**
4: $\text{transmit}(A)$
5: **end if**
6: $\mathcal{T} \leftarrow \{\}$
7: **for** $B \in (\mathcal{R} \setminus \{A\}) \text{ s.t. } \neg \text{invalid}(B)$ **do**
8: $\tilde{\mathbf{d}}_B^A \leftarrow$ local measurements of B w.r.t. A
9: $\langle \hat{z}_B^A, \hat{\alpha}_B^A, \hat{\beta}_B^A \rangle \leftarrow \hat{\kappa}_B^W - \hat{\kappa}_A^W$
10: $\langle x_B^A, y_B^A, \gamma_B^A \rangle \leftarrow$ optimize Eq. (9)
11: $\mathbf{T}_B^A \leftarrow \mathbf{T}(x_B^A, y_B^A, \hat{z}_B^A, \hat{\alpha}_B^A, \hat{\beta}_B^A, \gamma_B^A)$
12: $\mathcal{T} \leftarrow \mathcal{T} \cup \{\mathbf{T}_B^A\}$
13: **end for**
14: **return** \mathcal{T}
15: **end function**

- $\hat{\kappa}_R = [\hat{z}_R^W \quad \hat{\alpha}_R^W \quad \hat{\beta}_R^W]^\top$ is a vector of *constant* global state constraints on R with respect to world frame W . In other words, \hat{z}_R^W is the commanded altitude, while $\hat{\alpha}_R^W$ and $\hat{\beta}_R^W$ are commanded relative roll and pitch.
- $\tilde{\kappa}_R = [\tilde{z}_R^W \quad \tilde{\alpha}_R^W \quad \tilde{\beta}_R^W]^\top$ is a vector of *constant* absolute (i.e., \pm) tolerances on global constraints $\hat{\kappa}_R$.
- $\mathcal{P}_R = \{\mathbf{T}_{R_1}^R, \mathbf{T}_{R_2}^R, \dots, \mathbf{T}_{R_{N_R}}^R\}$ is an ordered set of N_R *constant* relative poses from R 's body frame to R 's k th antenna's frame. From $\mathbf{T}_{R_k}^R$, we can succinctly denote the location of R 's k th antenna in R 's frame as point \mathbf{p}_k^R , where $\tilde{\mathbf{p}}_k^R$ is the equivalent homogeneous point.

Since W is defined with respect to a level ground and gravity vector, the parameters $\langle \hat{z}_R^W, \hat{\alpha}_R^W, \hat{\beta}_R^W \rangle$ – unlike $\langle x_R^W, y_R^W, \gamma_R^W \rangle$ – are *instantaneously* observable and *directly* measurable from R (e.g., via altimeter or downward facing LiDAR and onboard IMU). Thus, the directly measured values $\tilde{\kappa}_R = [\tilde{z}_R^W \quad \tilde{\alpha}_R^W \quad \tilde{\beta}_R^W]^\top$ can be continuously monitored by R locally (i.e., without relying on swarm communication). After an initial transmission of \mathcal{I}_R to all agents \mathcal{R} , agent R only needs to transmit again if (1) it wants to change its commanded altitude/roll/pitch $\hat{\kappa}_R$, or (2) it locally observes a violation of its previously communicated constraints, specifically:

$$\underbrace{|\hat{\kappa}_R - \tilde{\kappa}_R|}_{\text{element-wise abs deviation}} \leq \underbrace{\tilde{\kappa}_R}_{\text{constraint tolerance}} \quad (1)$$

where $|\cdot|$ and \leq are performed element-wise. In other words, all agents \mathcal{R} can treat $\hat{\kappa}_R$ as *known* states during estimation, unless R communicates otherwise.

C. Inter-Agent Robot Definitions

Consider any arbitrary pair of robots $A \in \mathcal{R}$ and $B \in \mathcal{R}$. With a slight abuse of notation, let A and B also denote the robots respective reference frames, making \mathbf{T}_B^A denote the relative pose between robots A and B . From our known information tuples \mathcal{I}_A and \mathcal{I}_B we know each agent has N_A and N_B antennas at relative poses \mathcal{P}_A and \mathcal{P}_B respectively. This in turn allows us to define our observation model as just a function of \mathbf{T}_B^A . In other words, the distance between A 's i th antenna and B 's j th antenna is defined as:

$$d_{B_j}^{A_i}(\mathbf{T}_B^A) \triangleq \|\mathbf{T}_B^A \tilde{\mathbf{p}}_j^B - \tilde{\mathbf{p}}_i^A\|_2 \quad (2)$$

where $d_{B_j}^{A_i}(\mathbf{T}_B^A) \in \mathbb{R}_{\geq 0}$, as distances are non-negative, and $\|\cdot\|_2$ denotes the L_2 norm. Similarly, let $\tilde{d}_{B_j}^{A_i}$ be the current (i.e., most recent) noisy measurement between A 's i th antenna and B 's j th antenna as measured *locally* by A .⁸ For convenience, $\tilde{\mathbf{d}}_B^A \in \mathbb{R}^{N_A N_B}$ denotes the stacked vector of all current pairwise $\tilde{d}_{B_j}^{A_i}$ measurements.

D. Noise Model Definition

Let $\mathbf{T}_{B_j}^{A_i}$ denote the relative pose between A 's i th antenna and B 's j th antenna. Given a known \mathcal{P}_A and \mathcal{P}_B , and an arbitrary \mathbf{T}_B^A , we have:

$$\mathbf{T}_{B_j}^{A_i} = (\mathbf{T}_{A_i}^A)^{-1} \mathbf{T}_B^A \mathbf{T}_{B_j}^B = \mathbf{T}_A^A \mathbf{T}_B^A \mathbf{T}_{B_j}^B \quad (3)$$

so that $\mathbf{T}_{B_j}^{A_i}$ is a function of \mathbf{T}_B^A given known \mathcal{P}_A and \mathcal{P}_B . Let $\bar{d}(\mathbf{T}_{B_j}^{A_i}) \in \mathbb{R}$ be some provided (e.g., learned) sensor mean bias model as a function of $\mathbf{T}_{B_j}^{A_i}$ (i.e., antenna measurement bias is modeled as some function of the relative position and/or orientation between a pair of ranging antennas). We can then equivalently define:

$$\bar{d}_{B_j}^{A_i}(\mathbf{T}_B^A) \triangleq \bar{d}(\mathbf{T}_A^A \mathbf{T}_B^A \mathbf{T}_{B_j}^B) \quad (4)$$

where $\bar{d}_{B_j}^{A_i}(\mathbf{T}_B^A) \in \mathbb{R}$ is a state-dependent mean bias estimate for measurement $\tilde{d}_{B_j}^{A_i}$.

E. Optimization Definition

Consider the following error residual function between the measured and expected distances for antenna pair A_i and B_j with respect to some \mathbf{T}_B^A , a relative pose between A and B (and our eventual decision variable):

$$e_{B_j}^{A_i}(\mathbf{T}_B^A) \triangleq \underbrace{\left(\tilde{d}_{B_j}^{A_i} - \bar{d}_{B_j}^{A_i}(\mathbf{T}_B^A) \right)}_{\text{bias adjusted measurement}} - \underbrace{d_{B_j}^{A_i}(\mathbf{T}_B^A)}_{\text{observation model}} \quad (5)$$

where $\tilde{d}_{B_j}^{A_i}$ is the current measurements, $\bar{d}_{B_j}^{A_i}(\mathbf{T}_B^A)$ is a measurement bias correction function (current implementation in Section V-B), and $d_{B_j}^{A_i}(\mathbf{T}_B^A)$ the known observation model.

Using our error residual function $e_{B_j}^{A_i}(\mathbf{T}_B^A)$, and some loss function ℓ , we can calculate an instantaneous estimate of \mathbf{T}_B^A by

⁸Note that while $d_{B_j}^{A_i}(\mathbf{T}_B^A) = d_{A_i}^{B_j}(\mathbf{T}_B^A)$ (i.e., the observation model is symmetric), since each agent *locally* and *separately* measures antenna distances, measurement noise makes it that generally $\tilde{d}_{B_j}^{A_i} \neq \tilde{d}_{A_i}^{B_j}$.

minimizing the sum of loss of all error residuals for all inter-agent antenna pairs with respect to \mathbf{T}_B^A . That is:

$$\min_{\mathbf{T}_B^A \in SE(3)} \sum_{i=1}^{N_A} \sum_{j=1}^{N_B} \ell \left(e_{B_j}^{A_i}(\mathbf{T}_B^A) \right) \quad (6)$$

Recall that \mathbf{T}_B^A is a 6-DoF variable parameterized by $\langle x_B^A, y_B^A, z_B^A, \alpha_B^A, \beta_B^A, \gamma_B^A \rangle$. From known information tuples \mathcal{I}_A and \mathcal{I}_B , we have $\hat{\kappa}_A$ and $\hat{\kappa}_B$, which provide us with constraints $\langle \hat{z}_A^W, \hat{\alpha}_A^W, \hat{\beta}_A^W \rangle$ and $\langle \hat{z}_B^W, \hat{\alpha}_B^W, \hat{\beta}_B^W \rangle$. Furthermore, since these constraints are with respect to a common world frame W , see that: $\hat{z}_B^A = \hat{z}_B^W - \hat{z}_A^W$, $\hat{\alpha}_B^A = \hat{\alpha}_B^W - \hat{\alpha}_A^W$, and $\hat{\beta}_B^A = \hat{\beta}_B^W - \hat{\beta}_A^W$. Thus, given $\hat{\kappa}_A$ and $\hat{\kappa}_B$, we have the following constrained optimization:

$$\begin{aligned} \min_{\mathbf{T}_B^A \in SE(3)} & \sum_{i=1}^{N_A} \sum_{j=1}^{N_B} \ell \left(e_{B_j}^{A_i}(\mathbf{T}_B^A) \right) \\ \text{s.t.} & z_B^A = \hat{z}_B^A = \hat{z}_B^W - \hat{z}_A^W \\ & \alpha_B^A = \hat{\alpha}_B^A = \hat{\alpha}_B^W - \hat{\alpha}_A^W \\ & \beta_B^A = \hat{\beta}_B^A = \hat{\beta}_B^W - \hat{\beta}_A^W \end{aligned} \quad (7)$$

which can then be simplified to an equivalent 3-DoF unconstrained optimization problem:

$$\min_{\substack{x_B^A, y_B^A \in \mathbb{R} \\ \gamma_B^A \in [-180^\circ, 180^\circ]}} \sum_{i=1}^{N_A} \sum_{j=1}^{N_B} \ell \left(e_{B_j}^{A_i} \left(\mathbf{T} \left(\underbrace{x_B^A, y_B^A}_{\text{free}}, \underbrace{\hat{z}_B^A, \hat{\alpha}_B^A, \hat{\beta}_B^A}_{\text{constrained}}, \underbrace{\gamma_B^A}_{\text{free}} \right) \right) \right) \quad (8)$$

which can be rewritten as f to clearly see the dependencies:

$$\min_{\substack{x_B^A, y_B^A \in \mathbb{R} \\ \gamma_B^A \in [-180^\circ, 180^\circ]}} f \left(x_B^A, y_B^A, \gamma_B^A \mid \tilde{\mathbf{d}}_B^A, \mathcal{I}_A, \mathcal{I}_B, \ell, \tilde{\mathbf{d}}_B^{A_i} \right) \quad (9)$$

Using (9), Algorithm 1 outlines a procedure independently followed by each agent A in a swarm \mathcal{R} to individually estimate \mathbf{T}_B^A between itself and all other agents $B \in (\mathcal{R} \setminus \{A\})$. Observe that *event-based* communication (i.e., `transmit`) only occurs on Line 4, when an agent *locally* detects it has violated its *a priori* constraints – other agents use this information on Line 7 (i.e., `invalid`) to skip calculating \mathbf{T}_B^A for agents with violated constraints.

V. EXPERIMENTAL RESULTS

We present two sets of experiments: UAV (Section V-A) and UGV (Section V-B). We first verify the 3D noise properties, flight hardware, and show altitude/roll/pitch constraints can be sufficiently *locally* monitored on a UAV. Once this is verified, we attach surrogate UAVs at different altitudes to UGVs – this allows us to efficiently collect large amounts of data to evaluate our approach.


A. UAV Experiments

Hardware: Two agents, A and B , were used. A , a surrogate quadcopter, was a stationary Turtlebot2 UGV with 4 UWB sensors evenly distributed in the xy -plane 0.37 m from the center (i.e., same robot in prior work [10]). B is a custom hexrotor UAV equipped with 6 UWB sensors evenly distributed in the xy -plane 0.31 m from the center (Fig. 1). Additionally, B is equipped with

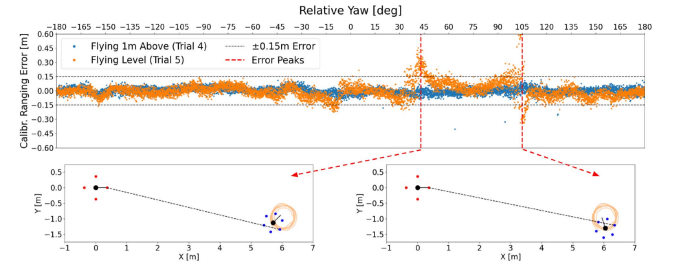
TABLE III
EVALUATION OF POSITIONAL ERROR IN METERS WITH DIFFERENT z -CONSTRAINTS

Abs Position Error [m]		Scenarios (Section V-A)								
		Trial 1		Trial 2		Trial 3				
color	alt constraint	Mean	Max	Std	Mean	Max	Std	Mean	Max	Std
■	z_{free}	1.17	2.83	1.06	1.04	1.53	0.24	0.45	0.94	0.19
■	z_{comm}	0.26	0.45	0.09	0.28	0.87	0.14	0.25	0.64	0.13
■	z_{meas}	0.26	0.44	0.09	0.27	0.87	0.14	0.24	0.66	0.14
■	z_{true}	0.25	0.44	0.10	0.26	0.87	0.15	0.22	0.64	0.14

In practice, the second row (i.e., *a priori* constraints) will be used, but the small difference between z_{comm} , z_{meas} , and z_{true} shows these constraints are not a major source of error. See Section V-A.



(a) Shows Table III's Trial 2 alongside the commanded, measured, and ground truth altitude, where the colors correspond with Table III. While the TeraRanger Evo does not perfectly track ground truth, it stays within the expected commanded tolerance of $\pm 0.10\text{m}$, providing a reasonable *local* measurement of altitude.



(b) Top plot shows calibrated ranging error (i.e., mean bias subtracted) with respect to relative yaw for the *1m Above* and *Level* flights (i.e., Trial 4 and 5 respectively). Observe the error for *1m Above* flight remains consistent, while the *Level* flight has anomalous error peaks, mostly within the $\pm 0.15\text{m}$ range, corresponding to specific relative yaw. The lower left and right diagrams show the relative pose of the agents at the two most prominent peaks (i.e., with error greater than 0.45m), where the path between antennas A_1 and B_5 is eclipsed by antennas B_3 and B_1 respectively. While this NLOS effect is consistent with our previous 2D work [10], we note the more complex 3D relationship resolves the occlusion with varying elevation.

Fig. 4. Plots from the UAV flight experiments. See Section V-A.

a downward facing TeraRanger Evo 15 m LiDAR for altitude measurements [24].

Trials: Five experimental trials were conducted in a large $10\text{ m} \times 10\text{ m}$ mocap space. For all trials, A was placed on a table for additional height and kept stationary, while B flew around. In the first three trials, B took off and flew simple line patterns at different altitudes (see Table III and Fig. 4(a)). For the latter two trials, B flew tight circles in the center of the room – first *1 m Above* and then *Level* with A (see Fig. 4(b)). Each trial had a max speed of 1 m/s and lasted approximately 3–5 min (i.e., the length of a battery charge).

Discussion of Results: Results are shown in Table III and Fig. 4. The key takeaways are as follows: 1) Demonstrated UWB measurements were not corrupted by proximity to propeller motors (e.g., electromagnetic interference). This gives us confidence measurements in Section V-B are comparable. 2) Showed the 2D NLOS concerns of prior work [10] are mitigated when level 3D agents are not in the same z -plane (Fig. 4(b)). 3) Confirmed the onboard downward facing LiDAR and IMU can *locally* monitor altitude/roll/pitch accurately with only small variation from the mocap ground truth. Fig. 4(a) shows data from Table III's Trial 2, demonstrating the TeraRanger Evo provides a reasonable *local* measurement of altitude. 4) Confirmed varying within our *a priori* tolerances are not a major source of APE error, supported by the small APE differences between the latter

TABLE IV
EVALUATION OF POSITIONAL ERROR IN METERS BETWEEN ALGORITHMS

Abs Position Error [m]				Experimental Scenarios (Section V-B)														
				Trial 1			Trial 2			Trial 3			Trial 4			Trial 5		
color	e _l _bias	z_fixed	Huber	Mean	Max	Std	Mean	Max	Std	Mean	Max	Std	Mean	Max	Std	Mean	Max	Std
■	✓	✓	✓	2.21	4.40	1.17	0.89	4.50	0.88	2.37	4.82	1.13	2.26	4.58	1.12	2.77	4.32	0.59
				2.55	4.95	0.64	0.89	4.46	0.97	1.32	3.94	0.95	1.33	3.97	0.97	1.55	3.71	0.90
				0.45	2.87	0.38	0.42	1.50	0.28	0.57	2.97	0.52	0.45	2.40	0.36	0.40	1.16	0.23
				0.34	2.86	0.43	0.34	1.49	0.32	0.53	2.96	0.55	0.38	2.41	0.39	0.30	1.16	0.28
■ ■ ■ ■	✓	✓	✓	1.14	3.29	1.12	2.15	3.29	1.03	1.24	3.60	1.01	2.30	3.57	1.02	1.30	3.24	1.16
				0.44	1.64	0.30	0.95	2.58	0.90	0.85	3.10	0.75	0.85	2.69	0.82	1.44	2.88	0.90
				0.29	0.91	0.14	0.27	1.02	0.14	0.32	1.03	0.20	0.28	0.78	0.14	0.29	0.71	0.13
				0.22	0.90	0.17	0.21	0.98	0.16	0.30	0.97	0.21	0.23	0.75	0.15	0.22	0.63	0.14

The final row (red) represents the proposed approach. See Section V-B. The bold entries represent the best values for a given column.

TABLE V
EVALUATION OF YAW ERROR IN DEGREES BETWEEN ALGORITHMS

Abs Heading Error [deg]				Experimental Scenarios (Section V-B)														
				Trial 1			Trial 2			Trial 3			Trial 4			Trial 5		
color	e _l _bias	z_fixed	Huber	Mean	Max	Std	Mean	Max	Std	Mean	Max	Std	Mean	Max	Std	Mean	Max	Std
■	✓	✓	✓	13.4	133.0	18.8	10.0	139.9	12.2	15.0	146.0	19.3	15.5	129.6	19.2	13.1	162.2	14.0
				13.4	134.8	18.8	10.0	139.9	12.1	15.1	146.0	19.5	15.6	132.9	19.4	13.3	162.1	14.0
				13.5	132.8	18.8	10.1	139.9	12.1	15.0	146.5	19.3	15.5	131.9	19.3	13.2	152.4	13.5
				13.8	133.0	18.7	10.4	139.0	12.2	15.0	146.2	19.3	16.8	170.1	23.1	14.0	166.7	14.4
■ ■ ■ ■	✓	✓	✓	8.0	47.5	5.5	7.7	177.7	12.8	8.5	57.5	6.1	10.1	111.4	10.7	10.4	47.5	7.6
				7.8	47.7	5.7	7.8	177.3	12.8	8.5	57.3	6.1	10.3	111.2	11.0	10.4	46.5	7.6
				7.6	42.4	5.1	7.6	172.6	12.5	8.3	58.4	6.1	10.4	122.0	11.2	10.6	48.5	7.6
				8.6	46.7	5.8	8.4	172.8	12.8	8.7	59.8	6.5	10.7	111.8	11.2	11.0	42.3	7.7

The final row (red) represents the proposed approach. See Section V-B. The bold entries represent the best values for a given column.

three rows of Table III. Specifically, this table shows the results of the proposed algorithm when z is free (z_{free}), or fixed to the commanded (z_{comm}), LiDAR measured (z_{meas}), or true (z_{true}) value. Together, takeaways (3) and (4) show that hardware realistically allows for an altitude tolerance of ± 0.10 m, meaning under realistic flight conditions, our minimal communication model introduces only a small amount of additional error (i.e., compare z_{comm} and z_{true} in Table III). Furthermore, all the takeaways together indicate our extensive UGV tests are sufficient to fairly evaluate our approach in Section V-B.

B. UGV Experiments

Hardware: Three agents, A , B and C , were used (see Fig. 3(a)). All are Turtlebot2 UGV with 6 UWB sensors evenly distributed in the xy -plane 0.32 m from the center. Agent A , B , and C 's sensors are mounted 1.75 m, 0.5 m, and 0.5 m above the ground respectively (i.e., $z_B^A = -1.25$ m).

Trials: A full 22 datasets were collected with all three agents, each with 6 antennas, moving in various ways within the mocap space. Ranging measurements are performed by each pair of antennas between agents at 25 Hz. Together the datasets total to nearly 6 hours (effectively creating over 200 hours of pairwise measurements). These datasets were used for Fig. 2(a) and (b). Tables IV and V show results for the 5 trials where all three agents continuously moved in arbitrary patterns with max positional and angular velocities of 1 m/s and 1 rad/s respectively. Additionally, two outdoor trials were conducted without mocap to verify the system generalizes to non-laboratory spaces and can operate for extended periods of time.

Parameters: 1) Eq. (9) is solved by `trust-constr` as provided by `scipy`'s `minimize` function [25].⁹ State

estimation is run at a fixed 1 Hz.¹⁰ 2) Our mean bias correction, $\bar{d}_{B_j}^{A_i}(\mathbf{T}_B^A)$, is the learned 6-degree polynomial with respect to elevation shown in Fig. 2(b). 3) The current range measurements, $\bar{d}_{B_j}^{A_i}$, and pose estimates are smoothed by a 1 s and 4 s moving average filter respectively. 4) Our loss function, ℓ , is selected to be the Huber loss $\rho_\delta(a)$ with $\delta = 0.06$. As with [9], [12], [13], Huber loss was selected due to its reduced outlier sensitivity [26].

Discussion of Results: Results are shown in Tables IV and V, Figs. 5 and 6. To demonstrate the value of individual algorithmic decisions, the tables toggle `el_bias`, `z_fixed`, `Huber`, where the final row (red) represents the proposed approach. Specifically, a check indicates:

- `el_bias`: $\bar{d}_{B_j}^{A_i}(\mathbf{T}_B^A)$ is the learned mean bias correction (see Fig. 2(b)), otherwise 0.
- `z_fixed`: z is constrained to z_B^A , otherwise free.
- `Huber`: $\ell(a) \triangleq \rho_{0.06}(a)$, otherwise squared error loss.

Our data shows the proposed approach (red) provides a 9× improvement over a direct NLLS trilateration (purple), similar to that used by [15]. The addition of the `el_bias` improves our approach's mean APE by an average of 19% (i.e., red vs orange). Overall, constraining z leads to the largest APE gains (as predicted by Table II). Unlike APE, there is not a clear best approach for AHE, but the proposed (red) is only at most 1° behind the best for any given trial. Additionally, Fig. 6 justifies our minimal communication protocol, demonstrating that even when operating with *maximal* altitude error within the experimentally demonstrated error tolerances (i.e., ± 0.10 m from Fig. 4(a)), red outperforms other formulations. Finally, the two outdoor trials, highlighted in the video, demonstrated comparable position accuracy to the laboratory trials as well as the long-term stability and reliability of the system.

⁹As in [10], initialized with previous estimate if available, otherwise $(\text{avg}(\bar{\mathbf{d}}_B^A), 0, 0)$. Here $\text{avg}(\bar{\mathbf{d}}_B^A)$ is the mean of the current measurements.

¹⁰Each robot's onboard i7-4710HQ processor can run estimation at ~ 5 Hz serially and > 25 Hz in parallel (i.e., faster than the UWB sensor rate).




Fig. 5. Error vs Time plots for Trial 2 in Tables IV and V respectively. The line color corresponds with the algorithm’s color column in the associated tables. Each colored dashed line represent that algorithm’s overall mean in that trial.




Fig. 6. Shows the corresponding absolute position error for Trial 1 when each approach (colors correspond with Table IV) is run with the incorrect altitude z -constraint $\hat{z}_B^A \triangleq -1.25 \text{ m} + \Delta z$, where -1.25 m is the true relative altitude between A and B (Fig. 3(a)) and Δz is relative altitude error (i.e., the x -axis value). The green/blue lines are level since z is unconstrained in these methods, and thus not depended on the x -axis. Although purple also has z unconstrained, without Huber loss, it is hyper-sensitive to altitude initialization (i.e., x -axis value). Comparing our proposed (red) to the approach without `el_bias` (orange), we see red outperforms orange even when operating with maximal altitude error within the reasonable error tolerances (i.e., $\pm 0.10 \text{ m}$) demonstrated in Fig. 4(a).

VI. CONCLUSION

This letter presents a UWB-ranging inter-agent relative pose estimation system that minimizes communication load. Our work outperforms the most-similar recent work [15] (APEs between 0.40 m and 1 m). We also remain competitive with other state-of-the-art approaches that have significantly higher communication costs (APE and AHE on the order of 0.10 m and 1° respectively), whose heavy reliance on *continuously* transmitted measurements has scaling issues with increased team size and/or decreased communication throughput.

By comparison, this work achieves a mean APE and AHE of 0.24 m and 9.5° respectively across our experimental trials, nearly a $9\times$ improvement over a direct NLLS trilateration approach **while requiring no continuously transmitted measurements**. Our learned measurement mean bias error correction, $\hat{d}_{B_j}^{A_i}(\mathbf{T}_B^A)$, greatly improves our APE by an average of 19% across our experimental trials. In the future, this bias correction can be further improved as a neural network, potentially dependent on more than just relative elevation – that being said, a key takeaway is that **most UWB systems would benefit from a pose-dependent UWB ranging error model**. Furthermore, future work will include additional flights on multiple larger drones in expansive outdoor environments, integration of this system into a larger resource-aware distributed SLAM pipeline, and locally synchronizing UWB hardware as in [27].

REFERENCES

- [1] A. Alarifi et al., “Ultra wideband indoor positioning technologies: Analysis and recent advances,” *Sensors*, vol. 16, no. 5, 2016, Art. no. 707.
- [2] Nooploop, “LinkTrack–Model Selection,” 2023. [Online]. Available: <https://www.nooploop.com/en/linktrack/>
- [3] N. Smaoui, O. Gnowali, and K. Kim, “Study and mitigation of platform related UWB ranging errors,” in *2020 IEEE Int. Conf. Commun. Syst. Netw.*, Bengaluru, India, Jan. 2020, pp. 346–353.
- [4] A. Ledergerber and R. D’Andrea, “Ultra-wideband range measurement model with Gaussian processes,” in *Proc. IEEE Conf. Control Technol. Appl.*, Mauna Lani Resort, HI, USA, Aug. 2017, pp. 1929–1934.Y
- [5] A. Ledergerber and R. D’Andrea, “Calibrating away inaccuracies in ultra wideband range measurements: A maximum likelihood approach,” *IEEE Access*, vol. 6, pp. 78719–78730, 2018.
- [6] M. Hamer and R. D’Andrea, “Self-calibrating ultra-wideband network supporting multi-robot localization,” *IEEE Access*, vol. 6, pp. 22292–22304, 2018.
- [7] Z. Cao et al., “Relative localization of mobile robots with multiple ultra-wideband ranging measurements,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 5857–5863.
- [8] H. Xu, L. Wang, Y. Zhang, K. Qiu, and S. Shen, “Decentralized visual-inertial-UWB fusion for relative state estimation of aerial swarm,” in *2020 IEEE Int. Conf. Robot. Automat.*, Paris, France, May 2020, pp. 8776–8782.
- [9] H. Xu et al., “Omni-swarm: A decentralized omnidirectional visual-inertial-UWB state estimation system for aerial swarm,” *IEEE Trans., Robot.*, vol. 38, no. 6, pp. 3374–3394, Dec. 2022.
- [10] A. Fishberg and J. P. How, “Multi-agent relative pose estimation with UWB and constrained communications,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Kyoto, Japan, Oct. 2022, pp. 778–785.
- [11] K. Guo, X. Li, and L. Xie, “Ultra-wideband and odometry-based cooperative relative localization with application to multi-UAV formation control,” *IEEE Trans. Cybern.*, vol. 50, no. 6, pp. 2590–2603, Jun. 2020.
- [12] S. Zheng et al., “UWB-VIO fusion for accurate and robust relative localization of round robotic teams,” *IEEE Robot. Automat. Lett.*, vol. 7, no. 4, pp. 11950–11957, Oct. 2022.
- [13] H. Zhang et al., “Range-aided drift-free cooperative localization and consistent reconstruction of multi-ground robots,” *IEEE Robot. Automat. Lett.*, vol. 8, no. 4, pp. 2094–2101, Apr. 2023.
- [14] Y. Qi, Y. Zhong, and Z. Shi, “Cooperative 3-D relative localization for UAV swarm by fusing UWB with IMU and GPS,” *J. Phys.: Conf. Ser.*, vol. 1642, no. 1, 2020, Art. no. 012028.
- [15] Y. Xianjia, L. Qingqing, J. P. Queralta, J. Heikkonen, and T. Westerlund, “Cooperative UWB-Based localization for outdoors positioning and navigation of UAVs aided by ground robots,” in *2021 IEEE Int. Conf. Auton. Syst.*, Montreal, QC, Canada, 2021, pp. 1–5.
- [16] Z. Xun et al., “CREPES: Cooperative relative pose estimation system,” in *2023 IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 5274–5281.
- [17] A. Goudar et al., “Continuous-time range-only pose estimation,” in *Proc. 20th Conf. Robots Vis.*, 2023, pp. 29–36.
- [18] A. Mohanty, A. Wu, S. Bhamidipati, and G. Gao, “Precise relative positioning via tight-coupling of GPS carrier phase and multiple UWBs,” *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 5757–5762, Apr. 2022.
- [19] W. Wen, X. Bai, G. Zhang, S. Chen, F. Yuan, and L.-T. Hsu, “Multi-agent collaborative GNSS/camera/INS integration aided by inter-ranging for vehicular navigation in urban areas,” *IEEE Access*, vol. 8, pp. 124323–124338, 2020.
- [20] Y. Xianjia, L. Qingqing, J. P. Queralta, J. Heikkonen, and T. Westerlund, “Applications of UWB networks and positioning to autonomous robots and industrial systems,” in *Proc. IEEE 10th Mediterranean Conf. Embedded Comput.*, Budva, Montenegro, Jun. 2021, pp. 1–6.
- [21] H. B. Lee, “Accuracy limitations of range-range (spherical) multilateration systems,” Lincoln Lab., Lexington, MA, USA, Tech. Rep. DOT/TSC-RA-3-8-(1), 1973.
- [22] B. K. P. Horn, “Closed-form solution of absolute orientation using unit quaternions,” *JOSA A*, vol. 4, no. 4, pp. 629–642, 1987.
- [23] R. B. Langley, “Dilution of Precision,” *GPS world*, vol. 10, no. 5, pp. 52–59, 1999.
- [24] R. H. Fabre. *TeraRanger Evo 15 m Specification Sheet*. 2021.
- [25] P. Virtanen et al., “Scipy 1.0: Fundamental algorithms for scientific computing in python,” *Nat. methods*, vol. 17, no. 3, 2020, Art. no. 261.
- [26] R. Hartley et al. *Multiple View Geometry in Computer Vision*. New York, NY, USA: Cambridge Univ. Press, 2003.
- [27] J. Sadowski and J. Stefanski, “Position and orientation estimation in radio network with groups of locally synchronized nodes,” *IEEE Access*, vol. 11, pp. 140981–140992, 2023.