Optimal Geometry for Ultra-wideband Localization using Bayesian Optimization

Wenda Zhao * **Marijan Vukosavljev** * **Angela P. Schoellig** *

* *Institute for Aerospace Studies, University of Toronto, Canada (e-mail: {wenda.zhao, mario.vukosavljev, angela.schoellig}@robotics.utias.utoronto.ca).*

Abstract: This paper introduces a novel algorithm to find a geometric configuration of ultra-wideband sources in order to provide optimal position estimation performance with Time-Difference-of-Arrival measurements. Different from existing works, we aim to achieve the best localization performance for a user-defined region of interest instead of a single target point. We employ an analysis based on the Cramer-Rao lower bound and dilution of precision to formulate an optimization problem. A Bayesian optimization-based algorithm is proposed to find an optimal geometry that achieves the smallest estimation variance upper bound while ensuring source placement constraints. The approach is validated through simulation and experimental results in 2D scenarios, showing an improvement over a naive source placement.

Keywords: Optimal source geometry, ultra-wideband, time difference of arrival, Bayesian optimization, robot localization.

1. INTRODUCTION

Accurate indoor localization remains an open research problem. Ultra-wideband (UWB) radio technology is a promising solution (Hamer and D'Andrea (2018)) which is lightweight and computationally inexpensive compared to alternative state-of-the-art approaches such as simultaneous localization and mapping. Two-way ranging (TWR) and time difference of arrival (TDoA) are the major localization approaches for UWB positioning systems and have many promising applications in autonomous robotics (Mueller et al. (2015); Hoeller et al. (2017)). Compared with TWR, TDoA does not require two-way communication between a source and a target, thus enabling localization of a large number of targets. However, the performance of TDoA-based localization relies heavily on the UWB source geometry due to its highly nonlinear measurement model (Wang et al. (2019); Kaune (2012)). In this work, we restrict our focus to TDoA-based localization, as it is better scalable to large robotic swarms.

In TDoA-based UWB localization systems, UWB modules on robots (called tags) receive radio signals from different UWB sources (also called anchors) passively, compute the difference of the arrival times (TDoA measurements) and localize themselves based on the measurements. Compared with the traditional centralized TDoA method involving a single reference anchor (Meng et al. (2016)), another common TDoA method, which we refer to as decentralized TDoA, computes the TDoA measurements between certain anchor pairs and consequently does not suffer from communication constraints and single-anchor failure (Meng et al. (2013), Ennasr et al. (2016)). Hence, we focus on decentralized TDoA localization systems.

Extensive research has been performed regarding the optimal TDoA anchor geometry for a single target point in both centralized and decentralized TDoA systems. The work in Wang et al. (2018) provided the derivation of the anchor-to-tag angle rules for the centralized TDoA approach through the Cramer-Rao lower bound (CRLB), resulting in the well-known uniform angular array (UAA) requirements. In Yang and Scheuing (2005), the optimal geometry theory was extended to 3D for Platonic solids. For the decentralized TDoA approach, the angle rules for both static and moving target points were derived in Meng et al. (2016). Most of the existing work on optimal TDoA geometry (Wang et al. (2018); Meng et al. (2016)) focuses on the performance of a single point. However, a localization system should enable robots to localize themselves

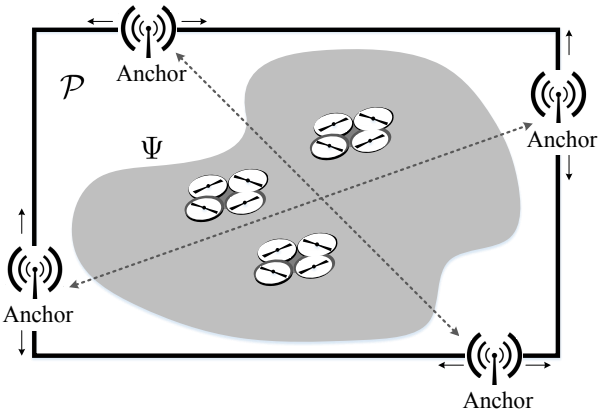

Fig. 1. The goal of this work is to find the optimal anchor positions for decentralized TDoA localization on the boundary of the rectangular space $\mathcal{P}$ to optimize the localization performance in the area Ψ. In decentralized TDoA localization systems, pairs of UWB anchors synchronize their clocks through communication (dotted arrows) and provide TDoA measurements.

inside an entire region. Therefore, we aim to find the optimal geometry for a region of interest. Extending the existing optimal geometry approaches to find the optimal geometry for an entire region is not trivial since the optimization problem becomes complicated and highly nonlinear. Instead of having strong geometric assumptions, an alternative approach for sensor placement uses a Gaussian process (GP) model from the area of machine learning to approximate the spatial phenomena of sensor performance. Then the sensor positions can be optimized by predicting the effect of placing sensors at particular locations through the learned GP model (Garnett et al. (2010)). Inspired by this approach, in this paper, we present a Bayesian optimization-based algorithm to find an optimal TDoA anchor placement that provides satisfactory localization performance for an entire region of interest. A decentralized TDoA localization system is shown in Fig. 1.

This paper provides both theoretical localization performance analysis and practical installation guidance during the localization system design. With respect to this objective, this paper has three main contributions. First, we introduce a novel optimal geometry problem that characterizes localization performance inside an entire region and formulate it as an optimization problem with soft constraints. Due to the complexity of dealing with an entire region (rather than a single point), we only consider a 2D problem as a preliminary effort. Second, we model the localization performance through a Gaussian process and achieve an optimal UWB anchor placement through Bayesian optimization. Finally, we compare the localization performance inside two pre-defined regions between baseline and optimal anchor placements in simulation and experiment. To the best of our knowledge, our method is the first to solve an optimal geometry problem with respect to a pre-defined region.

2. PROBLEM STATEMENT

The goal is to achieve an optimal anchor geometry for a decentralized TDoA localization system with respect to a pre-defined region while considering anchor placement constraints. In this section, we first define the decentralized TDoA measurement model, region of interest, and anchor placement constraints. Then we state the optimal geometry problem.

2.1 Measurement Model

TDoA-based localization systems provide a position estimate for tags based on the differences of signal arrival times. Without loss of generality, we assume that there is a single tag. For a localization system with an even number m of UWB anchors, the position of anchor i is denoted as $\boldsymbol{a}_i = [x_i, y_i]^T \in \mathbb{R}^2$, $i = 1, \cdots, m$, and the position of the tag is denoted as $\mathbf{p} = [x, y]^T \in \mathbb{R}^2$. For convenience, we define a vector $\boldsymbol{a} = \left[\boldsymbol{a}_1^T, \cdots, \boldsymbol{a}_m^T\right]^T \in \mathbb{R}^{2m}$ that contains all the anchor positions. To facilitate our analysis, we assume that the time of arrival (TOA) measurements from each anchor i to the tag have a zero mean Gaussian noise $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, where the variance σ_i^2 can typically be measured experimentally. For decentralized TDoA localization systems, the set of m anchors are divided into disjoint pairs and are denoted as $\Gamma = \{(1,2), \cdots, (m-1, m)\}$. Given an anchor placement $\boldsymbol{a}$, the TDoA measurement model for an anchor pair $(i, j) \in \Gamma$ becomes

$$\begin{aligned} d_{i,j}(\mathbf{p}, \boldsymbol{a}) &= \|\mathbf{p} - \boldsymbol{a}_i\| - \|\mathbf{p} - \boldsymbol{a}_j\| + \epsilon_i + \epsilon_j \\ &= r_{i,j}(\mathbf{p}, \boldsymbol{a}) + n_{i,j}, \end{aligned} \tag{1}$$

where $r_{i,j}(\mathbf{p}, \boldsymbol{a})$ indicates the distance difference of the tag and anchors i and j given the anchor placement $\boldsymbol{a}$, $n_{i,j} \sim \mathcal{N}(0, \sigma_i^2 + \sigma_j^2)$ is the measurement noise, and $\|\cdot\|$ is the Euclidean norm.

For a decentralized TDoA localization system with m UWB anchors, stacking the TDoA measurements in (1) into a column vector yields the measurement model

$$\mathbf{d}(\mathbf{p}, \boldsymbol{a}) = \mathbf{r}(\mathbf{p}, \boldsymbol{a}) + \mathbf{n}, \ \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\mathbf{d}), \tag{2}$$

where $\boldsymbol{\Sigma}_\mathbf{d}$ is the covariance matrix of the TDoA measurement errors. Assuming that the decentralized TDoA measurements are independent and uncorrelated, $\boldsymbol{\Sigma}_\mathbf{d}$ is of the form

$$\boldsymbol{\Sigma}_\mathbf{d} = \operatorname{diag}\left(\sigma_1^2 + \sigma_2^2, \cdots, \sigma_{m-1}^2 + \sigma_m^2\right). \tag{3}$$

For convenience, we drop the functional dependence and indicate the TDoA measurements and distance differences as $\mathbf{d}$ and $\mathbf{r}$, respectively.

2.2 Region of Interest

The region of interest (ROI) $\Psi \subset \mathbb{R}^2$, indicated as the gray region in Fig. 1, is expected to have good estimation performance through the placement of anchors. Considering obstacles inside the space, Ψ could be irregularly shaped, for example, being non-convex or not connected. However, we assume Ψ is compact and that the tag is inside the ROI, $\mathbf{p} \in \Psi$. As a preliminary effort, our measurement model (2) also assumes line-of-sight propagation in the region of interest and does not consider biases of UWB measurements.

2.3 Anchor Placement Constraints

In most practical settings, the positions at which the anchors can be installed can be restricted due to factors such as accessibility and the physical geometry of the space. In this paper, we assume the space is rectangular, contains the ROI, and that the anchors must be placed on the boundary of the space. More formally, the space has the form $\mathcal{P} = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}] \subset \mathbb{R}^2$, where $x_{\min}, x_{\max}, y_{\min}, y_{\max} \in \mathbb{R}$ denote the space limits in the plane. We denote the boundary of $\mathcal{P}$ as $\partial\mathcal{P}$. Thus we require that $\Psi \subset \mathcal{P}$ and $\boldsymbol{a}_i = [x_i, y_i]^T \in \partial\mathcal{P}$, $i = 1, \ldots, m$. For the placement of m anchors, we define a set $\mathcal{B} \subset \mathbb{R}^{2m}$ containing all possible geometric configurations of m anchors such that $\boldsymbol{a} \in \mathcal{B}$ if and only if $\boldsymbol{a}_i \in \partial\mathcal{P}$, $i = 1, \ldots, m$.

2.4 Optimal Geometry Problem for a Region of Interest

In this section, we provide a metric to quantify the performance of the tag's position estimate using the TDoA localization system. We use this metric to find an optimal anchor geometry for a pre-defined ROI given placement constraints. We propose that the optimal geometry of UWB anchors corresponds to the smallest upper bound of the estimation variance anywhere within Ψ. Given an

anchor placement $\boldsymbol{a}$ and tag position $\mathbf{p} \in \Psi$, an estimator provides the position estimate $\hat{\mathbf{p}} = [\hat{x}, \hat{y}]^T \in \mathbb{R}^2$, which possesses a corresponding covariance matrix $\boldsymbol{\Sigma}(\hat{\mathbf{p}}|\mathbf{p}, \boldsymbol{a})$. The estimation is assumed to be unbiased. The specific estimator we used in our experiments is introduced in Sec. 5.2 (see (28)), although others may be used. In order to evaluate the localization performance, we define an A-optimal (Fedorov (2013)) type performance metric to denote the localization variance at the position $\mathbf{p}$, which converts the covariance matrix $\boldsymbol{\Sigma}(\hat{\mathbf{p}}|\mathbf{p}, \boldsymbol{a})$ into a scalar using the diagonal entries only

$$V(\hat{\mathbf{p}}|\mathbf{p}, \boldsymbol{a}) = \sqrt{\mathrm{Var}(\hat{x}|\mathbf{p}, \boldsymbol{a}) + \mathrm{Var}(\hat{y}|\mathbf{p}, \boldsymbol{a})}, \tag{4}$$

where $\mathrm{Var}(\hat{x}|\mathbf{p}, \boldsymbol{a})$ and $\mathrm{Var}(\hat{y}|\mathbf{p}, \boldsymbol{a})$ are the estimation variances of $\hat{x}$ and $\hat{y}$. Hence, the upper bound of localization variance in Ψ can be expressed as

$$M(\boldsymbol{a}) = \max_{\mathbf{p} \in \Psi} V(\hat{\mathbf{p}}|\mathbf{p}, \boldsymbol{a}), \tag{5}$$

given an anchor placement $\boldsymbol{a}$. We aim to find the optimal placement of UWB anchors $\boldsymbol{a}^\star$ that reaches the smallest $M(\boldsymbol{a})$ given the anchor placement constraint:

$$\boldsymbol{a}^\star = \underset{\boldsymbol{a} \in \mathcal{B}}{\operatorname{argmin}}\, M(\boldsymbol{a}). \tag{6}$$

3. OPTIMAL ANCHOR PLACEMENT

In this section, we propose a Bayesian optimization-based approach (Fig. 2) to solve the optimization problem formulated in Sec. 2. Through the analysis of the Cramer-Rao lower bound (CRLB) and dilution of precision (DOP), we introduce a set of angle constraints for each anchor pair and compute the localization performance metric for an anchor placement $\boldsymbol{a}$. In order to analyze the highly complex objective function, we model the relationship between an anchor placement and its estimation performance with a Gaussian process and achieve an optimal solution through Bayesian optimization.

3.1 Angle Constraints of Anchor Placement

Prior to solving the optimization problem in (6), we introduce a set of angle constraints in this subsection to reduce the solution space and thereby allow us to more efficiently optimize the anchor placement. We assume that the UWB localization system provides an unbiased estimate $\hat{\mathbf{p}}$ of the tag position $\mathbf{p}$ through multilateration using a TDoA measurement $\mathbf{d}$. The CRLB is the most commonly used benchmark for unbiased estimators (Wang et al. (2018)). The variance of an unbiased estimate $\hat{\mathbf{p}}$ is lower bounded by the CRLB

$$\mathbf{I}^{-1}(\mathbf{d}|\mathbf{p}) \leq \mathbb{E}\left[(\mathbf{p} - \hat{\mathbf{p}})(\mathbf{p} - \hat{\mathbf{p}})^T\right], \tag{7}$$

where $\mathbb{E}(\cdot)$ indicates the expectation and $\mathbf{I}(\mathbf{d}|\mathbf{p})$ is the Fisher information matrix (FIM). With the definition in (2), we denote the probability density function of $\mathbf{d}$ by $f(\mathbf{d}|\mathbf{p})$. The FIM was originally derived in Chan and Ho (1994) and can be expressed as

$$\begin{aligned}\mathbf{I}(\mathbf{d}|\mathbf{p}) =& \mathbb{E}\left[\frac{\partial}{\partial \mathbf{p}} \ln(f(\mathbf{d}|\mathbf{p})) \frac{\partial}{\partial \mathbf{p}} \ln(f(\mathbf{d}|\mathbf{p}))^T\right] \\ =& \frac{\partial \mathbf{r}^T}{\partial \mathbf{p}} \boldsymbol{\Sigma}_{\mathbf{d}}^{-1} \frac{\partial \mathbf{r}}{\partial \mathbf{p}}.\end{aligned} \tag{8}$$

with $\boldsymbol{\Sigma}_{\mathbf{d}}$ defined in (3). For simplification, we define

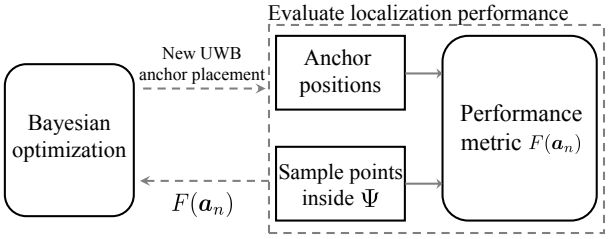

Fig. 2. An optimal anchor placement is achieved through a Bayesian optimization-based approach. In the nth iteration, the localization performance $F(\boldsymbol{a}_n)$ of a selected anchor placement $\boldsymbol{a}_n$ is evaluated at the sampled points inside the region of interest Ψ. A Bayesian optimization algorithm optimizes anchor placements iteratively based on the performance at selected anchor positions.

$$\mathbf{Q} = \frac{\partial \mathbf{r}^T}{\partial \mathbf{p}} = [q_{1,2}, \cdots, q_{m-1,m}],\ q_i = \begin{bmatrix}\cos(\theta_i(\mathbf{p}, \boldsymbol{a}))\\ \sin(\theta_i(\mathbf{p}, \boldsymbol{a}))\end{bmatrix}, \tag{9}$$

where $q_{i,j} = q_i - q_j$, $\cos(\theta_i(\mathbf{p}, \boldsymbol{a})) = (x_i - x)/(\|\mathbf{p} - \boldsymbol{a}_i\|)$, $\sin(\theta_i(\mathbf{p}, \boldsymbol{a})) = (y_i - y)/(\|\mathbf{p} - \boldsymbol{a}_i\|)$. The angle of arrival of the measurement from anchor i to tag $\mathbf{p}$ is $\theta_i(\mathbf{p}, \boldsymbol{a}) \in [0, 2\pi]$ and later we indicate it as $\theta_{i\mathbf{p}}(\boldsymbol{a})$ for brevity. Then (8) can be expressed as

$$\mathbf{I}(\mathbf{d}|\mathbf{p}) = \mathbf{Q}\boldsymbol{\Sigma}_{\mathbf{d}}^{-1}\mathbf{Q}^T. \tag{10}$$

According to the A-optimal type performance metric in (4), we seek to find the anchor geometry that minimizes the trace of the inverse of the FIM. By direct calculation, the trace of the FIM can be expressed as

$$\begin{aligned}\mathrm{Tr}(\mathbf{I}(\mathbf{d}|\mathbf{p})) =& \sum_{(i,j)\in\Gamma} \frac{2}{\sigma_i^2 + \sigma_j^2}\left[1 - \cos\left(\theta_{i\mathbf{p}}(\boldsymbol{a}) - \theta_{j\mathbf{p}}(\boldsymbol{a})\right)\right] \\ \leqslant& \sum_{(i,j)\in\Gamma} \frac{4}{\sigma_i^2 + \sigma_j^2},\end{aligned} \tag{11}$$

where $\mathrm{Tr}(\cdot)$ represents the trace of a matrix. The equality holds if and only if $\theta_{i\mathbf{p}}(\boldsymbol{a}) - \theta_{j\mathbf{p}}(\boldsymbol{a}) = \pm\pi, \forall(i,j) \in \Gamma$. One can easily show that there exists the relationship

$$\mathrm{Tr}(\mathbf{I}^{-1}(\mathbf{d}|\mathbf{p})) \geqslant \frac{4}{\mathrm{Tr}(\mathbf{I}(\mathbf{d}|\mathbf{p}))}. \tag{12}$$

Therefore, the optimal anchor placement $\boldsymbol{a}^\star$ should maximize the trace of the FIM to reach the CRLB and the angle measurements from each anchor pair should satisfy the angle rule $\theta_{i\mathbf{p}}(\boldsymbol{a}) - \theta_{j\mathbf{p}}(\boldsymbol{a}) = \pm\pi, \forall(i,j) \in \Gamma$.

Considering an entire region, it is not feasible for all points inside the region to satisfy the angle rule. Therefore, the estimation inside the region can usually not reach the theoretical lower bound in (7). However, we can use relaxed angle constraints as prior knowledge to reduce the solution space of the optimization problem in (6). In Fig. 3, a simple anchor pair geometric configuration is demonstrated. We represent the tag position as the red star and the anchor positions as solid circles with $\theta^* = \theta_{2\mathbf{p}} - \theta_{1\mathbf{p}}$ and $\tilde{\theta} = \theta'_{2\mathbf{p}} - \theta_{1\mathbf{p}}$. Given the rectangular boundary anchor placement constraint, for an anchor pair placement $\boldsymbol{a}_1$ and $\boldsymbol{a}'_2$ with $\cos(\tilde{\theta}) > 0$, there always exists a better geometric configuration of $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$ with $\cos(\theta^*) < 0$ that provides a lower estimation variance (according to (11)). Hence, we aim to guarantee $\cos\left(\theta_{i\mathbf{p}}(\boldsymbol{a}) - \theta_{j\mathbf{p}}(\boldsymbol{a})\right) \leqslant 0$ for the points $\mathbf{p} \in \Psi$, which reduces the allowable set of anchor positions. We define a restricted set of possible anchor placements as

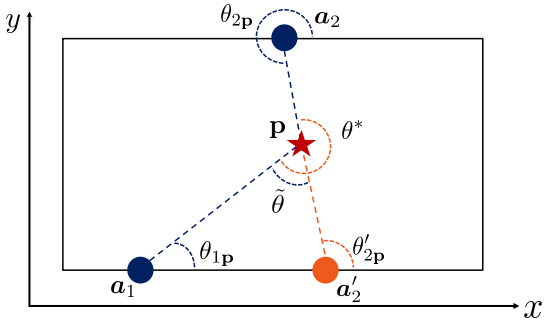

Fig. 3. An example of an anchor configuration with tag $\mathbf{p}$ and anchors $\boldsymbol{a}_1$, $\boldsymbol{a}_2$ and $\boldsymbol{a}'_2$. The angle of arrival of the measurements are indicated as $\theta_{1\mathbf{p}}, \theta_{2\mathbf{p}}$ and $\theta'_{2\mathbf{p}}$, respectively.

$$\mathcal{A} = \left\{ \boldsymbol{a} \in \mathcal{B} \mid \frac{\pi}{2} \leqslant |\theta_{i\mathbf{p}}(\boldsymbol{a}) - \theta_{j\mathbf{p}}(\boldsymbol{a})| \leqslant \frac{3\pi}{2}, \forall (i,j) \in \Gamma \right\}. \tag{13}$$

This set of angle constraints reduces the solution space and allows us to more efficiently optimize the anchor geometry problem.

3.2 Performance Metric of TDoA Anchor Geometries

Given the restricted set of possible anchor configurations, we now derive the A-optimal type performance metric in (4) through dilution of precision (DOP). DOP is a localization performance metric commonly used in satellite navigation (Santerre et al. (2017)). Intuitively, DOP establishes a connection between localization accuracy and measurement accuracy. Based on the concept of DOP, we derive an approximation for the localization performance metric $M(\boldsymbol{a})$ in this subsection for an arbitrary unbiased estimator $\hat{\mathbf{p}}$. This approach follows Ledergerber et al. (2015), but we derive the covariance matrix of the position estimate through the measurement noise.

Suppose that the anchor placement $\boldsymbol{a}$ is fixed. With an unbiased estimate $\hat{\mathbf{p}} = [\hat{x}, \hat{y}]^T$, we linearize (2) around the state $\mathbf{p} = [x, y]^T$ to obtain a first order approximation between the deviation error $\Delta\mathbf{p} = \hat{\mathbf{p}} - \mathbf{p}$ and the measurement noise $\mathbf{n}$:

$$\mathbf{n} = \mathbf{H}\Delta\mathbf{p}. \tag{14}$$

The matrix $\mathbf{H}$ is straightforward to derive and has the form

$$\mathbf{H} = \begin{bmatrix} \frac{\partial d_{1,2}}{\partial x}(\mathbf{p}) & \frac{\partial d_{1,2}}{\partial y}(\mathbf{p}) \\ \vdots & \vdots \\ \frac{\partial d_{m-1,m}}{\partial x}(\mathbf{p}) & \frac{\partial d_{m-1,m}}{\partial y}(\mathbf{p}) \end{bmatrix}. \tag{15}$$

Since we are interested in the variance of the estimated position, we invert (14) and solve for $\Delta\mathbf{p}$:

$$\Delta\mathbf{p} = \mathbf{H}^+\mathbf{n}, \tag{16}$$

where $\mathbf{H}^+ = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T$ is the pseudo inverse of $\mathbf{H}$. With the expression of (16), the covariance matrix of estimated position can be expressed as

$$\boldsymbol{\Sigma}(\hat{\mathbf{p}}|\mathbf{p}, \boldsymbol{a}) = \mathbb{E}\left[\Delta\mathbf{p}\Delta\mathbf{p}^T\right] = \mathbf{H}^+\boldsymbol{\Sigma}_{\mathbf{d}}(\mathbf{H}^+)^T, \tag{17}$$

where $\boldsymbol{\Sigma}_{\mathbf{d}}$ is the covariance matrix of measurement noise defined in (3). The performance metric $M(\boldsymbol{a})$ in (5) for the anchor placement $\boldsymbol{a}$ becomes

$$M(\boldsymbol{a}) = \max_{\mathbf{p}\in\Psi} \sqrt{\mathrm{Tr}\left(\boldsymbol{\Sigma}(\hat{\mathbf{p}}|\mathbf{p}, \boldsymbol{a})\right)}. \tag{18}$$

To further approximate the computation in (18), we select sample points inside Ψ and evaluate the estimation variance of each sample point $\mathbf{p}_s = [x_s, y_s]^T \in \Psi$ to represent the localization performance within Ψ.

In order to incorporate the angle requirements derived in Sec. 3.1, we introduce a penalty term into the performance metric to penalize the anchor configuration $\boldsymbol{a} \notin \mathcal{A}$. Then the performance metric becomes

$$F(\boldsymbol{a}) = M(\boldsymbol{a}) + P(\boldsymbol{a}), \tag{19}$$

where

$$P(\boldsymbol{a}) = \begin{cases} 0, & \boldsymbol{a} \in \mathcal{A} \\ c, & \boldsymbol{a} \notin \mathcal{A} \end{cases} \tag{20}$$

is the penalty term and $c > 0$ is an adjustable constant parameter. In summary, the original optimal geometry problem in (6) has been transformed into

$$\boldsymbol{a}^\star = \underset{\boldsymbol{a}\in\mathcal{B}}{\operatorname{argmin}}\, F(\boldsymbol{a}). \tag{21}$$

The formulation in (21) results in an optimization problem with $2m$ decision variables.

3.3 Bayesian Optimization for Optimal Anchor Geometry

For a given anchor placement $\boldsymbol{a}$, the performance metric $F(\boldsymbol{a})$ is obtained through computing the highly nonlinear function (19) for each sample point. Hence, it is very computationally expensive to solve the optimization problem (21) through brute-force search. Therefore, we propose a Bayesian optimization-based approach (see Fig. 2) to find an optimal anchor placement $\boldsymbol{a}^\star$ that minimizes $F(\boldsymbol{a})$.

First, we model the relationship between the anchor placement $\boldsymbol{a}$ and the performance metric $F(\boldsymbol{a})$ as a GP,

$$g(\boldsymbol{a}) \sim \mathcal{GP}(\mu(\boldsymbol{a}), \kappa(\boldsymbol{a}, \boldsymbol{a}')). \tag{22}$$

where $\mu(\cdot)$ and $\kappa(\cdot,\cdot)$ are the prior mean and the prior covariance function, also known as kernel, of the GP. Using the GP framework, the function value $g(\boldsymbol{a})$ can be predicted at an arbitrary input $\tilde{\boldsymbol{a}}$ with a mean $\mu(\tilde{\boldsymbol{a}})$ and a covariance $\sigma^2(\tilde{\boldsymbol{a}})$ based on a set of n past observations (the training set) $\mathcal{D}_n = \{\boldsymbol{a}_i, F(\boldsymbol{a}_i)\}_{i=1}^n$ (Rasmussen and Williams (2005)). The expressions of the posterior mean and variance at $\tilde{\boldsymbol{a}}$ are

$$\begin{aligned} \mu(\tilde{\boldsymbol{a}}) &= \mathbf{k}^T(\tilde{\boldsymbol{a}})\mathbf{K}^{-1}\hat{\mathbf{g}} \\ \sigma^2(\tilde{\boldsymbol{a}}) &= \kappa(\tilde{\boldsymbol{a}}, \tilde{\boldsymbol{a}}) - \mathbf{k}^T(\tilde{\boldsymbol{a}})\mathbf{K}^{-1}\mathbf{k}(\tilde{\boldsymbol{a}}), \end{aligned} \tag{23}$$

where $\mathbf{K}$ is the covariance matrix with the entries $\mathbf{K}_{i,j} = \kappa(\boldsymbol{a}_i, \boldsymbol{a}_j)$, $\hat{\mathbf{g}}$ is a column vector of the observed function values $F(\boldsymbol{a}_i)$, and the vector $\mathbf{k}(\tilde{\boldsymbol{a}}) = [\kappa(\tilde{\boldsymbol{a}}, \boldsymbol{a}_1), \ldots, \kappa(\tilde{\boldsymbol{a}}, \boldsymbol{a}_n)]^T$ contains the covariance between the new test point $\tilde{\boldsymbol{a}}$ and the observed data in $\mathcal{D}_n$. Here, we select a common squared exponential kernel for the GP

$$\kappa(\boldsymbol{a}, \boldsymbol{a}') = \exp\left(-\frac{\|\boldsymbol{a} - \boldsymbol{a}'\|^2}{2l^2}\right) \tag{24}$$

with the length scale $l > 0$.

We apply Bayesian optimization to find the minimum of $F(\boldsymbol{a})$ by iteratively sampling anchor placements inside the set $\mathcal{B}$. As a sample-efficient global optimization algorithm, Bayesian optimization proposes sample points through an acquisition function. Expected Improvement (EI) is a common acquisition function to balance the exploration and

exploitation of the unknown function. In the nth iteration, given past observation data $\mathcal{D}_n = \{\boldsymbol{a}_i, F(\boldsymbol{a}_i)\}_{i=1}^n$, we fit a GP to approximate (19) and denote the model as $g_n(\boldsymbol{a})$ with the mean $\mu_n(\boldsymbol{a})$ and the variance $\sigma_n^2(\boldsymbol{a})$. Defining $g_n^\star := \min_i F(\boldsymbol{a}_i), i = 1, \ldots, n$, the EI acquisition function can be expressed as

$$\mathrm{EI}_n(\boldsymbol{a}) = (g_n^\star - \mu_n(\boldsymbol{a}))\Phi(Z) + \sigma_n(\boldsymbol{a})\phi(Z), \tag{25}$$

where $Z := \frac{g_n^\star - \mu_n(\boldsymbol{a})}{\sigma_n(\boldsymbol{a})}$, and Φ and ϕ are the cumulative distribution and probability density functions of the standard normal distribution (Shahriari et al. (2015)). Then the next sample anchor placement is

$$\boldsymbol{a}_{n+1} = \underset{\boldsymbol{a}\in\mathcal{B}}{\operatorname{argmax}}\ \mathrm{EI}_n(\boldsymbol{a}). \tag{26}$$

Comparing to the original objective function $F(\boldsymbol{a})$ in (21), $\mathrm{EI}_n(\boldsymbol{a})$ is smooth and allows easy evaluation of first and second derivatives. Hence, the optimization problem in (26) can be solved through a continuous first- or second-order optimization method (Frazier (2018)).

4. OPTIMAL GEOMETRY ALGORITHM

The optimal anchor positions are found through an iterative process: we select the sample anchor placement $\boldsymbol{a}$ based on the acquisition scheme in (26), compute the performance metric $F(\boldsymbol{a})$ in (19) using sample points within Ψ, refit the GP with the new observation, and repeat the process. After a finite number of iterations (Srinivas et al. (2009)), we can find an optimal anchor placement $\boldsymbol{a}^\star$ which is the global minimum of the GP approximated performance metric:

$$\boldsymbol{a}^\star = \underset{\boldsymbol{a}\in\mathcal{B}}{\operatorname{argmin}}\, \mu(\boldsymbol{a}), \tag{27}$$

where $\mu(\boldsymbol{a})$ indicates the posterior mean of the fitted GP function. In order to reduce the dimension of the GP, we used a mapping function to describe an anchor position on the boundary of the space $\mathcal{P}$ during the implementation. Hence, the GP is of m dimensions in the proposed algorithm. The optimal geometry algorithm for decentralized TDoA-based UWB localization systems is outlined in Algorithm 1.

Algorithm 1: Optimal anchor geometry for TDoA localization system.

Input : Number of anchors m, anchor placement constraints $\mathcal{B}$, and ROI Ψ
Hyperparameters of the GP
Output: Optimal placements of anchors $\boldsymbol{a} = [\boldsymbol{a}_1, \cdots, \boldsymbol{a}_m]^T$

1 $\{\mathbf{p}_s\} \longleftarrow$ sample points inside Ψ
2 $\mathcal{A} \longleftarrow$ the soft angle constraint (13)
3 $\boldsymbol{a}_0 \longleftarrow$ initial guess of anchor placement
4 $F_0 \longleftarrow \max_{\mathbf{p}_s} \sqrt{\mathrm{Tr}\left(\boldsymbol{\Sigma}\left(\hat{\mathbf{p}}_s|\mathbf{p}_s, \boldsymbol{a}_0\right)\right)} + P(\boldsymbol{a}_0)$, $n \leftarrow 1$
5 Initialize GP with $(\boldsymbol{a}_0, F_0)$
6 **do**
7 $\quad \boldsymbol{a}_n \longleftarrow \operatorname{argmax}_{\boldsymbol{a}\in\mathcal{B}} \mathrm{EI}_{n-1}(\boldsymbol{a})$
8 $\quad$ Calculate observation $F(\boldsymbol{a}_n)$
9 $\quad$ Update GP with $(\boldsymbol{a}_n, F(\boldsymbol{a}_n))$
10 $\quad$ n $\longleftarrow$ n+1
11 **while** $\|\boldsymbol{a}_n - \boldsymbol{a}_{n-1}\| > \varepsilon$ ***and*** $n < \max_{iter}$
12 **return** $\boldsymbol{a}^\star$ corresponding to $\min_{\boldsymbol{a}} F(\boldsymbol{a})$

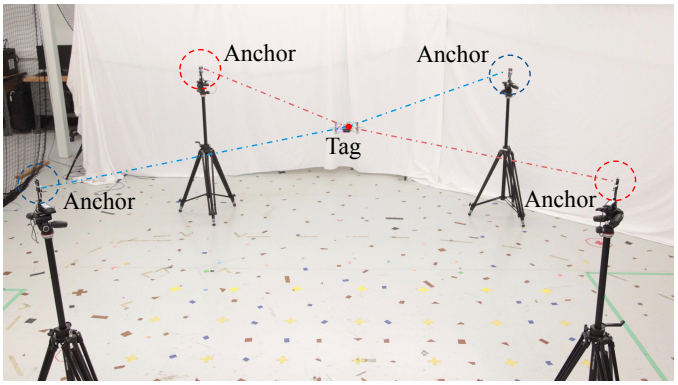

Fig. 4. Ultra-wideband indoor localization system using decentralized TDoA. Ranging measurements of the same anchor pair are indicated with the same color (red or blue).

The main loop (lines 6 – 11) repeatedly samples the anchor placements through the acquisition function (25) and updates the objective function based on the GP framework. The algorithm is terminated if either (a) the distance between two consecutive sample points is smaller than the set parameter ε or (b) the iteration n is larger than $\max_{\text{iter}}$.

5. SIMULATION AND EXPERIMENTAL RESULTS

This section contains the simulation and experimental results of the proposed optimal geometry algorithm. The algorithm is implemented in Python using the GPy framework (GPy (since 2012)). In what follows, we consider two pre-defined ROIs, denoted as Ψ_1 and Ψ_2 (see Fig. 5). We demonstrate that the simulation predictions agree with the experimental results. Then we compare the improvement in localization performance using the optimal placement over a baseline placement for each ROI.

5.1 Experimental Setup

For 2D localization, we used 2 pairs of TDoA anchors and placed the UWB anchors on tripods at a height of 1.55 m during the experiments. The anchor positions were calibrated with a Leica Total Station and the installation errors were within 1 centimeter. The experimental setup is shown in Fig. 4. We used the Crazyflie 2.0 quadrotor, UWB tags and anchors from Bitcraze's Loco Positioning System. A motion capture system is used as the ground truth. The variances of UWB measurements, $\sigma_i^2 = 0.03\ \mathrm{m}^2$ for $i = 1, \ldots, m$, were obtained from experimental data using this system.

Without any further analysis, the naive anchor placement approach is to place anchors in the corners of the room. Therefore, we selected this configuration as the baseline placement for both ROIs and compared the localization performance with the optimal geometry.

5.2 Localization Performance Evaluation

The localization performance of the baseline and optimal anchor placements are evaluated in both simulation and experiment. In simulation, we uniformly sampled 2400

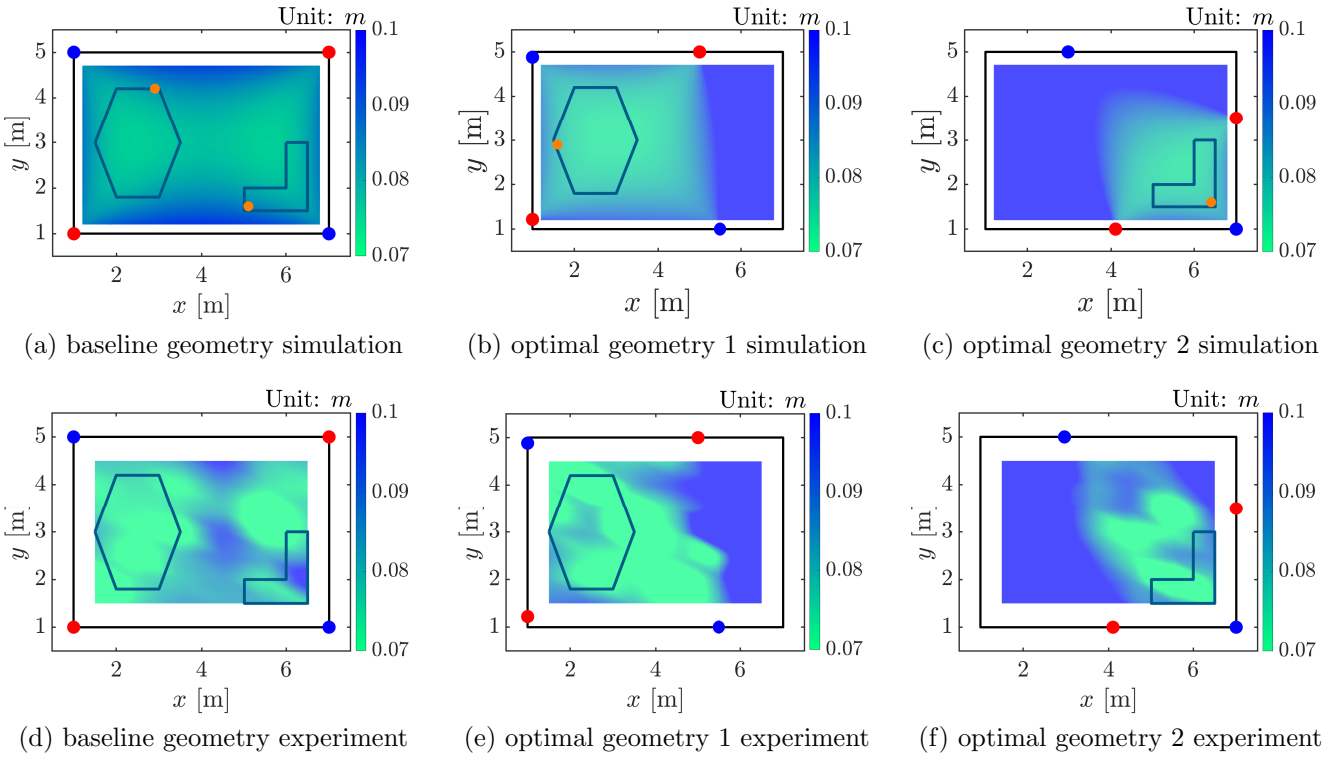

(a) baseline geometry simulation (b) optimal geometry 1 simulation (c) optimal geometry 2 simulation

(d) baseline geometry experiment (e) optimal geometry 1 experiment (f) optimal geometry 2 experiment

Fig. 5. Localization performance comparisons between simulation (a,b,c) and experiment (d,e,f). The shapes inside the 4m×6m space indicate the region of interest Ψ_1 (left) and Ψ_2 (right). Different anchor pairs are indicated with different colors (red or blue). The orange point represents the sample point with the largest estimation variance during simulation. The color plots show $V(\hat{\mathbf{p}}|\mathbf{p}, \boldsymbol{a})$ for baseline and optimal geometries. With the optimal anchor placements, the localization variances are smaller (indicated as lighter color) inside both the regions of interest.

points inside the 4m × 6m space $\mathcal{P}$ and computed the estimation variance (17) at each sample point $\mathbf{p}_s = [x_s, y_s]^T \in \mathcal{P}$ for both the baseline and optimal anchor geometries for Ψ_1 and Ψ_2. During the experiments, we uniformly selected 42 sample points with the same height as the anchors inside the space. For testing the localization performance of UWB measurements, we fixed the Crazyflie with the UWB tag at each sample point $\mathbf{p}_s$ and collected the TDoA measurements and ground truth data at the baseline and optimal anchor placements for each ROI.

For the 2D localization problem, we used the multilateration method as an unbiased offline estimator (Betke and Gurvits (1997)). Given an anchor placement $\boldsymbol{a}$ corresponding to Ψ_j and the TDoA measurement $\mathbf{d}$, the estimated position $\hat{\mathbf{p}}_s$ of sample point $\mathbf{p}_s \in \Psi_j$ can be computed as

$$\hat{\mathbf{p}}_s = \underset{\bar{\mathbf{p}}_s \in \mathbb{R}^2}{\operatorname{argmin}} \sum_{(i,j)\in\Gamma} \left[\|\bar{\mathbf{p}}_s - \boldsymbol{a}_i\| - \|\bar{\mathbf{p}}_s - \boldsymbol{a}_j\| - d_{i,j}(\mathbf{p}_s, \boldsymbol{a}) \right]^2, \tag{28}$$

by using Newton's method for optimization. Since (28) is a convex problem, an unique minimum exists.

To obtain an empirical estimate of the distribution of $\hat{\mathbf{p}}_s$, we collected N independent TDoA measurements $\left\{\mathbf{d}^{(1)}, \cdots, \mathbf{d}^{(N)}\right\}$ and solved for the corresponding position estimate $\left\{\hat{\mathbf{p}}_s^{(1)}, \cdots, \hat{\mathbf{p}}_s^{(N)}\right\}$ based on (28), where the superscripts denote the measurement indexes. Before sending the raw measurements into the estimator, we performed outlier rejection and lowpass filtering to remove high-frequency noise. The empirical value for (4) is obtained

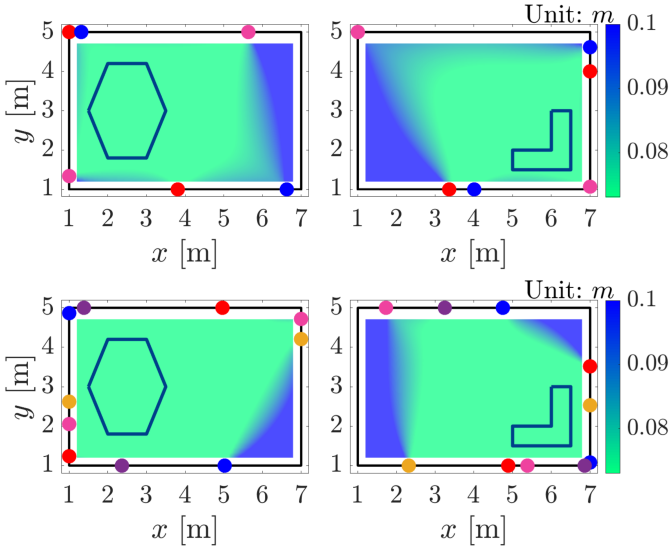

Fig. 6. The localization performance of optimal geometries with 6 anchors (in the first row) and 10 anchors (in the second row) in simulation. The anchor pairs are indicated with the same color and the range of the color bars are the same as Fig. 5.

by

$$V_{\text{emp}}(\hat{\mathbf{p}}_s) = \sqrt{\text{Var}_{\text{emp}}(\hat{x}_s) + \text{Var}_{\text{emp}}(\hat{y}_s)}, \tag{29}$$

with $\text{Var}_{\text{emp}}(\hat{x})$ being the empirical variance of the error values $(\hat{x}_s^{(1)} - x_s, ..., \hat{x}_s^{(N)} - x_s)$ and analogously for y. The

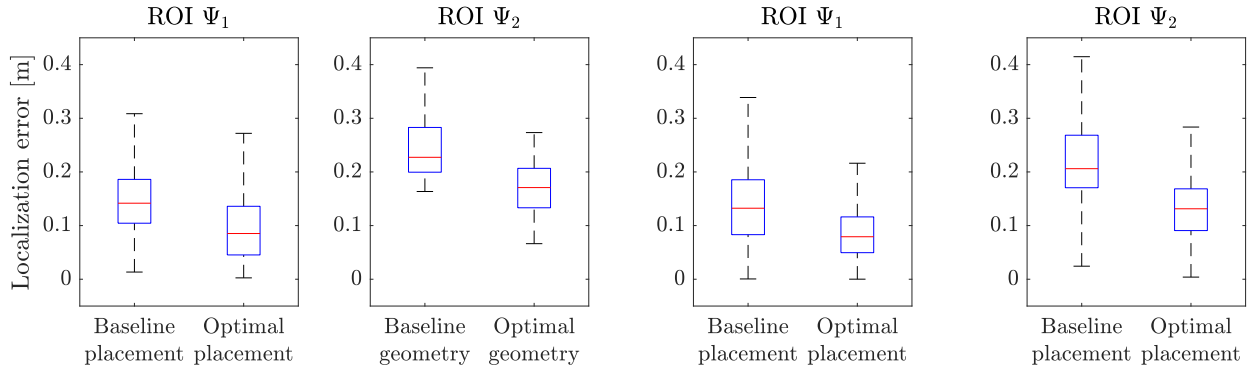

(a) the worst estimation performance (b) the average estimation performance

Fig. 7. A summary of the localization error with baseline and optimal UWB anchor placements. (a) Comparison of the worst localization performance for Ψ_1 and Ψ_2. (b) Comparison of the average localization error over 20 sample points. With optimal anchor placements, both the largest and average estimation variances in Ψ_1 and Ψ_2 are smaller as compared to the baseline anchor placement.

Table 1. Localization Performance for the Baseline Geometry and Optimal Geometries in Simulation.

	Baseline 4 anchors	Optimal 4 anchors	Optimal 6 anchors	Optimal 10 anchors
F $[m]$ for Ψ_1	0.082	0.080	0.070	0.053
F $[m]$ for Ψ_2	0.085	0.079	0.073	0.055

(x_s, y_s) is obtained from the ground truth. The comparison of the simulation and experimental results with 4 anchors are demonstrated in Fig. 5. It can be observed that the simulation and experimental results show similar trends. Hence, it is reasonable to use the proposed algorithm to find the optimal anchor placement in simulation. With the baseline placement, the localization performance inside a pre-defined region is not optimized. With the optimal anchor geometry, however, the worst-case estimation variance inside the pre-defined region is minimized. Consequently, the region is guaranteed to achieve the best estimation performance given the anchor installation constraint.

We also show how the results vary with the number of anchors in simulation. Fig. 6 shows the localization performance for the optimal geometries with 6 and 10 anchors. Table 1 summarizes these simulation results together with the baseline geometry with 4 anchors shown in Fig. 5 (a). We drop the function dependence and indicate the performance metric as F for convenience. Compared to the optimal geometries in Fig. 5 (b) and (c), the localization system provides better performance as more anchors are available, but the optimization problem requires more iterations to converge as one would expect. We summarize the localization performance reductions (% Red.) of optimal geometries with 6 and 10 anchors as compared to the optimal geometry with only 4 anchors in Table 2.

Table 2. Summary of Localization Performance Reductions of Optimal Geometries in Simulation.

	Optimal 6 anchors, Ψ_1	Optimal 10 anchors, Ψ_1	Optimal 6 anchors, Ψ_2	Optimal 10 anchors, Ψ_2
% Red.	12.5%	33.7%	7.6%	30.4%

Incorporating the soft angle constraint in (19)-(21) accelerates the optimization process. Without the soft angle constraint, the Bayesian optimization algorithm requires more iterations to converge to an optimal geometry.

5.3 Performance Evaluation inside the ROI

Finally, to evaluate the actual localization performance inside the ROI, we selected 20 points $\mathbf{p}_s$ inside each region Ψ_j and collected UWB measurements and the ground truth position data from the motion capture system.

We compute the RMS localization error at each sample point and for each of the N measurements

$$\mathbf{e}(\mathbf{p}_s) = \left[e_s^{(1)}, \cdots, e_s^{(N)}\right]^T, \tag{30}$$

where $e_s^{(i)} = \sqrt{(\hat{x}_s^{(i)} - x_s)^2 + (\hat{y}_s^{(i)} - y_s)^2}$, and compare the localization performance between baseline and optimal geometries. Using the vector entries, we compute the empirical mean and variance for each sample point. We indicate the largest variance over the selected sample points as the worst localization performance and the average variance of all sample points as the average localization performance. In Fig. 7, we show the worst and average localization performance through box plots where the center red lines indicate the median and the lower and upper boundary are at the 25 and 75 percentiles of the errors. We note that, in practice, the TDoA measurements may have a non-zero bias (as seen in Fig. 7); in future work, we aim to explore learning-based approaches to compensate for the non-zero mean bias in TDoA measurements (Ledergerber and D'Andrea (2017)).

We present the performance metric $F(\boldsymbol{a})$, and the largest and average localization variances of the RMS error for each anchor placement in Table 3. In experiments, we compute the performance metric as

$$F(\boldsymbol{a}) = \max_{\mathbf{p}_s} V_{\text{emp}}(\hat{\mathbf{p}}_s) \tag{31}$$

Compared to the baseline placement, we can see that the performance metric $F(\boldsymbol{a})$, and maximum and average variances of the localization error are smaller in the optimal anchor geometries for both ROIs in experiment. With the optimal anchor placement, the reductions of the average

Table 3. Localization Performance in Experiments

	F $[m]$	Maximum RMS Variance $[m^2]$	Average RMS Variance $[m^2]$
Optimal geometry 1	0.1142	0.0039	0.0027
Baseline geometry 1	0.1200	0.0046	0.0043
Optimal geometry 2	0.1082	0.0045	0.0033
Baseline geometry 2	0.1255	0.0072	0.0062

estimation variances for Ψ_1 and Ψ_2 are 41.30% and 46.77% respectively. These results reveal that robots can achieve a significantly better localization performance in the entire region of interest with the optimal anchor placement. In addition, the results from this algorithm reach the smallest upper bound of the estimation variance inside the region given the anchor placement constraints. Based on this analysis, we can determine the localization quality of a specific anchor configuration prior to physically installing the anchors, thus saving time and effort.

6. CONCLUSION

The algorithm developed in this paper finds an optimal anchor configuration of a decentralized TDoA-based UWB localization system with respect to a pre-defined region of interest. Leveraging the analysis based on the CRLB and DOP, we imposed a set of soft angle constraints for each UWB anchor pair and benchmarked the localization performance of the region through an A-optimal performance metric, given an anchor placement. The optimal anchor placement was found through the proposed Bayesian optimization-based algorithm. We showed in experiment that the optimal anchor placement reduced the average localization variance and provided a better localization performance inside the entire pre-defined region of interest.

REFERENCES

Betke, M. and Gurvits, L. (1997). Mobile robot localization using landmarks. *IEEE Transactions on Robotics and Automation*, 13(2), 251–263.

Chan, Y.T. and Ho, K. (1994). A simple and efficient estimator for hyperbolic location. *IEEE Transactions on Signal Processing*, 42(8), 1905–1915.

Ennasr, O., Xing, G., and Tan, X. (2016). Distributed time-difference-of-arrival (TDOA)-based localization of a moving target. In *Proc. of the Conf. on Decision and Control (CDC)*, 2652–2658.

Fedorov, V.V. (2013). *Theory of Optimal Experiments*. Elsevier.

Frazier, P.I. (2018). A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*.

Garnett, R., Osborne, M.A., and Roberts, S.J. (2010). Bayesian optimization for sensor set selection. In *Proc. of the ACM/IEEE Intl. Conf. on Information Processing in Sensor Networks*, 209–219.

GPy (since 2012). GPy: A Gaussian process framework in Python. `http://github.com/SheffieldML/GPy`.

Hamer, M. and D'Andrea, R. (2018). Self-calibrating ultra-wideband network supporting multi-robot localization. *IEEE Access*, 6, 22292–22304.

Hoeller, D., Ledergerber, A., Hamer, M., and D'Andrea, R. (2017). Augmenting ultra-wideband localization with computer vision for accurate flight. *IFAC-PapersOnLine*, 50(1), 12734–12740.

Kaune, R. (2012). Accuracy studies for TDOA and TOA localization. In *Proc. of the IEEE Intl. Conf. on Information Fusion*, 408–415.

Ledergerber, A. and D'Andrea, R. (2017). Ultra-wideband range measurement model with Gaussian processes. In *Proc. of the IEEE Conf. on Control Technology and Applications (CCTA)*, 1929–1934.

Ledergerber, A., Hamer, M., and D'Andrea, R. (2015). A robot self-localization system using one-way ultra-wideband communication. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 3131–3137.

Meng, W., Xie, L., and Xiao, W. (2013). Decentralized TDOA sensor pairing in multihop wireless sensor networks. *IEEE Signal Processing Letters*, 20(2), 181–184.

Meng, W., Xie, L., and Xiao, W. (2016). Optimal TDOA sensor-pair placement with uncertainty in source location. *IEEE Transactions on Vehicular Technology*, 65(11), 9260–9271.

Mueller, M.W., Hamer, M., and D'Andrea, R. (2015). Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadrocopter state estimation. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 1730–1736.

Rasmussen, C. and Williams, C. (2005). *Gaussian process for machine learning*. MIT Press.

Santerre, R., Geiger, A., and Banville, S. (2017). Geometry of GPS dilution of precision: revisited. *GPS Solutions*, 21(4), 1747–1763.

Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., and De Freitas, N. (2015). Taking the human out of the loop: A review of Bayesian optimization. *Proc. of the IEEE*, 104(1), 148–175.

Srinivas, N., Krause, A., Kakade, S.M., and Seeger, M. (2009). Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*.

Wang, R., Chen, Z., and Yin, F. (2019). DOA-based three-dimensional node geometry calibration in acoustic sensor networks and its Cramér–Rao bound and sensitivity analysis. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27(9), 1455–1468.

Wang, W., Bai, P., Li, H., and Liang, X. (2018). Optimal configuration and path planning for UAV swarms using a novel localization approach. *Applied Sciences*, 8(6):1001, 1–18.

Yang, B. and Scheuing, J. (2005). Cramér–Rao bound and optimum sensor array for source localization from time differences of arrival. In *Proc. of the IEEE Intl. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, volume 4, 961–964.