

Calibration and Uncertainty Characterization for Ultra-Wideband Two-Way-Ranging Measurements

Mohammed Ayman Shalaby, Charles Champagne Cossette, James Richard Forbes, Jerome Le Ny

Abstract—Ultra-Wideband (UWB) systems are becoming increasingly popular for indoor localization, where range measurements are obtained by measuring the time-of-flight of radio signals. However, the range measurements typically suffer from a systematic error or bias that must be corrected for high-accuracy localization. In this paper, a ranging protocol is proposed alongside a robust and scalable antenna-delay calibration procedure to accurately and efficiently calibrate antenna delays for many UWB tags. Additionally, the bias and uncertainty of the measurements are modelled as a function of the received-signal power. The full calibration procedure is presented using experimental training data of 3 aerial robots fitted with 2 UWB tags each, and then evaluated on 2 test experiments. A localization problem is then formulated on the experimental test data, and the calibrated measurements and their modelled uncertainty are fed into an extended Kalman filter (EKF). The proposed calibration is shown to yield an average of 46% improvement in localization accuracy. Lastly, the paper is accompanied by an open-source UWB-calibration Python library, which can be found at https://github.com/decargroup/uwb_calibration.

I. INTRODUCTION

Robotic localization and mapping applications typically require a means of acquiring position information relative to a reference point with known location. Global Navigation Satellite System (GNSS) provides accurate and precise positioning information outdoors; however, localization performance degrades significantly in obstructed or indoor environments [1, 2]. An attractive indoor localization option that has been increasingly gaining traction is the use of *ultra-wideband* (UWB) radio signals between transceivers, or *tags*, as a means of ranging. UWB transceivers, such as the DWM1000 module provided by Decawave [3], are typically inexpensive, consume little power, and provide a means for data transfer between robots, thus deeming them particularly useful for a variety of robotic applications [4–6].

UWB-based ranging typically relies on measuring the *time-of-flight* (ToF) of radio signals from one tag to another. This requires estimating the offset between the clock on each tag. Furthermore, the clocks often run at different rates due to physical imperfections in the individual clock’s crystal oscillator, causing the offset to be time-varying. The rate

This work was supported by the NSERC Alliance Grant program, by the CFI JELF program, and by the FRQNT.

M. A. Shalaby, C. C. Cossette, and J. R. Forbes are with the department of Mechanical Engineering, McGill University, Montreal, QC H3A 0C3, Canada. {mohammed.shalaby@mail.mcgill.ca, charles.cossette@mail.mcgill.ca, james.richard.forbes@mcgill.ca}. J. Le Ny is with the department of Electrical Engineering, Polytechnique Montreal, Montreal, QC H3T 1J4, Canada. {jerome.le-ny@polymtl.ca}.

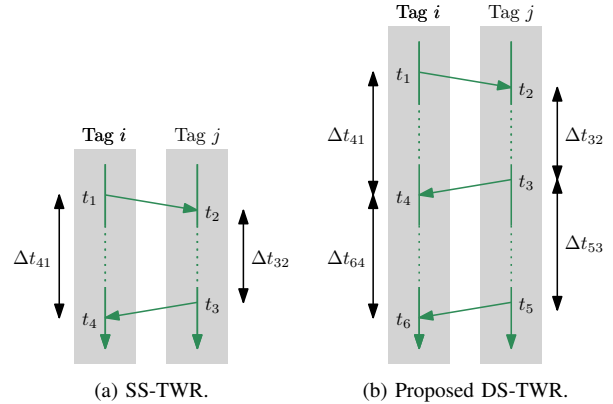


Fig. 1: Timeline schematics for two tags i and j representing the different TWR ranging protocols, where t_ℓ represents the ℓ^{th} timestamp for a TWR instance and $\Delta t_{\ell k} \triangleq t_\ell - t_k$.

of change of the clock offset is referred to as the clock *skew*. In order to negate the effect of the clock offset during ranging, different ranging protocols have been proposed, with the choice being dependent on the specific application and availability of tags [7], [8, Section 7.1.4]. A commonly used protocol is *two-way ranging* (TWR), which relies on averaging out the measured ToF between two signals to negate the clock offset. This form of TWR is referred to as *single-sided TWR* (SS-TWR), and is shown in Figure 1a.

Nonetheless, even after correcting for clock offsets, UWB range measurements typically suffer from a systematic error or *bias*. A significant contributor to this error is the skew between the clocks of the two ranging tags, as the different tags measure the passage of time in different units [9, 10]. This additional bias can be corrected by estimating the clock skew between the tags and embedding a skew-dependent correction factor when computing the range measurement, as proposed in [9]. However, this necessitates estimating the clock skew between all tags involved in ranging. Alternatively, [10] proposes a form of computing the range measurement utilizing *double-sided TWR* (DS-TWR), which is shown to mitigate clock-skew-dependent bias.

Another source of ranging bias stems from relative-pose-dependent antenna radiation pattern [11], where pose refers to both position and attitude. The varying signal strength can cause timestamping errors, and this effect is typically addressed using data-driven models. In [12], a simple experiment with pre-localized fixed tags or *anchors* is used to determine a relation between bias and the distance between ranging tags, while in [13], models are trained using the distance between the tags and 7 features extracted from the channel impulse response (CIR). In [14] and [15], a robot

is flown around in a room with UWB anchors to learn a model of the range bias as a function of the robot's pose. The main drawback of these methods is that the learned model is dependent on the relative poses of the ranging tags, which are typically unknown in real-time without the bias-corrected measurements in the first place. Additionally, the learnt models are trained and tested on the same anchor formations and are therefore not necessarily generalizable; calibration must occur for every new anchor formation. In [9], the former issue is addressed by finding a relation between the bias and the received *first-path power* (FPP) in *line-of-sight* (LOS) conditions with 2D motion.

Delays in communication between the embedded microchip and the UWB antenna are another source of bias [16]. This antenna delay is roughly the same for different UWB tags with the same physical design and is at least a few hundreds of nanoseconds [16], but can vary tenths of a nanosecond or more from tag to tag due to manufacturing inaccuracies. Given that a one-nanosecond timestamping error corresponds to 30 cm in ranging error, the need to perform antenna-delay calibration for every tag is critical. In [16], a basic TWR-based calibration procedure is suggested for calibrating antenna delays. However, the lack of motion introduces a risk of learning the aforementioned relative-pose-dependent bias as antenna delays. In [9], experiments involving a pair of tags at a time ranging with each other is used to fit what is referred to as a "pair-dependent constant". Therefore, the calibration procedure involves calibrating the relative delay between one pair at a time, which does not scale well to systems with many UWB tags.

This paper addresses the problem of calibrating UWB tags, and the main contributions are as follows.

- An alternative DS-TWR protocol is proposed and is shown to mitigate the clock-skew-induced bias.
- A scalable antenna-delay calibration algorithm is presented that is robust to outliers and pose-dependent bias.
- The bias-versus-FPP fit presented in [9] is extended to also address the uncertainty of the measurements as a function of FPP, and DS-TWR is utilized to overcome the need to estimate the clock skew.
- The proposed antenna-delay and bias-FPP calibration are evaluated on an aerial experiment with no anchors, where all the tags are fitted to moving robots.
- The code for the full calibration procedure is attached to this paper as an open-access online repository, which can be found at https://github.com/decargroup/uwb_calibration.

The remainder of this paper is organized as follows. The proposed DS-TWR is discussed in Section II, alongside a theoretical analysis of the clock-skew-dependent bias. In Section III, a robust antenna-delay calibration algorithm is presented, followed by the bias and uncertainty calibration as a function of FPP in Section IV. The calibration methods presented in Sections III and IV are introduced on the same experimental training data, and are then evaluated on 2 testing experiments in Section V.

II. THE RANGING PROTOCOL

UWB ranging relies on the time-of-flight (ToF) of signals between two tags in order to compute range measurements. The simplest way to do this is using SS-TWR, shown in Figure 1a, where the ToF measurement can be computed as

$$t_f = \frac{1}{2}(\Delta t_{41} - \Delta t_{32}). \quad (1)$$

However, different UWB tags have different clocks that are typically running at different rates, and this clock skew results in additional bias in the computed ToF measurement. In [10], an alternative DS-TWR-based ranging protocol is proposed to mitigate clock-skew-dependent bias. In this paper, the DS-TWR protocol shown in Figure 1b is proposed, which differs from [10] by having the responding tag instead of the initiating tag transmit the third signal. The ToF measurement can then be computed as

$$t_f = \frac{1}{2} \left(\Delta t_{41} - \frac{\Delta t_{64}}{\Delta t_{53}} \Delta t_{32} \right). \quad (2)$$

This protocol is motivated by the intuitive understanding that the additional correcting factor in (2) transforms Δt_{32} from time units of the receiver tag's clock to time units of the initiator tag's clock. Additionally, the proposed ranging protocol allows the initiating tag to process the range measurement by computing (2), without requiring additional signals for the responding tag to send Δt_{32} and Δt_{53} .

A. Analytical Bias Model

To demonstrate clock-skew-dependent bias, consider in SS-TWR the clock-skew-corrupted ToF measurement,

$$\tilde{t}_f^{\text{ss}} = \frac{1}{2} \left((1 + \gamma_i)(\Delta t_{41} + \eta_{41}) - (1 + \gamma_j)(\Delta t_{32} + \eta_{32}) \right), \quad (3)$$

where γ_i is the skew of Tag i 's clock relative to real time, $\eta_{kl} = \eta_k - \eta_l$, and $\eta_k, \eta_l \sim \mathcal{N}(0, R)$ are mutually-independent timestamping white noise associated with timestamps t_k and t_l , respectively. The ToF error is thus

$$\begin{aligned} e^{\text{ss}} &\triangleq \tilde{t}_f^{\text{ss}} - t_f \\ &= \frac{1}{2} (\gamma_i \Delta t_{41} + (1 + \gamma_i) \eta_{41} - \gamma_j \Delta t_{32} - (1 + \gamma_j) \eta_{32}), \end{aligned} \quad (4)$$

and the expected value of e^{ss} is

$$\begin{aligned} \mathbb{E}[e^{\text{ss}}] &= \frac{1}{2} (\gamma_i \Delta t_{41} - \gamma_j \Delta t_{32}) \\ &\stackrel{(1)}{=} \frac{1}{2} (\gamma_i (2t_f + \Delta t_{32}) - \gamma_j \Delta t_{32}) \\ &= \gamma_i t_f + \frac{1}{2} (\gamma_i - \gamma_j) \Delta t_{32}. \end{aligned} \quad (5)$$

The first component of (5) is negligible as skew is in the order of parts-per-million and ToF in nanoseconds. However, Δt_{32} is typically in hundreds of microseconds, meaning that clock-skew-dependent bias is not negligible.

Negating the second component of (5) is the motivation behind the proposed ranging protocol. Rewriting (2) as

$$t_f = \frac{\Delta t_{41} \Delta t_{53} - \Delta t_{64} \Delta t_{32}}{2 \Delta t_{53}},$$

and following the same steps as in (3)-(4), the ToF error for the proposed DS-TWR can be derived to be approximately

$$e^{\text{ds}} \triangleq \tilde{t}_f^{\text{ds}} - t_f \approx \gamma_i t_f + \frac{1}{2} (1 + \gamma_i) \left[\frac{\Delta t_{32}}{\Delta t_{53}} (\eta_{53} - \eta_{64}) + \eta_{41} - \eta_{32} \right], \quad (6)$$

where \tilde{t}_f^{ds} is the clock-skew-corrupted time-of-flight measurement. Deriving (6) relies on the assumptions that $\Delta t_{53} \gg \|\eta_{53}\|$, $\Delta t_{32} \gg \|\eta_{32}\|$. This is expected since the timestamping error is typically in the order of nanoseconds or less, and the delay intervals are in the order of hundreds of microseconds. The expected value of the error is therefore

$$\mathbb{E}[e^{\text{ds}}] = \gamma_i t_f,$$

which suffers from less bias when compared to the error of the SS-TWR protocol.

III. ANTENNA-DELAY CALIBRATION

The delay between a chip timestamping transmission and the antenna actually transmitting the signal is referred to as the antenna transmission delay d^t , while the delay between an antenna receiving a signal and the chip timestamping reception is the antenna reception delay d^r . Looking back at Figure 1b, the measured time-stamps are therefore

$$\tilde{t}_k = t_k + d^t, \quad k \in \{1, 3, 5\}, \quad (7)$$

$$\tilde{t}_\ell = t_\ell + d^r, \quad \ell \in \{2, 4, 6\}. \quad (8)$$


In this section, a scalable antenna-delay calibration procedure is presented that addresses the need for incorporating motion. In particular, a linear least-squares approach is presented utilizing DS-TWR, which is solved using robust least squares to accommodate for outliers [17, Section 5.3.2]. Both the antenna delays and the presented approach are environment independent, and therefore the antenna-delay calibration only needs to be performed once for new uncalibrated transceivers.

A. Least Squares Formulation

The goal of antenna-delay calibration is to find the best-fit delays based on some collected data. In order to perform this calibration procedure, the tags to be calibrated must be capable of DS-TWR. The effect of antenna delays on the DS-TWR ToF measurements is shown by substituting (7) and (8) into (2), thus yielding

$$t_f = \frac{1}{2} \left(\Delta \tilde{t}_{41} + \underbrace{d_i^t - d_i^r}_{d_i} - \frac{\Delta \tilde{t}_{64}}{\Delta \tilde{t}_{53}} (\Delta \tilde{t}_{32} - \underbrace{d_j^t + d_j^r}_{-d_j}) \right) \quad (9)$$

when Tag i initiates with Tag j , where d_i^t and d_i^r are the antenna transmission and reception delays of Tag i , respectively. In this case, transmission and reception delays can be combined into one delay variable d_i to be estimated for every tag, where $d_i = d_i^t - d_i^r$. This is sufficient for systems where only TWR is utilized, which is the focus of this paper. When other ranging protocols are implemented such as time-difference-of-arrival (TDoA) or time-of-arrival (ToA),



(a) Quadcopter equipped with two UWB tags. (b) Snapshot from training experiment.

Fig. 2: Experimental set-up for antenna-delay and bias-FPP calibration.

another antenna-delay calibration procedure is necessary to solve for d_i^t and d_i^r separately.

In the presence of n tags to be calibrated, let \mathcal{P} denote the ordered set of tuples representing all ranging pairs of tags. Consequently, the antenna delays are calibrated by formulating a linear least-squares problem as

$$\hat{\mathbf{d}} = \arg \min_{\mathbf{d} \in \mathbb{R}^n} \sum_{(i,j) \in \mathcal{P}} \sum_{k=1}^{m_{ij}} g(e_{ij}^k(\mathbf{d})), \quad (10)$$

where the error e_{ij}^k is defined from (9) as

$$e_{ij}^k(\mathbf{d}) = \frac{1}{2} (d_i + K^k d_j) - t_f^k + \frac{1}{2} (\Delta \tilde{t}_{41}^k - K^k \Delta \tilde{t}_{32}^k),$$

m_{ij} is the number of range measurements between Tags i and j , $\mathbf{d} = [d_1 \cdots d_n]^\top$, the superscript k denotes the k^{th} measurement, and $K^k \triangleq \Delta \tilde{t}_{64}^k / \Delta \tilde{t}_{53}^k$. Moreover, g is the loss function, and the choice of g is discussed in Section III-B.

If $n = 2$, the formulated least-squares problem would have 2 unknowns and only 1 pair of ranging tags, which results in non-uniqueness of the solution. Therefore, the calibration procedure should involve at least 3 tags, yielding 3 unknowns and 3 pairs of ranging tags.

B. Experimental Results on Training Data

To evaluate the proposed antenna-delay calibration procedure experimentally, three Uvify IFO-S quadcopters are equipped with 2 UWB tags each as shown in Figure 2a, for a total of 6 tags. In order to compute the theoretical ToF t_f for any measurement to formulate a similar problem to (10), a motion-capture system is used to get the ground truth distances between the ranging tags. A motion-capture system is chosen for its mm-accuracy, but any other localization approach would suffice, with the accuracy of the calibration depending on the accuracy of the localization algorithm. Unlike the static experiments suggested in [16], this allows a dynamic experiment where the quadcopters fly randomly in 3-dimensional space as shown in Figure 2b, which reduces the proneness to learning relative-pose-dependent biases. This dataset consists of 4 minutes of flight time and a total of 38000 range measurements. The calibration procedure can be done by fitting each drone with one tag, but it is common in localization problems to fit two tags to overcome the lack of bearing information [5, 18, 19]. Nonetheless, the ranging schedule does not directly involve TWR measurements between any two tags on the same drone. Therefore, there are 6 unknown delays and 12 pairs of ranging tags.

Typically, problems of the form (10) are solved by finding \mathbf{d} that minimizes the squared error (i.e., choosing g to be

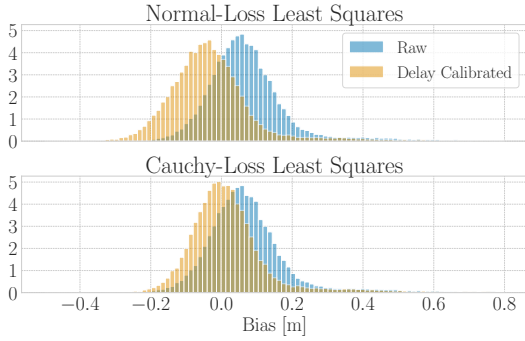


Fig. 3: Histograms showing the effect on the ranging bias post-antenna-delay calibration. (Top) Using L2 loss. (Bottom) Using Cauchy loss.

L2 loss), which is derived from an assumption that the underlying distribution of the noise is Gaussian. However, UWB measurements suffer from positive outliers due to multipath propagation and other sources of error, and are better modelled using Cauchy distributions [20]. It is therefore proposed for this particular application to minimize the Cauchy loss $g(x) = \log(0.5x^2 + 1)$ instead to solve (10) while reducing the effect of outliers [17, Section 5.3.2].

The use of the Cauchy loss is found to be useful for mitigating the effect of positive outliers as the mode of the bias distribution becomes 0, as shown in Figure 3. Moreover, the precision of the proposed approach is evaluated by comparing antenna-delay solutions derived from 2 datasets collected one month apart, and the delays converge to the same result within 0.03 ns of accuracy, which corresponds to less than 1 cm of ranging error.

C. Using Calibrated Tags to Calibrate New Tags

In [9], the sum of the antenna delays for every pair of tags (i, j) is calibrated as one constant c_{ij} , which for the DS-TWR protocol would be of the form $c_{ij} = d_i + Kd_j$. This approach requires that calibration is done for every pair of tags, which is tedious and not scalable. That is due to the fact that lumping the delay terms into one constant fails to utilize the constant delay terms that appear in different pairs.

By solving for the aggregate antenna delays d_i and d_j individually, this allows calibrating a new tag without collecting data between the new tag and all previously calibrated tags. In order to calibrate a new tag Tag j , only one calibrated tag Tag i is required, which then allows solving for d_j from (9),

$$\hat{d}_j = \frac{2t_f - \Delta\tilde{t}_{41} - \hat{d}_i}{K} + \Delta\tilde{t}_{32}.$$

However, it is still recommended to collect more data using a dynamic experiment with the two tags in order to improve robustness to noise and pose-dependent bias.

IV. POWER-CORRELATED CALIBRATION

Another source of error in UWB-based ranging is irregularities in the antenna radiation pattern and system design elements, such as PCB-induced losses. Typically, such losses introduce biases in the measurements that are pose-dependent and that are correlated with the received signal power [9,

11]. In this section, the experiments in [9] are extended in the following ways.

- 1) The proposed DS-TWR is used rather than SS-TWR, which overcomes the need to estimate the skew between all pairs of tags.
- 2) The results are shown to hold for experiments in three-dimensional space.
- 3) The results are shown to hold for experiments with some non-LOS measurements due to occlusions from the quadcopters' bodies.
- 4) The individual measurements are used in the data-fitting process rather than averaging out measurements from a discrete number of relative poses.

The last point is particularly important as it overcomes the need to remain static during data collection, which allows the calibration procedure to be a simple experiment of robots moving randomly and covering as many relative poses as possible in a relatively short period of time. Another advantage of using all the data in the calibration process is that there is no loss of variance information through averaging out similar measurements. Consequently, the relation between the variance of the measurements and the received signal power can then be analyzed.

A. Bias Calibration

The bias calibration procedure is similar to the one presented in [9]. The reception timestamp at Tag i is usually corrupted by an unknown function $\rho_i(\cdot)$ of the received FPP p^f ; therefore, from (2),

$$\begin{aligned} t_f &= \frac{1}{2} \left(\Delta\tilde{t}_{41} + \rho_i(p_4^f) \right. \\ &\quad \left. - \frac{\Delta\tilde{t}_{64}}{\Delta\tilde{t}_{53} + \rho_i(p_6^f) - \rho_i(p_4^f)} (\Delta\tilde{t}_{64} - \rho_j(p_2^f)) \right) \\ &\stackrel{(a)}{\approx} \frac{1}{2} \left(\Delta\tilde{t}_{41} - \frac{\Delta\tilde{t}_{64}}{\Delta\tilde{t}_{53}} \Delta\tilde{t}_{64} \right) + \frac{1}{2} (\rho_i(p_4^f) + \rho_j(p_2^f)) \\ &\triangleq \frac{1}{2} \left(\Delta\tilde{t}_{41} - \frac{\Delta\tilde{t}_{64}}{\Delta\tilde{t}_{53}} \Delta\tilde{t}_{64} \right) + f \left(\Psi \left(\frac{p_4^f + p_2^f}{2} \right) \right), \end{aligned} \quad (11)$$

where p_i^f is the FPP associated with timestamp t_i , and $\Psi(x) \triangleq 10^{(x-\alpha)/10}$ is the lifting function suggested in [9] with α as a normalization parameter. Moreover, $f(\cdot)$ is an unknown function to be learned from data, defined based on an experimentally-motivated assumption that the effects of power-correlated bias due to the individual tags can be aggregated into one function of the average received FPP that is common to all tags of similar design. The step (a) in the derivation involves the assumptions that $\rho_i(p_6^f) = \rho_i(p_4^f)$ and $\frac{\Delta\tilde{t}_{64}}{\Delta\tilde{t}_{53}} \rho_j(p_2^f) \approx \rho_j(p_2^f)$. The former assumption is due to the fact that the motion of the robots is negligible in the time window Δt_{64} and therefore the relative-pose between the two tags is similar, while the later assumption is due to $\rho_j(p_2^f)$ being in the order of tenths of nanoseconds, and therefore $\left(1 - \frac{\Delta\tilde{t}_{64}}{\Delta\tilde{t}_{53}}\right) \rho_j(p_2^f) \approx 0$.

Referring back to the experiment mentioned in Section III, the function f is learned by simply fitting a spline to the

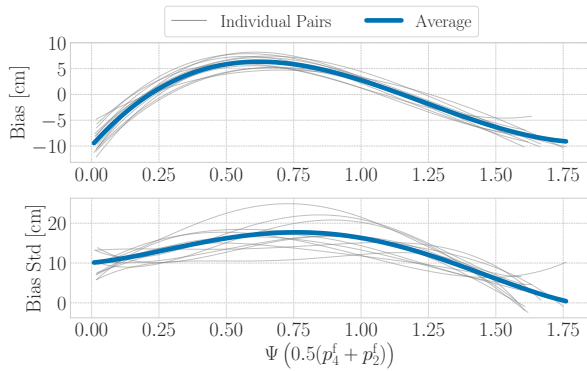


Fig. 4: The fitted bias and standard deviation curves as a function of the lifted average FPP, using a 4-minute long training experiment.

post-antenna-delay-calibration range bias as a function of the lifted average FPP. The result is shown at the top of Figure 4 for all individual pairs as well as for all the data, where the ToF bias is converted to range bias. As expected, when the antenna delay is corrected first, the bias-power curve is similar for all pairs as they all use the same antenna and PCB-board design. In this case, this additional calibration procedure can remove up to 10 cm of bias, but even though this calibration procedure is environment independent [9], this would vary for different tag designs and this process must be done separately for different antenna/board designs.

B. Variance Calibration

The variance of the measurements is also expected to vary as a function of the received-signal power, that is,

$$\mathbb{E}[(\tilde{t}_f - t_f)^2] \triangleq \sigma \left(\Psi \left(\frac{p_4^f + p_2^f}{2} \right) \right)^2.$$

Intuitively, it is expected that the receiver should be able to detect and timestamp the direct-path signal more accurately when the FPP is high as this indicates a high signal-to-noise ratio (SNR). Additionally, multipath and obstacle-attenuated signals typically have lower FPP, and therefore not a lot of variance is expected at higher received FPP. At lower received FPP, the SNR is lower, and the received measurements might have been corrupted with equally-powerful multipath and body-attenuated signals.

In order to analyze this experimentally using the training data, a similar procedure to the power-bias calibration step is proposed. The standard-deviation samples are generated by computing the standard deviation of the range bias of the measurements in a window of FPP. A spline is fitted to the standard-deviation samples, and the resulting curves are shown at the bottom of Figure 4. As expected, the lowest standard deviation is at the highest FPP, where the standard deviation of the range bias is as low as 2.5 cm. Additionally, the highest standard deviation of approximately 17 cm is in the mid-FPP region. This is potentially due to reflections off the ground being primarily in this region.

Even though there is a clear trend, the standard deviation curves appear to somewhat vary between different pairs. This is partly due to the training experiment being relatively short,




Fig. 5: Distribution of the testing-data range bias pre- and post-calibration, with collection bins indicating the amount of outliers beyond the axes.

and at some FPP values there is not enough data points to accurately compute the standard deviation. Additionally, despite the curve seeming to slope downwards at lower FPP values, it is expected that at some point as the received FPP value decreases beyond the lower detection threshold the variance will increase drastically. However, the lower detection threshold is chosen to be higher than the point where random meaningless signals would be detected, and therefore the point where the variance increases significantly does not appear in the recorded experiment.

V. EXPERIMENTAL RESULTS ON TESTING DATA

The proposed calibration procedure from Sections III and IV is evaluated on 2 testing experiments with the same set-up as the training experiment presented in Section III-B. While in the training experiment the quadcopters follow a more structured trajectory, the testing experiments involve the quadcopters flying around the 3-dimensional space randomly. Each testing experiment consists of 60 seconds of flight time and 10000 range measurements between the 12 pairs.

A. Bias Correction

The distributions of the resulting biases in the testing data pre- and post-calibration are shown in Figure 5. The mean of the bias of the raw measurements is reduced by 36% through antenna-delay calibration only and by a further 20% by fully calibrating the measurements, bringing the mean from 11.11 cm to 5.91 cm. The standard deviation of the measurements is barely affected by antenna-delay calibration, but is reduced approximately 6% through power-correlated calibration from 18.95 cm to 17.82 cm. Both the mean and standard deviation are affected by positive outliers potentially resulting from non-LOS and multipath propagation.

In order to reject outliers, the underlying distribution must be known. Through the variance calibration procedure, the range measurements are assumed to be corrupted with zero-mean Gaussian noise with a standard deviation given by (12). An outlier can be rejected if it does not satisfy the underlying distribution with a certain degree of confidence. For each individual measurement k with ground-truth-computed bias b_k and power-correlated standard deviation σ_k from (12), this can be done by performing the *chi-squared test* [21, Section 1.4.17]. Any measurement that does not satisfy the

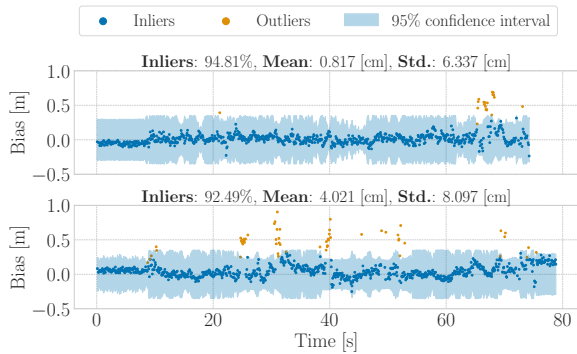


Fig. 6: The range bias for 2 ranging pairs in one testing experiment after calibration, with some statistical information.

inequality $\frac{b_k^2}{\sigma_k^2} \leq \gamma$ indicates that it is not from the underlying distribution to a certain degree of confidence. The threshold γ depends on the chosen degree of confidence, typically 95%.

The results of this standard outlier-rejection algorithm are shown in Figure 6. An indication of the calibrated variance being close to the actual underlying distribution would be that exactly 5% of measurements are rejected; however, due to other factors such as non-LOS and multipath propagation, more than 5% of measurements are rejected for some pairs. Even in this non-ideal scenario with non-LOS and multipath, the mean of the measurements after outlier rejection among all pairs reduces to a maximum of roughly 4 cm and a maximum standard deviation of approximately 8 cm.

B. Position Estimator

In real-world applications, the ground truth distance between tags is usually not known but is rather estimated, and the outlier rejection method is usually done using the *normalized-innovation-squared* (NIS) test [21, Section 5.4.2]. The NIS test is similar to the chi-squared test mentioned in Section V-A, but additionally accommodates for uncertainty in the state estimates.

To evaluate the variance calibration using the NIS test, the following simple localization problem is formulated using the testing data. Consider the problem of estimating \mathbf{r}_a^{1w} , the position of Robot 1 relative to some arbitrary reference point w , resolved in some inertial frame \mathcal{F}_a . There are two tags on the robot, and \mathbf{r}_1^{i1} represents the position of Tag i relative to the robot's reference point in the robot's own body frame \mathcal{F}_1 , which can be measured manually. Additionally, assume that the orientation of the robot given by a direction cosine matrix $\mathbf{C}_{1a} \in SO(3)$ is known, and that velocity measurements \mathbf{v}_a^{1w} are available from the motion-capture system. Assuming that poses and tag positions of the neighbouring robots $n \in \{2, 3\}$ are known, an extended Kalman filter (EKF) is used to estimate \mathbf{r}_a^{1w} , where the measurements are modelled as

$$y = \|\mathbf{r}_a^{1w} + \mathbf{C}_{1a}^T \mathbf{r}_1^{i1} - \mathbf{r}_a^{nw} - \mathbf{C}_{na}^T \mathbf{r}_n^{jn}\| + \nu$$

for the range measurement between Tag i on Robot 1 and Tag j on Robot n , and $\nu \sim \mathcal{N}(0, R)$ is white Gaussian noise. The NIS test is used in the filter for outlier rejection.

The performance of the filter is shown based on the *root-mean-squared-error* (RMSE) metric in Figure 7 for one




Fig. 7: Comparison of the position-estimate RMSE using raw and calibrated data for one testing-data scenario.

TABLE I: RMSE comparison for the raw and calibrated measurements in all 6 testing-data scenarios.

		RMSE [m]		
	Robot no.	Raw	Calibrated	Calibrated w/ Variance
Experiment 1	1	0.9245	0.5122	0.4243
	2	0.7577	0.5095	0.4437
	3	0.6731	0.3872	0.3467
Experiment 2	1	0.8664	0.4879	0.4932
	2	0.8174	0.5854	0.3980
	3	0.7837	0.5167	0.4788

experimental run and summarized in Table I for 6 different scenarios using 1) the raw measurements and fixed variance, 2) the calibrated measurements and fixed variance, and 3) the calibrated measurements and the calibrated variance-power curve. The 6 scenarios represent a variation of which of the three robots is the one with an unknown position to be estimated, and doing so in two different experimental runs. The choice of fixed measurement variance is decided experimentally based on what consistently yields the best performance. On average for the 6 different scenarios, the antenna-delay and bias-calibration procedures alone yield a 38% improvement in localization accuracy, while additionally utilizing the power-correlated variance calibration results in an average of 46% reduction in the RMSE, thus emphasizing the importance of calibrating UWB sensors and the added benefit of using the received FPP as an indication of the uncertainty of measurements.

VI. CONCLUSION

In this paper, the problem of calibrating UWB bias is addressed. To eliminate the need for estimating the clock states, a DS-TWR-based ranging protocol is presented and shown to theoretically mitigate the effect of clock-skew-dependent bias. Furthermore, a robust and scalable antenna-delay calibration procedure is presented and trained on data from an aerial experiment. A model is then learnt on the experimental data to find the relation between the remaining bias and the uncertainty of the measurements as a function of the received-signal power. The delays and models learnt are then applied to two testing experiments to evaluate the calibration procedure. A localization problem is then formulated using an EKF, and it is shown that an average of 46% improvement in localization accuracy can be achieved by using the corrected measurements and the modelled variance.

REFERENCES

- [1] P. D. Groves, "Shadow matching: A new GNSS positioning technique for urban canyons," *Journal of Navigation*, vol. 64, no. 3, pp. 417–430, 2011.
- [2] A. T. Irish, J. T. Isaacs, F. Quidin, J. P. Hespanha, and U. Madhoo, "Probabilistic 3D mapping based on GNSS SNR measurements," *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2390–2394, 2014.
- [3] Decawave, DW1000 Radio IC. <https://www.decawave.com/product/dw1000-radio-ic/>.
- [4] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultrawideband range measurements with accelerometers and rate gyroscopes for quadcopter state estimation," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 1730–1736, 2015.
- [5] B. Hepp, T. Nægeli, and O. Hilliges, "Omni-directional person tracking on a flying robot using occlusion-robust ultrawideband signals," *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, vol. 2016–November, pp. 189–194, 2016.
- [6] Y. Cao, C. Yang, R. Li, A. Knoll, and G. Beltrame, "Accurate position tracking with a single UWB anchor," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 2344–2350, 2020.
- [7] Z. Sahinoglu, S. Gezici, and I. Gvenc, *Ultra-wideband Positioning Systems: Theoretical Limits, Ranging Algorithms, and Protocols*. Cambridge University Press, 2008.
- [8] P. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems, Second Edition*. Artech House Publishers, 2013.
- [9] J. Cano, G. Pages, E. Chaumette, and J. Le Ny, "Clock and Power-Induced Bias Correction for UWB Time-of-Flight Measurements," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 2, pp. 2431–2438, 2022, ISSN: 23773766.
- [10] D. Neiryck, E. Luk, and M. McLaughlin, "An alternative double-sided two-way ranging method," *13th Workshop on Positioning, Navigation and Communication, WPNC 2016*, pp. 16–19, 2017.
- [11] Decawave, "Sources of error in DW1000 based two-way ranging (TWR) schemes," pp. 1–21, 2014.
- [12] J. González, J. L. Blanco, C. Galindo, A. Ortiz-de Galisteo, J. A. Fernández-Madrigal, F. A. Moreno, and J. L. Martínez, "Mobile robot localization based on Ultra-Wide-Band ranging: A particle filter approach," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 496–507, 2009.
- [13] H. Wymeersch, S. Marano, W. M. Gifford, and M. Z. Win, "A machine learning approach to ranging error mitigation for UWB localization," *IEEE Trans. on Communications*, vol. 60, no. 6, pp. 1719–1728, 2012.
- [14] A. Ledergerber and R. D'Andrea, "Calibrating Away Inaccuracies in Ultra Wideband Range Measurements: A Maximum Likelihood Approach," *IEEE Access*, vol. 6, pp. 78 719–78 730, 2018.
- [15] J. Zhu and S. S. Kia, "Bias Compensation for UWB Ranging for Pedestrian Geolocation Applications," *IEEE Sensors Letters*, vol. 3, no. 9, pp. 7–10, 2019.
- [16] Decawave, "APS014: DW1000 Antenna Delay Calibration Version 1.2," 2018.
- [17] T. Barfoot, *State estimation for Robotics*. Cambridge University Press, 2017.
- [18] T. M. Nguyen, A. H. Zaini, C. Wang, K. Guo, and L. Xie, "Robust Target-Relative Localization with Ultra-Wideband Ranging and Communication," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 2312–2319, 2018.
- [19] M. Shalaby, C. C. Cossette, J. R. Forbes, and J. Le Ny, "Relative Position Estimation in Multi-Agent Systems Using Attitude-Coupled Range Measurements," *IEEE Robotics and Automation Letters (RA-L)*, vol. 6, no. 3, pp. 4955–4961, 2021, ISSN: 23773766.
- [20] M. Kok, J. D. Hol, and T. B. Schon, "Indoor positioning using ultrawideband and inertial measurements," *IEEE Trans. on Vehicular Technology*, vol. 64, no. 4, pp. 1293–1303, 2015, ISSN: 00189545.
- [21] Y. Bar-Shalom, T. Kirubarajan, and X.-R. Li, *Estimation with Applications to Tracking and Navigation*. USA: John Wiley & Sons, Inc., 2002.