ALADIn: Autonomous Linear Antenna Delay Inference on Resource-Constrained Ultra-Wideband Devices

Patrick Rathje[1], Olaf Landsiedel[1, 2]
[1]Kiel University, Kiel, Germany
[2]Chalmers University of Technology, Gothenburg, Sweden
{pra,ol}@informatik.uni-kiel.de

Abstract

Enabling precise indoor localization in a cheap and small package, Ultra-Wideband (UWB) transceivers bring decimetre-accurate ranging to resource-constrained IoT devices. Due to hardware-induced signal processing delays, device-specific antenna calibration enables the most accurate ranging results. This work introduces ALADIn for estimation and calibration of antenna delays in an autonomous manner, removing the need for manual labor and external hardware or computation. Based on known geometry, our approach allows already deployed devices to utilize their ranging and computational capabilities to optimize delays and reduce ranging errors autonomously. At its heart, ALADIn combines an efficient all-to-all ranging primitive with ordinary least squares inference. We conduct both extensive simulations and on-site evaluations. Our simulation results indicate that the proposed approach performs similarly to available calibration methods while being computationally less expensive. Deployed on three testbeds, we analyze the calibration performance on up to 14 DWM1001 devices. For one, the proposed calibration reduces mean absolute error alongside the standard deviation: from uncalibrated 10.7 (7.0 *SD*) cm to 6.7 (4.2 *SD*) cm, which is also lower than the 8.1 (5.7 *SD*) cm of error induced by factory-calibrated values. In addition, our results highlight the quality of measurements, i.e., pairwise variances, and further exhibit the potential of excluding multipath-affected links from the estimation process. Our implementation builds on Zephyr RTOS and is released as open-source.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*; C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems

General Terms

Algorithms, Design, Measurement

Keywords

Ultra-Wideband, All-to-All Ranging, Two-Way-Ranging, Antenna Delay Calibration, Indoor Localization

1 Introduction

IoT devices are ubiquitous and blend into our smart homes, hospitals, and factories. They simultaneously open up the opportunity for location-dependent services in health services, disaster management, security, or tracking [24]. Demanding high accuracy, indoor localization for IoT devices builds on alternative technologies, as GNSS multilateration is generally unavailable indoors. Since walls result in reflections that often induce multipath fading, Ultra-wideband (UWB) radio signals offer a practical solution due to their inherent multipath resistance resulting from short pulses in a wide frequency span. With their low price and small form factor, UWB transceivers allow decimetre-accurate ranging even on resource-constrained IoT devices. Employing the Two-Way Ranging (TWR) protocol [15], devices depend on precise timestamps of sent and received packets; each nanosecond of deviation results in approximately 30 cm of error in distance estimations. However, device-specific delays in the signal processing hardware skew timestamps and demand an individual calibration of each antenna's delays for optimal accuracy. For one, in the case of the widely-used DWM1001 UWB module, calibrating the antenna delays reduces errors from around 30 cm to 4.5 cm [6].

When factory calibration is unavailable, nodes require manual calibration: Based on the known geometry of devices, i.e., their inter-distances and pairwise TWR ranging, existing works employ computationally expensive algorithms for optimization like Particle-Swarm-Optimization or Gauss-Newton, solving nonlinearities [4, 17, 20]. Other works utilize a custom Time-Difference of Arrival (TDoA) approach to isolate single antenna delays [18]. In all cases, the calibration requires additional hardware and external computation capabilities to store and analyze the data.

Instead, we envision executing the inference on the devices, resulting in an autonomous calibration of antenna delays. Hence, our goal is to utilize the ranging and computing capabilities of resource-constrained IoT devices, allowing deployments to continuously self-calibrate in place, si-

multaneously removing the need for labor-intensive, manual calibration. However, UWB measurements are inherently noisy [11], and clock drifts require proper mitigation [12]. Further, resource constraints restrict the execution of complex algorithms.

This work presents and evaluates ALADIn, a procedure for self-calibration of antenna delays in UWB localization systems. We construct a linear ranging model which captures the connection of antenna delays and measurement errors. Solved on-device using least squares optimization, this linear system relates available pairwise measurements and node delays. On the system side, our work tailors all-to-all ranging to mitigate clock drifts and aggregate measurements, i.e., their mean and variance, of all device pairs. We implement and evaluate the protocol for constrained DWM1001 UWB devices equipped with 64 kByte of memory. Using remote UWB testbeds in Lille [1] and Trento [13], which offer ground truth positions and known factory-calibrated antenna delays, we study the calibration capabilities.

Our results indicate that our derived linear approach executes efficiently on devices, decreasing errors of pairwise measurement with estimated delays. Compared to uncalibrated devices, the proposed calibration reduces the mean absolute error (MAE) of 7 nodes in Trento from 10.7 cm to 6.7 cm, simultaneously decreasing the standard deviation from 7.0 cm down to 4.2 cm. Further, the results suggest benefits even for already established localization systems, as errors are lowered compared to factory-calibrated values with an error of 8.1 (5.7 *SD*) cm.

The major contributions of this work summarize as follows:

- Autonomous calibration of UWB antenna delays on resource-constrained devices using efficient all-to-all ranging, named ALADIn.
- Derivation of a linear model that combines antenna delays with inter-device distances in the TWR setting.
- Open-Source implementation ALADIn's ranging and on-device calibration algorithm [1].
- Evaluation in simulations and on three deployments in Lille and Trento, showing comparable estimation performance to non-linear approaches in simulations and reduction of mean pairwise ranging errors in testbed deployments by up to 4.0 cm compared to uncalibrated and 1.4 cm relative to factory-calibration, also lowering the standard deviation by 1.2 cm.

The outline of this work is as follows: Section 2 presents the background on UWB, Two-Way Ranging, and Linear Least Squares estimation. Section 3 lists related work and puts this work into context. Section 4 contains the theoretical part of this work, including the linear ranging model and the formulation of the linear least squares problem. Section 5 lays out the systems view, including the all-to-all-ranging protocol and implementation details. Section 6 holds the evaluation results of the simulations and practical testbed deployments. Finally, Section 7 concludes this work.

[1] Available at `https://github.com/ds-kiel/aladin-uwb`

2 Background

This section introduces the central concept of ranging with Ultra-Wideband and briefly summarizes Linear Least Square Estimation.

2.1 Two-Way Ranging over Ultra-Wideband

Standardized by IEEE 802.15.4a [10], Ultra-Wideband (UWB) transmissions consist of radio frequency pulses with extremely short durations over a large frequency bandwidth. UWB transceivers, like the IEEE 802.15.4 compliant DWM1000, capture transmission and reception timestamps with a precision of 15.65 ps. Hence, a simple back-and-forth packet exchange, called Two-Way Ranging (TWR), between two devices allows the estimation of the time-of-flight (ToF) of a wireless transmission, which directly relates to the distance when multiplied by the speed of light.

However, even minor clock deviations between the devices challenge the resulting estimation accuracy, with each nanosecond of clock deviation resulting in approximately 30 cm of error. Estimating these relative drifts allows mitigation of this error. For one, the offset in the carrier frequency relates to the relative drift [7]. Another method is Alternative Double-Sided Two-Way Ranging (Alt-DS-TWR) which adds another message for mitigation [15]. Figure 1 displays the message exchange for traditional Single-Sided-TWR and DS-TWR.

In addition to device-specific clock drifts, hardware-induced antenna delays contribute to this error. Ideally, devices individually undergo delay calibration in a controlled environment, i.e., within the factory, eliminating external influences. In the case of the DWM1001 platform, antenna delays vary slightly with temperature, skewed by 2.15mm / °C, and with battery voltage by 5.35 cm per V [6]. Calibration of antenna delays, however, reduces the related ranging error from 30 cm to 4.5 cm in 99.7% on 2000 sample devices [6]. Lastly, measurements are further skewed by a range bias: strong signals are processed faster by the signal processing hardware. Correction values are available for the DWM1001's DW1000 chip [6, 5].

2.2 Linear Least Squares Estimation

The Linear Least Squares estimation minimizes the sum of squared residuals in a stochastic linear model of the form:

$$Y = X\beta + \varepsilon$$

with measurement matrix $Y \in \mathbb{R}^{m\times 1}$, design matrix $X \in \mathbb{R}^{m\times p}$, parameter matrix $\beta \in \mathbb{R}^{p\times 1}$ and independent, multivariate centered noise $\varepsilon \sim \mathcal{N}(0, \sigma^2 \mathcal{I}_m)$.

For a full rank matrix X, i.e. $rank(X) = p$, we get an unbiased estimation of parameters $\hat{\beta} \in \mathbb{R}^{p\times 1}$ that minimizes the sum of squared residuals as follows:

$$\hat{\beta} := (X^T X)^{-1} X^T Y$$

3 Related Work

The device-specific antenna delay directly influences all active rangings of a device. However, proper calibration reduces the influence and boosts the overall accuracy of measurements [6].

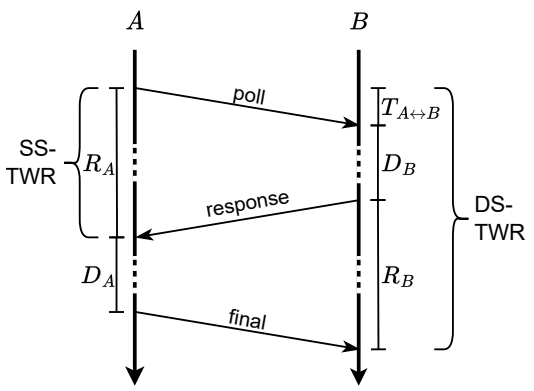

Figure 1: The message exchange in Two-Way Ranging (TWR) protocols: Measuring the round times and delays, involved parties *A* and *B* estimate the time-of-flight and their distance despite clock offsets. The addition of a third (final) message extends the traditional Single-Sided TWR (SS-TWR) to Double-Sided TWR (DS-TWR) and mitigates the effect of relative clock drifts.

In the best case, the geometry of the anchors is known, allowing self-calibration of antenna delays using the difference in measured and actual distances. Using the Alt-DS-TWR protocol, Gui et al. derive a non-linear system relating the error between the actual and the measured distances to the antenna delays [8]. Similar to the instruction of the manufacturer [6], measurements are first collected between all devices. Knowing the resulting error but having a non-linear dependency, they employ Particle-Swarm-Optimization (PSO) and search for an acceptable estimation of the combined antenna delay of both devices per connection, i.e., the best candidates of combined delays for each node pair. Due to the nonlinearity, all measurements are recorded and retained in the optimization process. Then, least squares optimization distributes the aggregate delays to the individual devices. Piavanini et al. improve the computational complexity by using Gauss-Newton (GN) optimization resulting in superior performance compared to the PSO approach [16]. However, the PSO and the GN approach individually solve the non-linear systems for each of the $n(n-1)/2$ pairs, demanding substantial resources. Afterward, both cases use an ordinary least-squares approach to estimate individual antenna delays.

A different approach is pursued by Shah et al.: Based on their proposed Multi-Simultaneous ranging [19], the authors derive techniques for antenna delay calibration [18]. For every connected three nodes, one node remains passive, measuring the inter-arrival time between the TWR exchange of the other two. Based on known geometry, this approach then isolates the antenna delay of one node per ranging session. To estimate the complete network, a master node initiates several ranging sessions for each node utilizing a third node for support. Collected measurements of the participating nodes are then forwarded to the master node. The estimate of each antenna delay is computed externally as the average over the executions. Antenna delays are determined sequentially and individually for each node, hence not with respect to the complete network.

Contemporary works facilitate the self-localization of anchor networks without knowledge of the entire geometry, i.e., nodes infer relative positions based on inter-device measurements. Piavanini et al. employ antenna delay estimation as the first step before the self-localization of anchors. However, the nonlinearity in TWR measurements with antenna delays further resides in this work on self-localization [17]. Based on known tag positions, Shah et al. estimate unknown anchor positions and their antenna delays using non-linear Linear Squares with Taylor Series linearization [20]. Corbalán et al. focus on self-calibration of anchor positions with sparse connectivity [3]. While inference of positions is a significant step towards fully autonomous calibration, it introduces an inherent nonlinearity restraining on-device execution.

In contrast, ALADIn uses known geometry, i.e., the distances between nodes, to calibrate the antenna delays on the devices, enabling an autonomous self-calibration of antenna delays. Utilizing efficient all-to-all TWR ranging, nodes capture the exchange of all pairs at once, reducing air time occupancy. Our linear model reduces computational complexity and enables the aggregation of measurements for all node pairs, resulting in a calibration suited for resource-constrained devices. Hence, we create and solve an adaptive linear system on the devices. As nodes overhear all rangings, we execute the inference on each device separately, removing the need to coordinate the inference process using backbone communication or a central instance. By recording the standard deviation for each connection, we purify and filter out any inaccurate measurements that may have been recorded.

4 Design: Linear Antenna Delay Calibration

To enable autonomous calibration, we devise ALADIn, a calibration scheme that fully executes on resource-constrained devices, allowing continuous self-calibration without the need for external hardware. Performing the calibration protocol on the device further obtains the calibration values on the respective machine, even allowing the calibration of remote installations.

Yet, this self-calibration relies on the measurements gained by the ranging process between the UWB devices, and those measurements are inherently noisy [11]. Further, clock inconsistencies induce a relative drift that demands proper compensation, e.g., using a double-sided protocol [12]. While the Alt-DS-TWR protocol negates relative clock drifts in the ranging, in its original form, it introduces a nonlinearity between the antenna delays and the retrieved distance measurements [15]. This nonlinearity hinders efficient calibration and execution on constrained devices [8, 16, 17]. On the systems level, the devices' limited memory and floating-point accuracy present additional challenges. Further challenges reside in transmission schedules, finite packet sizes, and, ultimately, the constrained computational capabilities of those small devices.

In this section, we devise our linear model of the rang-

ing process with respect to individual antenna delays and hence reformulate the non-linear Alternative-Double-Sided Two-Way Ranging equation [15]. Based on the known geometry of nodes, we show that this linear connection permits efficient inference even on resource-constrained nodes. As Section 5 presents, we address the measurement process with a round-robin, all-to-all-ranging protocol. This protocol efficiently aggregates measurements for all node pairs and allows exclusion based on the sampled standard deviation. After the ranging, each sensor device derives a linear system for calibration and solves it separately using linear least squares. Calibration values are thus available on the spot and allow direct integration on the devices.

We first derive the linear ranging model, mitigating clock drifts concerning antenna delays. We then combine the individual measurements into a common linear system that allows efficient solving using ordinary least squares. A brief discussion of our model's limitations finalizes this section.

4.1 Linear TWR Ranging Model

Fundamentally, the measured durations of the Two-Way Ranging and antenna delays employ a direct linear dependency. We model one of those exchanges to capture this relation in the following. We assume that nodes are static, which generally holds for networks of anchor devices. Without loss of generality, let A initiate a TWR process with B.

4.1.1 Ranging

For the single-sided ranging of two distinct participants A and B, we denote the actual round times as R_A and the response delay D_B (cf. Figure 1). Let $T_{A\leftrightarrow B}$ denote the symmetric time of flight (ToF) between participants A and B. Moreover, we assume that messages are timestamped before their transmission; a downstream message could deliver the missing timestamps otherwise. The ToF relates to the actual round and delay times as follows:

$$2T_{A\leftrightarrow B} = R_A - D_B \tag{1}$$

4.1.2 Clock Drifts

Clocks tick with different frequencies due to clock imperfections are likely to drift apart while executing the ranging protocol. Hence, the recorded durations deviate from the actual ones. We denote the clock drift factor for a node A as k_A. As those drift factors mainly depend on the hardware and external factors, such as the temperature, we assume constant drift factors in our model, respectively for the short duration of the exchange. We denote X^Y as the skewed value of X according to the clock of Y, i.e.:

$$X^Y := k_Y X$$

4.1.3 Antenna Delays

Measured durations are further skewed by the delays in signal processing [6]. Let C_A denote the hardware-specific, combined antenna delay for node A, i.e., the sum of the delays for both the transmission and the reception of packets. As a result, recorded durations for the round times are increased while measured response delays are reduced by this combined delay. We assume that delays are constant in our model as they change only slightly with external factors such as temperature. Devices compensate range biases resulting from different signal strengths at reception using known values and are not considered in this model.

4.1.4 Reception Noise

We further assume timestamps of received packets are shifted by independent, zero mean additive Gaussian noise with variance σ^2_{RX} for the reception timestamps:

$$\varepsilon_{R_A}, \varepsilon_{D_B} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2_{RX})$$

We mitigate the non-linear range bias based on the received signal strength [6].

4.1.5 Measurement Model

We define $\hat{X}^Y$ as the measured duration or computed value of X as seen by node Y [15]. Hence, $\hat{R}^A_A$ describes the raw duration of R_A as measured on node A. Combining the delay, clock drifts, and reception noises, we assume the following relations for line-of-sight measurements:

$$\hat{R}^A_A = k_A(R_A + C_A + \varepsilon_{R_A}) = R^A_A + C^A_A + \varepsilon^A_{R_A} \tag{2}$$

$$\hat{D}^B_B = k_B(D_B - C_B + \varepsilon_{D_B}) = D^B_B - C^B_B + \varepsilon^B_{D_B} \tag{3}$$

4.1.6 Relative Clock Drift Mitigation

We additionally assume that relative drift factors, denoted as $\frac{k_A}{k_B}$ for nodes A, B, are known. This factor can be estimated based on the Carrier Frequency Offset (CFO) on message reception [7]. As an alternative, the symmetry within DS-TWR for static devices, i.e., $R_A + D_A = D_B + R_B$, (see Figure 1) can be exploited, requiring an additional message but no CFO estimation [15]:

$$\frac{k_A}{k_B} \approx \frac{\hat{R}^A_A + \hat{D}^A_A}{\hat{R}^B_B + \hat{D}^B_B} \tag{4}$$

This only holds approximately due to the inherent noise in the reception timestamps. Note that the aggregated delay shifts the round duration in the positive and the response delay in a negative direction, thus canceling out in (4).

With knowledge of the relative drift factor, node A converts the measured value $\hat{D}^B_B$ to $\hat{D}^A_B$, mitigating the effect of the relative clock drift:

$$\hat{D}^A_B = \frac{k_A}{k_B}\hat{D}^B_B \approx \frac{\hat{R}^A_A + \hat{D}^A_A}{\hat{R}^B_B + \hat{D}^B_B}\hat{D}^B_B \tag{5}$$

Note that this results in an alternative formulation of the Alt-DS-TWR protocol.

4.1.7 Linear Model

Combining (2), (3) and (1), we relate the actual ToF and antenna delays to the measured values (including mitigation of relative clock drifts) as seen by node A:

$$2T^A_{A\leftrightarrow B} = R^A_A - D^A_B = \hat{R}^A_A - C^A_A - \varepsilon^A_{R_A} - (\hat{D}^A_B + C^A_B - \varepsilon^A_{D_B})$$

Hence, we obtain a linear dependency between the measured values, the ToF and the combined delays of the antenna values:

$$\hat{R}^A_A - \hat{D}^A_B = 2T^A_{A\leftrightarrow B} + C^A_A + C^A_B + \varepsilon^A_{R_A} - \varepsilon^A_{D_B} \tag{6}$$

4.1.8 Repeated Measurements

The influence of reception noise is typically reduced by repeating the TWR process. Due to the linear relationship, we can average the measurements without affecting the calibration. Let $\hat{T}_{A\leftrightarrow B}$ denote the raw ToF estimate between the nodes A and B averaged over k repetitions, i.e., for $k=1$ calculated as: $\hat{T}^A_{A\leftrightarrow B} = 0.5\hat{R}^A_A - 0.5\hat{D}^A_B$

For arbitrary k, we get the following relation:

$$2\hat{T}^A_{A\leftrightarrow B} = 2T^A_{A\leftrightarrow B} + C^A_A + C^A_B + \varepsilon^A_{T_{A\leftrightarrow B}} \tag{7}$$

With $\varepsilon_{T_{A\leftrightarrow B}}$ being the expected noise, for which holds:

$$\varepsilon_{T_{A\leftrightarrow B}} \overset{\text{iid}}{\sim} \mathcal{N}(0, 2\sigma^2_{RX}k^{-1})$$

4.2 Deriving the Linear System

Our linear TWR ranging model allows relating the antenna delays of uncalibrated devices, relying on available average measurements and the anchor geometry. We now derive a complete linear system based on the individual equations.

Let n with $n \geq 3$ denote the number of deployed static anchor nodes with unknown antenna delays and let $\mathcal{A}$ represent the index set $\mathcal{A} := \{0,...,n-1\}$. We assume that distances between the anchor nodes under calibration are known, i.e., $T_{I\leftrightarrow J}$ is known for all pairs $I,J \in \mathcal{A} : I \neq J$. The goal now is to find a vector of antenna delays $\hat{\mathcal{C}} = (\hat{C}_0, \hat{C}_1, ..., \hat{C}_{n-1})^T$ which minimize the squared errors between the average distance measurements skewed by the antenna delays and the actual distances for all pairs, i.e., antenna delays which capture this difference:

$$\hat{\mathcal{C}}^M = \underset{(\hat{C}^M_0,...,\hat{C}^M_{n-1})^T}{\arg\min} \sum_{\substack{I,J\in\mathcal{A}\\ I<J}} (\hat{C}^M_I + \hat{C}^M_J + 2T^M_{I\leftrightarrow J} - 2\hat{T}^M_{I\leftrightarrow J})^2 \tag{8}$$

Using the average estimated ToF for each link, we derive a linear system of equations using (7). This system consists of $m := n(n-1)/2$ equations, one for each device pair. It features the antenna delays of all devices as unknowns $C_0,...,C_{n-1}$. This system is linear and overdetermined, so we use ordinary least squares to solve it. We, therefore, define a measurement matrix $Y^{m\times 1}$ using the difference between the average measured ToFs and the true ToFs. The design matrix $X^{m\times n}$ contains the linear influence of the individual delays for every pair:

$$Y^{m\times 1} := \begin{pmatrix} 2\hat{T}^M_{1\leftrightarrow 0} - 2T^M_{1\leftrightarrow 0} \\ 2\hat{T}^M_{2\leftrightarrow 0} - 2T^M_{2\leftrightarrow 0} \\ 2\hat{T}^M_{2\leftrightarrow 1} - 2T^M_{2\leftrightarrow 1} \\ \vdots \end{pmatrix}, X^{m\times n} := \begin{pmatrix} 1 & 1 & 0 & \\ 0 & 1 & 1 & \cdots \\ 1 & 0 & 1 & \\ & \vdots & & \end{pmatrix} \tag{9}$$

Hence, based on the averages of the measurements, we directly retrieve the estimation for the antenna calibration as the result of the following matrix multiplications:

$$\hat{\mathcal{C}}^M = (X^T X)^{-1} X^T Y \tag{10}$$

With our assumption of normally distributed noise, this ordinary least squares estimate further corresponds to the maximum-likelihood estimation for the antenna delays.

4.3 Limitations of the Linear Model

As introduced in this paper, our linear system assumes a direct dependency between the antenna delays, inter-distances, and noisy, recorded timestamps. Consequently, this simplification excludes known (non-linear) influences like the range bias as the result of differences in received signal strengths [6, 14]. In addition, the model misses the orientation of devices as another major factor of ranging errors [11]. We correct for the range bias based on the RSSI on reception using available correction values [5] but do not consider the errors caused by orientation. Instead, our linear approach prioritizes low resource requirements for on-device execution.

Moreover, this work does not model non-line-of-sight conditions [9] or not fully-connected scenarios [3]. Yet, the calibration procedure remains applicable to every fully connected group of nodes within anchor networks.

5 System Design

As introduced in the previous section, linearly optimizing the hardware-specific antenna delays enables us to move the calibration onto the sensor devices. This section covers the technical details necessary for on-device antenna delay calibration in ALADIn.

We optimize the antenna delays based on the known geometry and collected measurements between the anchor devices. The optimization procedure executes as follows:

1. Execution of an all-to-all-ranging protocol
2. Measuring and mitigating relative clock drifts
3. Aggregation of measurements
4. Efficient inference of the estimated antenna delays using linear least squares

This optimization incorporates the respective clock drifts and relates the antenna delay, inter-distances, and recorded noisy timestamps. Each device solves this optimization independently based on the exchanged measurements and known geometry, resulting in readily-available estimations of the antenna delays on each device.

5.1 All-To-All Ranging

We leverage many-to-many ranging [10, 21, 22] and execute consecutive ranging rounds. Instead of using the channel-impulse-response [2] or clock synchronization, we define a simple order for transmission and let nodes transmit in sequence just after the reception of their predecessors. Hence, based on assigned numbers, throughout every round, one after the other, every node sends a UWB packet to the network. This packet contains the transmission timestamp and all reception timestamps by other nodes received since its last transmission, cf. Figure 2. This allows all-to-all ranging between at least 20 nodes, scaling further with, e.g., a proprietary packet length of 1023 Bytes of the DWM1001 platform.

Consequently, we utilize the ability of Alt-DS-TWR to cope with long response delays for simultaneous rangings between all participating devices [21]. Devices execute rangings to all other devices and simultaneously overhear the ranging process between all pairs. Yet, for best accuracy, we limit the response to the same round, i.e., devices with a

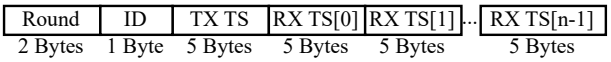

Figure 2: Packet Structure of ALADIn's All-To-All Ranging: Nodes broadcast a single ranging packet in each round. Nodes listen and record reception timestamps, estimating the measured distance for all pairs in the network.

lower number initiate the ranging process, and devices with higher numbers respond in the same round. Hence, each device collects $n(n-1)/2$ measurements per round.

5.2 Relative Clock Drift Mitigation

Fundamentally, we employ a Double-Sided TWR protocol. Similar to the Alt-DS-TWR protocol, we use the previous round to estimate the relative drift. Nodes then locally compute the measured ToF (which is skewed by antenna delays) for all pairs using (5). For example, node M that records the exchange between nodes A and B approximates $\hat{T}^M_{A\leftrightarrow B}$ for a single exchange as follows:

$$\hat{T}^M_{A\leftrightarrow B} = 0.5\hat{R}^M_A - 0.5\hat{D}^M_B = 0.5\frac{k_M}{k_A}\hat{R}^A_A - 0.5\frac{k_M}{k_B}\hat{D}^B_B \quad (11)$$

The relative drift factors k_M/k_A and k_M/k_B are approximated using (4) based on the timestamps of the current and the previous round. We assume constant relative drift for each round but not the overall ranging procedure; note that this assumption relies on the short execution time of the ranging exchange (cf. [6]).

As we are limited in the precision of native floating point numbers, particular care is required to handle the clock drift mitigation. Hence, instead of following (11) directly, devices calculate the offset of the relative drift, i.e., device M approximates this offset to device A as follows:

$$\frac{k_M}{k_A} - 1 \approx \frac{\hat{R}^M_M + \hat{D}^M_M - (\hat{R}^A_A + \hat{D}^A_A)}{\hat{R}^A_A + \hat{D}^A_A} \quad (12)$$

Based on (12), M computes and applies respective correction for the round and delay durations.

5.3 Aggregation of Measurements

For $k+1$ rounds, the all-to-all ranging protocol results in $kn(n-1)/2$ measurements. However, the average of the measurements is a sufficient statistic based on the assumed linear model in (7). Thus, devices can efficiently store and handle numerous measurements to reduce the effect of the inherent reception noise of single measurements, i.e., devices average measurements for each pair over multiple rounds. Hence, for each of the $n(n-1)/2$ estimates, we reduce storage requirements and aggregate the mean and variance using Welford's algorithm [23], allowing us to compute the mean value and variance for the downstream inference process.

5.4 Variance-Based Exclusion

While computationally simple, the linear model assumes independent noise with equal variance over the pairs. However, this assumption is challenged by multipath reflections and other nonlinearities present in practical scenarios. With the aggregation of the mean value, devices further capture the variance of each link. We utilize this variance and let devices exclude potential outliers. The evaluation, as covered by Section 6, analyzes the efficacy of such exclusions.

5.5 Efficient On-Device Inference

Every node overhears and collects the measurements for all device pairs. A timeout ensures that all devices eventually execute the inference after the last successful ranging round. Using the ordinary least squares algorithm from Section 4, we execute the inference process on every device without further communication. Let M denote the inferring device. As the actual distances are known, M determines the matrices Y and X based on (9) using the recorded average measurements. Thus, M computes its local estimate of all antenna delays $\hat{C}^M$ using (10).

As a further optimization, the matrices involved in the inference process allow precomputation, reducing the estimation to a matrix-vector multiplication of an $n \times m$ matrix and an m-dimensional vector.

5.6 Implementation

Built on top of Zephyr RTOS [2], we provide our open-source implementation to the community. The implementation is adjusted for the DWM1001 platform by Qorvo and adds precise ranging capabilities. Being equipped with an nRF52832, the DWM1001 features a 64 MHz Arm Cortex-M4 CPU with FPU support and 64 kByte of memory. Distances are handled in ToF units of the DW1000 UWB module and transmitted with 40 bit precision. Due to its 64 GHz clock frequency, one unit relates to 15.65 ps, a ToF equivalent to 0.47 cm. Devices mitigate the range bias on reception based on the reported received signal strength using provided correction values [5].

After startup, devices initiate the all-to-all ranging protocol and execute a fixed amount of rounds. Initially, nodes are tuned to the default antenna delay value [6], i.e., the calibration only results in a relative change of delays. Each device aggregates the resulting measurement per device pair. Values are stored in memory, and the linear system is generated dynamically. CMSIS-DSP accelerates the inference's matrix multiplications employing the available onboard DSP hardware.

Started by the first device, a single ranging round for 14 nodes completes in approximately 35 ms, i.e., 2.5 ms per device. The ranging rounds can be executed end-to-end, allowing the execution of hundreds of rounds per minute. For 14 nodes, the estimation itself executes in about 205 ms but requires 37.75 kBytes of memory with 16-bit floating point precision as we adaptively derive the linear systems on the device. Caching the matrix reduces the computation to a single matrix-vector operation with a matrix of size $n \times m$, lowering the memory footprint to several kBytes.

6 Evaluation

We start the evaluation with a numerical, simulation-based analysis of the proposed protocol. This first part compares ALADIn to related work. The second part then analyzes the performance of the proposed protocol under real-life conditions in testbeds. Afterward, we study the ef-

[2] https://zephyrproject.org/

Table 1: Resources Requirements (n nodes, k repetitions, *holds with respect to packet length)

Approach	Records	Messages	Remarks
TDoA [18]	$\mathcal{O}(1)$	$\mathcal{O}(nk)$	Nodes calibrated individually
GN [16]	$\mathcal{O}(n^2+k)$	$\mathcal{O}(n^2k)$	Values required for each pair
ALADIn	$\mathcal{O}(n^2)$	$\mathcal{O}(nk)$*	All-to-all ranging, values aggregated for each pair, joint calibration.

Table 2: Simulation Parameters

Parameter	Value
Number of devices	8
RX noise SD	1.0 ns
RX/TX Antenna Delay Mean	0.516 ns
RX/TX Antenna Delay SD	0.06 ns
Clock Drift SD	10 ppm
Monte-Carlo Runs	100

fect of excluding noisy measurement pairs and finish with a brief discussion. Please note that values displayed in square brackets denote the standard deviation.

6.1 Numerical Results

We first compare ALADIn to the TDoA approach by Shah et al. [18] and the non-linear Gauss-Newton protocol by Piavanini et al. [16]. The TDoA approach calibrates devices individually using a TWR exchange between two devices and TDoA measurements by a third device. Proposed solutions for TWR ranging extract combined antenna delays using non-linear equations for each pair using Particle-Swarm-Optimization [8] or the Gauss-Newton algorithm [16, 17] and combine them afterward using a linear system. Our proposed linear model induces a system that allows efficient aggregation and direct inference using ordinary least squares.

Table 1 displays the expected resource requirements of all three approaches. We display the number of records and messages. Our proposed protocol executes ranging with all devices and estimates delays jointly and on the devices. This results in higher memory consumption, although measurements get aggregated. The employed all-to-all ranging greatly reduces the number of required messages (up to the available packet length).

6.1.1 Simulation Setup

Using Monte Carlo simulations, we analyze the expected quality of the estimation. For comparison, we simulate ranging rounds for 100 Monte-Carlo executions. We assume drifts to be constant for each run, independent, and zero-centered with a standard deviation of 10 ppm. The antenna delays for each device are split into transmission and reception delays, each modeled as a Gaussian random variable. Reception noise is modeled via zero-centered, independent Gaussian noise. We refer to Table 2 for the main parameters and to the original work by Piavanini et al. for further details like the anchor geometry [16]. We vary the number of simulated rounds, i.e., the number of exchanges for each pair of devices. The RMSE of all estimated antenna delays from the Gauss-Newton approach and our proposed solution is derived and averaged over all pairs for each run. For the TDoA and our proposed algorithm, calibration values are extracted regarding the clock drift of the first device. The GN optimization executes with 100 iterations. In this static scenario, all approaches identify and mitigate the influence of clock drift; the error results from the reception noise. In the same way, the exact positions and number of anchors have little impact on the calibration performance.

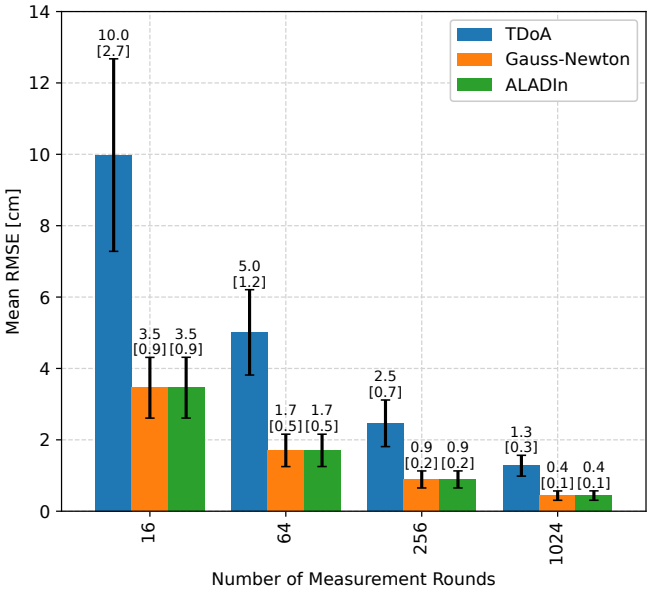

Figure 3: Simulation-based comparison of the proposed solution with a non-linear approach using Gauss-Newton, and a TDoA based: Two additional message receptions increase the amount of noise in the TDoA estimation. While computationally simpler than the Gauss-Newton, ALADIn offers comparable performance. Results are based on simulations as described in [16]. Error bars represent standard deviations. Table 2 contains the main parameters.

6.1.2 Simulation Results

The results of Figure 3 indicate that our proposed solution offers calibration performance comparable to the Gauss-Newton approach [16]. The average RMSE and its standard deviation stay comparable across the tested number of measurement rounds. The TDoA approach shows a little worse estimation performance, likely due to the reception noise of the incorporated two TDoA packets and individual calibration of nodes. Reducing the mean RMSE by half requires around four times the amount of repetitions: with $k = 16$ repetitions, GN and the proposed approach estimate with a mean RMSE of 3.4 cm, dropping to 0.4 cm at $k = 1024$ repetitions, just below the timestamp resolution of the DWM1001 platform.

Overall, the simulations indicate that the proposed solution is a viable alternative for delay estimation in terms of

accuracy. Our linear model, moreover, allows efficient aggregation and on-device inference, which the next section examines under practical conditions. Further, when we compare the GN and the proposed approach, the estimation in the node's drifted clock does not result in a noticeable shift in estimation quality as node drifts skew values in parts-per-million, hence a miniature effect. Moreover, even though the GN approach involves a non-linear formulation, observed gradients are constant, indicating a fundamentally linear dependency.

6.2 Testbed Results

While the simulations show the efficacy of ALADIn, the overall goal is to execute the estimation on actual devices. Hence, we deploy the protocol on three different public testbeds: We execute the calibration on two installations of DWM1001 devices with 7 and 14 nodes within the CLOVES Testbed in Trento [13] and on another 14 nodes residing in the FIT IoT-LAB in Lille [1]. As Figure 4 presents, all layouts allow LOS conditions between the devices. In Trento, the UWB nodes are attached to the ceiling, whereas in Lille, the nodes on the sides further reside on the walls. We rely on the specified geometry of the testbed for ground-truth distances between devices.

6.2.1 Testbed Setup

With the geometry embedded in the firmware, we deploy and run the ranging and calibration in all three scenarios on DWM1001 UWB nodes. We reuse the order of nodes and repeat the ranging $k = 1000$ times, resulting in $k+1$ ranging rounds. Devices transmit on channel 5 with a PRF of 64 MHz and a preamble length of 128 symbols. We initially run 200 warm-up rounds to limit the clock drift at device startup. Measurements are aggregated on each device separately and fed into the linear least squares algorithm with the known geometry. As presented in Section 4, all devices execute the inference individually and retrieve estimated antenna delays. Then, nodes calculate distances based on the mean measurements and the estimated delays. For comparison, nodes infer distances using the default antenna delay and the factory-calibrated delay values (as present in OTP memory). For evaluation purposes, values are logged using a serial backbone connection. If not noted otherwise, displayed values are extracted from the first device of each installation. Yet, due to the LOS conditions, nodes commonly receive transmissions from all other nodes. And, while the relative drift results in drifted measurements and hence skewed estimated delays, values are only skewed by parts of a million; the influence of relative drift is thus negligible. Note that nodes can derive values in terms of another clock.

6.2.2 Testbed Results

We compare the effect of the estimated delays to the default delays, i.e., uncalibrated and factory-calibrated delays, in terms of their change in the mean absolute error (MAE) of all pairs. Figure 6 displays the results: In all three scenarios, the MAE lowers with factory-calibrated delays and further decreases with the proposed calibration algorithm. Compared to the uncalibrated case, the MAE reduces with our proposed calibration by 3.8 cm in CLOVES-7, 2.8 cm in CLOVES-14, and 3.0 cm in IoT-LAB-14. As the MAE is less sensitive to extremes, it inherits a high standard deviation. Yet, the proposed calibration lowers the standard deviations in all scenarios.

Table 3: Estimated and Factory Antenna Delays Offsets for CLOVES-7: Estimated delay values differ from factory calibrated values.

Node	Estimated [cm]	Factory [cm]	Diff. [cm]
1	1.88	10.34	-8.46
2	20.68	4.23	16.45
3	9.40	10.34	-0.94
4	29.61	10.34	19.27
5	1.41	3.76	-2.35
6	-3.76	4.70	-8.46
7	10.81	4.70	6.11

Using CLOVES-7 as an example, we analyze the effect of the calibration on the individual pairs in detail. Table 3 lists the estimated and factory delays (offset by the default antenna delay): The estimated delays differ on average from their factory-calibrated values by 8.9 cm. Most notably, the approach adds 19.27 cm of delay to device 4. We display the effect of the estimated delays on the individual pairs in Figure 7. While the error of pairs mostly declines, a high offset in Pair 7-4 results in an overcompensation of node 4's delay, simultaneously impeding other rangings. As the individual measurements of pair 7-4 indicate, the pair suffers from high noise in the measurements, affecting the aggregated offset and hence the calibration. Likely the result of multipath components, Figure 8 presents the individual measurements of this noisy pair 7-4 and, for comparison, Pair 7-3 throughout the ranging process.

6.3 Exclusion of Noisy Links

The proposed calibration reduces ranging errors even in the presence of multipath components which are apparent in CLOVES-7 and IoT-LAB-14. Yet, our results indicate that the variance is similar for most pairs [16]; Figure 5 presents variances of the pairs aggregated on the devices using channel 5. Experimental runs with channel 2 indicate increased This section analyzes the potential of excluding measurements from the estimation process, boosting the calibration's robustness against prominent noise in measurements.

6.3.1 Filtered Setup

In addition to the unfiltered estimation, devices now exclude pairs with a standard deviation of 10 cm, based on the on-device aggregates. Empirically, those pairs indeed suffer from multipath reflections. Hence, we exclude Pair 7-4 for CLOVES-7 and pairs 5-3, 8-2, and 11-4 for IoT-LAB-14 in the error estimation.

6.3.2 Filtered Results

Figure 9 depicts the resulting MAE over all pairs. Because outliers are excluded, the MAE decreases in all cases. Yet, removing noisy measurements from the estimation indicates additional improvement in the resulting error and its standard deviation. The gain decreases with the number of

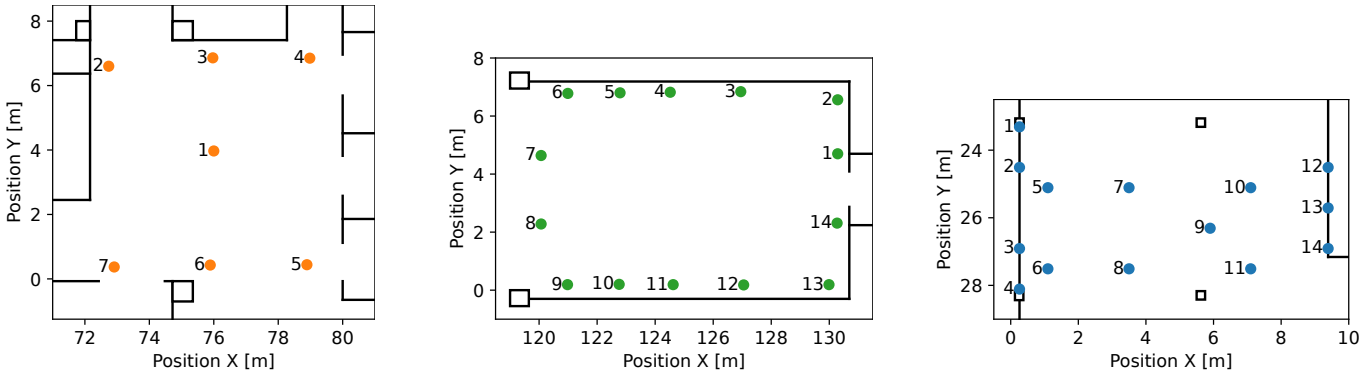

(a) CLOVES installation with 7 nodes (b) CLOVES installation with 14 nodes (c) FIT IoT-LAB with 14 nodes

Figure 4: The practical experiments run on three public UWB installations: Two installations within the CLOVES Testbed in Trento [13] and the FIT IoT-LAB in Lille [1] contain 7 and 14 DWM 1001 UWB devices. Their configuration allows for LOS transmissions between all nodes.

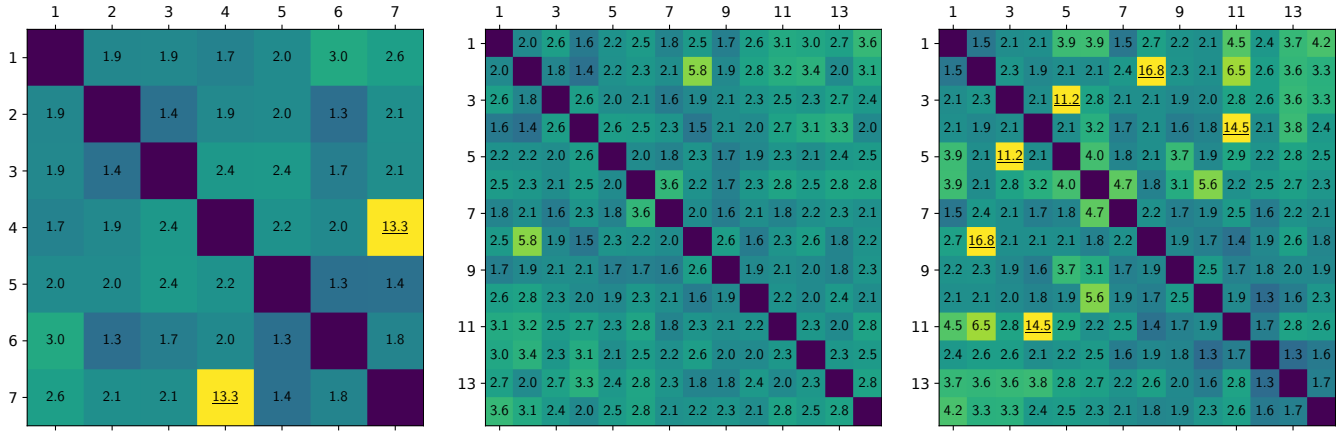

(a) CLOVES-7 (b) CLOVES-14 (c) IoT-LAB-14

Figure 5: The three scenarios exhibit different variances in the pairwise measurements: Displayed is the standard deviation aggregated using channel 5 on node 1 for each pair in cm. CLOVES-14 shows reduced measurement noise in comparison. Underlined values indicate empirical multipath propagation: Pair 7-4 in CLOVES-7 and pairs 5-3, 8-2, as well as 11-4 in IoT-LAB-14, empirically experience multipath reflections.

nodes involved in the calibration process as the effect of outliers is inherently reduced. Overall, excluding multipath-impacted links reduces the error in CLOVES-7 by 4.0 cm with a difference of 2.8 cm in its standard deviation. When we compare the raw estimated delay values to the factory delays, the filtration moves estimated delays closer to factory values, bringing the mean absolute deviation from 8.9 cm in the unfiltered down to 6.8 cm with enabled filtering. The impact on IoT-Lab-14, however, appears insignificant. This difference presumably originates from the increased number of pairs (i.e., 21 for CLOVES-7 and 91 for IoT-Lab-14), reducing outliers' influences on the calibration process.

Instead of removing noisy pairs, the sample variance of pairwise measurements encourages dynamic weighting, i.e., weighted least squares. As this only requires scaling the individual equations, the overall inference stays linear and executable on the device.

6.4 Discussion

ALADIn minimizes the squared errors between actual distances and UWB measurements, designed to fit linear antenna delays. Yet, results notably differ between simulations and practical deployments: Non-linear influences or NLOS conditions presumably result in inherent errors not handled by the approach. While the available variance allows devices to filter out noisy links, the proposed algorithm retains the potential of overfitting to the specific layout of installations, like relative orientations and distances of devices, presumably resulting in the recorded difference between estimated and factory delays. Consequently, in the best case, the calibration process includes numerous nodes and, further, nodes in characteristic positions, i.e., areas of nodes that should be localized. However, our evaluation results indicate that

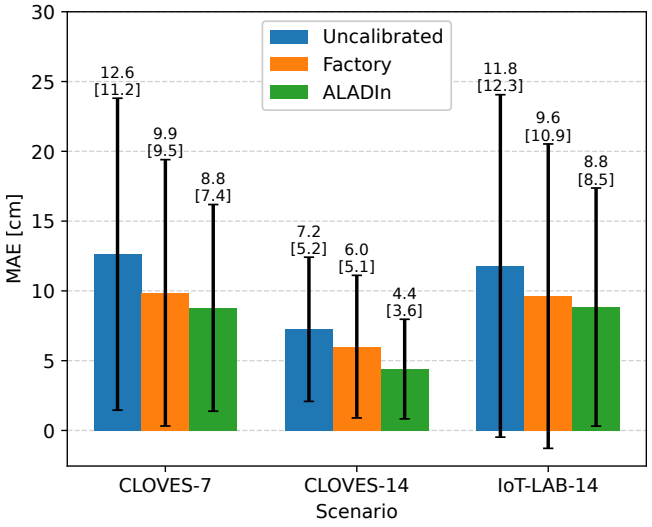

Figure 6: Calibration Effect on the MAE of Pairs Without Filtering: The proposed linear calibration reduces ranging errors between all pairs in all tested installations. Devices carry out calibration fully autonomously. Our calibration also decreases the standard deviation.

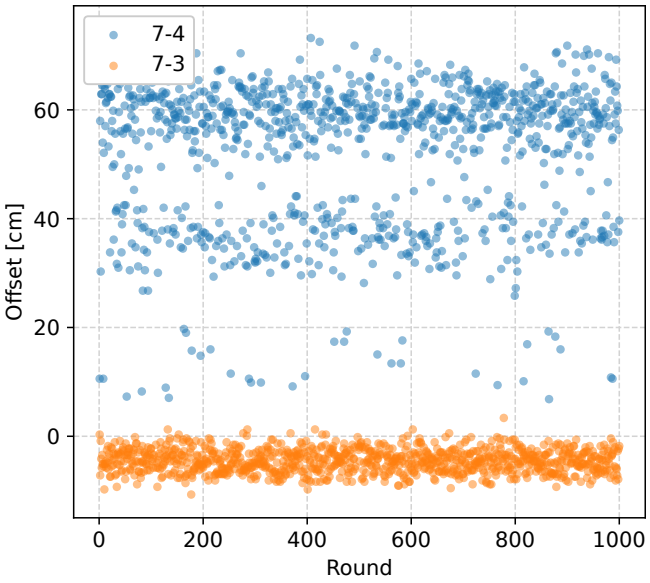

Figure 8: Scatter Graph of Individual Measurements in CLOVES-7: Throughout the ranging rounds, the pair 7-3 shows expected reception noise, whereas the pair 7-4 presumably suffers from multipath propagation, resulting in a higher standard deviation and skewed mean offset (cf. Figure 5).

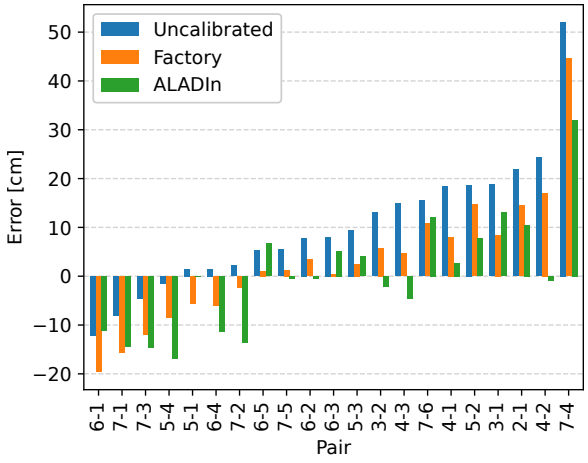

Figure 7: Calibration Effect on the Distance Estimation Error of Individual Pairs: The model excludes nonlinearities and multipath propagation. Noisy pairs such as 7-4 can result in overcompensation. Pairs are ordered by their uncalibrated error. The linear approach is unable to mitigate all perceived ranging inaccuracies.

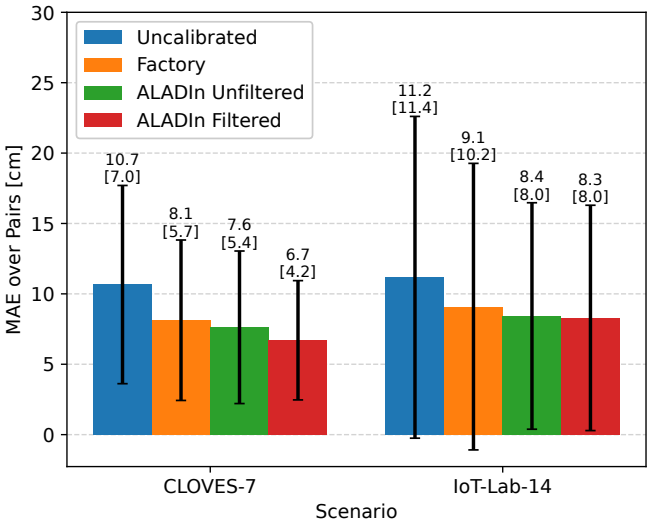

Figure 9: Calibration Results Excluding High Variance Pairs: For a more robust calibration, devices exclude high variance pairs (SD of at least 10 cm, presumably suffering from multipath effects) from the error calculation. Excluding noisy connections from the calibration further enhances calibration accuracy but experiences diminishing returns with the number of devices.

even this simple linear model reduces ranging errors below the factory calibration; without any need for external computation or hardware. Energy usage mostly depends on the ranging procedure, as the inference on the devices executes within a second.

7 Conclusion

UWB enables precise indoor localization on IoT devices. Yet, calibration of antenna delays remains an essential step for accurate estimations. Enabling autonomous calibration, this work presents ALADIn, an approach for antenna delay calibration on resource-constrained UWB devices. As we derive the linear least squares approach from a linear dependency of antenna delays, we lower computational complexity and move the calibration to the devices, reducing manual labor and the need for external hardware and computation. The combination with the all-to-all ranging procedure allows efficient aggregation of TWR measurements between all devices, decreasing calibration time to minutes. In simulations, the approach behaves similarly to existing approaches while lowering computational complexity. Remotely deployed and evaluated on three public UWB testbeds with once 7 and twice 14 DWM1001 nodes, the open-source implementation decreases ranging errors and standard deviation, reducing the MAE by up to 4.0 cm in the CLOVES deployment with 7 devices based on known geometry. Further, our results show that excluding pairs based on their variance results in a more robust calibration procedure. While ALADIn is inherently limited to linear factors and LOS conditions, our results indicate the potential to reduce ranging errors even in already deployed UWB ranging networks.

8 References

[1] C. Adjih, E. Baccelli, E. Fleury, G. Harter, N. Mitton, T. Noel, R. Pissard-Gibollet, F. Saint-Marcel, G. Schreiner, J. Vandaele, et al. Fit iot-lab: A large scale open experimental iot testbed. In *2015 IEEE 2nd World Forum on Internet of Things (WF-IoT)*, pages 459–464. IEEE, 2015.

[2] P. Corbalán and G. P. Picco. Concurrent ranging in ultra-wideband radios: Experimental evidence, challenges, and opportunities. In *EWSN*, pages 55–66, 2018.

[3] P. Corbalán, G. P. Picco, M. Coors, and V. Jain. Self-localization of ultra-wideband anchors: From theory to practice. *IEEE Access*, 2023.

[4] A. De Preter, J. Anthonis, J. Swevers, and G. Pipeleers. Experiment design for ultra-wideband sensor node calibration. In *Proceedings of the 2018 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2018.

[5] Decawave. Aps011 application note: Sources of error in dw1000 based two-way ranging (twr) schemes, 2014.

[6] DecaWave. Aps014 application note: Antenna delay calibration dw1000-based products and systems, version 1.2, 2018.

[7] I. Dotlic, A. Connell, and M. McLaughlin. Ranging methods utilizing carrier frequency offset estimation. In *2018 15th Workshop on Positioning, Navigation and Communications (WPNC)*, pages 1–6, 2018.

[8] X. Gui, S. Guo, Q. Chen, and L. Han. A new calibration method of uwb antenna delay based on the ads-twr. In *2018 37th Chinese Control Conference (CCC)*, pages 7364–7369. IEEE, 2018.

[9] I. Guvenc, C.-C. Chong, and F. Watanabe. Nlos identification and mitigation for uwb localization systems. In *2007 IEEE Wireless Communications and Networking Conference*, pages 1571–1576. IEEE, 2007.

[10] IEEE. Ieee standard for low-rate wireless networks–amendment 1: Enhanced ultra wideband (uwb) physical layers (phys) and associated ranging techniques. *IEEE Std 802.15.4z-2020 (Amendment to IEEE Std 802.15.4-2020)*, pages 1–174, 2020.

[11] A. Ledergerber and R. D'Andrea. Ultra-wideband range measurement model with gaussian processes. In *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pages 1929–1934. IEEE, 2017.

[12] C. Lian Sang, M. Adams, T. Hörmann, M. Hesse, M. Porrmann, and U. Rückert. Numerical and experimental evaluation of error estimation for two-way ranging methods. *Sensors*, 19(3):616, 2019.

[13] D. Molteni, G. P. Picco, M. Trobinger, and D. Vecchia. Cloves: A large-scale ultra-wideband testbed. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, pages 808–809, 2022.

[14] S. Monica and G. Ferrari. Improving uwb-based localization in iot scenarios with statistical models of distance error. *Sensors*, 18(5):1592, 2018.

[15] D. Neirynck, E. Luk, and M. McLaughlin. An alternative double-sided two-way ranging method. In *2016 13th workshop on positioning, navigation and communications (WPNC)*, pages 1–4. IEEE, 2016.

[16] M. Piavanini, L. Barbieri, M. Brambilla, M. Cerutti, S. Ercoli, A. Agili, and M. Nicoli. A calibration method for antenna delay estimation and anchor self-localization in uwb systems. In *2022 IEEE International Workshop on Metrology for Industry 4.0 & IoT (MetroInd4.0&IoT)*, pages 173–177. IEEE, 2022.

[17] M. Piavanini, L. Barbieri, M. Brambilla, M. Cerutti, S. Ercoli, A. Agili, and M. Nicoli. A self-calibrating localization solution for sport applications with uwb technology. *Sensors*, 22(23):9363, 2022.

[18] S. Shah, K. Chaiwong, L.-O. Kovavisaruch, K. Kaemarungsi, and T. Demeechai. Antenna delay calibration of uwb nodes. *IEEE Access*, 9:63294–63305, 2021.

[19] S. Shah and T. Demeechai. Multiple simultaneous ranging in ir-uwb networks. *Sensors*, 19(24):5415, 2019.

[20] S. Shah, L.-o. Kovavisaruch, K. Kaemarungsi, and T. Demeechai. Node calibration in uwb-based rtlss using multiple simultaneous ranging. *Sensors*, 22(3):864, 2022.

[21] F. Shan, J. Zeng, Z. Li, J. Luo, and W. Wu. Ultra-wideband swarm ranging. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*, pages 1–10. IEEE, 2021.

[22] M. Stocker, J. Kowalczyk, C. A. Boano, and K. Römer. Towards secure multicast ranging with ultra-wideband systems. In *Proceedings of the 2022 INTERNATIONAL CONFERENCE ON EMBEDDED WIRELESS SYSTEMS AND NETWORKS*, pages 256–261, 2022.

[23] B. P. Welford. Note on a method for calculating corrected sums of squares and products. *Technometrics*, 4(3):419–420, 1962.

[24] F. Zafari, A. Gkelias, and K. K. Leung. A survey of indoor localization systems and technologies. *IEEE Communications Surveys & Tutorials*, 21(3):2568–2599, 2019.