ISRiCH: Infrastructure-light Semi-passive Ranging for Collaborative Hybrid Localization

Benoît Perroux
IRIT, Université de Toulouse, CNRS, Toulouse INP, Alliantech
Toulouse, France
benoit.perroux@irit.fr

Jérôme Ermont
IRIT, Université de Toulouse, CNRS, Toulouse INP
Toulouse, France
jerome.ermont@irit.fr

Katia Jaffrès-Runser
IRIT, Université de Toulouse, CNRS, Toulouse INP
Toulouse, France
katia.jaffres-runser@irit.fr

***Abstract*—Wireless instrumentation has led to greater flexibility in the deployment of sensors in industry, from dense scenarios to temporary ones. Many of the new use cases, however, need knowledge of the sensor's positions to be able to process the acquired data and model the physical phenomenon of interest across the environment. Wireless localization solutions achieving centimeter-level accuracy leverage Ultra-wideband ranging technology. It requires rolling out at least 3 *anchors*, which are fixed wireless nodes whose positions are known. Recently introduced, Single-Anchor Localization (SAL) facilitates temporary deployments but at the cost of reduced accuracy. This paper presents ISRiCH, a SAL algorithm tailored for localizing a dense deployment of sensors with reduced overhead and improved accuracy. It relies on a semi-passive ranging mechanism overhearing additional data (for free) that is then accounted for in a multidimensional scaling (MDS) computation of sensor positions. Results show that accuracy is improved by a factor of 3 with 16 sensors or more compared to classical SAL, down to around 25 cm. Furthermore, ISRiCH can match with certainty sensors corresponding to known points of interest in 2 to 5 runs for most regular grid deployments, when classical SAL never can, while still being easily deployable for temporary instrumentation.**

***Index Terms*—Wireless sensor networks; Ultra-wideband; Localization; Single-anchor localization; Passive ranging**

I. Introduction

The recent developments in industrial instrumentation show a clear inclination to remove all non-essential wires to reduce weight and to simplify the sensors' deployment. This shift is driven by the growing need for flexibility in the measurement setups and the rise of low-power wireless communication technologies. As a result, the switch to wireless sensor networks has paved the way for new use cases, from temporary deployments to instrumentation of low accessibility environments.

Temporary deployments are particularly of use for the study of mechanical wave propagation, such as the acoustic or vibratory analysis of an environment. An example of these contexts is the modal analysis of an aircraft wing. In this setup, tens to hundreds of accelerometers are placed on the wing, often in a regular grid pattern, to record the deformation of the structure when put under stress. The use of wireless sensors avoids the risk for the wires to interact with the structure and thus impact the measurement quality. Similarly, the deployment of a distributed acoustic antenna in an aircraft cabin or an amphitheater implies using a large number of wireless microphones densely positioned (e.g. one per seat).

In both cases, the sensors' installation procedure is made easier and allows for more temporary installations. Once the network is deployed and the accelerometer or microphone data is acquired, it is necessary to know with good accuracy (around $20\,\mathrm{cm}$) where each sensor was placed to perform the physical analysis. Such analysis can range from the modeling of the mechanical deformation waves of the wing to applications such as source localization [1].

Additionally, mechanical wave propagation measurement often implies that sensors be placed at specific locations (e.g. on seats in an aircraft) or in a regular pattern such as a grid. In these cases, we know a priori where a sensor can be placed, and the challenge is to associate the actual sensor to one of the predefined positions among the tens or hundreds that could be in its vicinity. Localization and matching methods need to be able to resolve any uncertainty and avoid sensor permutations that could be damaging to the data reconstruction.

As the use of wireless sensors allows for more temporary deployments, the localization system must be easy to deploy in a limited amount of time, not relying on a heavy infrastructure, but also adapt to the high sensor density in these setups. This paper thus focuses on this question of providing a precise localization service for a dense and temporary deployment of wireless sensor nodes.

Localization algorithms for wireless sensor networks have been a hot topic in recent years. Main methods can be classified into three categories sharing some common bases, such as the use of several fixed nodes known as *anchors* used as references to find the positions of possibly mobile *tags*. These three localization paradigms are fingerprinting [2], range-based [3], [4], and angular-based [5], [6]. While the range-based paradigm has become common with the development of the Ultra-wideband (UWB) radio technology, its rank as leader is being challenged by the angular-based methods [7], which can not only find the distance to an object but also the direction using antenna arrays. Both paradigms have led to the development of Single-Anchor Localization (SAL) [8], [9] that combines range and angular-based mechanisms to localize tags using only one fixed anchor. SAL is thus a good candidate to fit the reduced infrastructure needs of temporary deployments, though the accuracy of such methods is lower than more traditional multi-anchor range-based systems.

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

Our main contribution is to improve the performance of Single-Anchor Localization systems, especially for dense sensor deployments such as the one required by the modal analysis of a structure. To do so, we propose a new localization protocol, ISRiCH, that improves SAL using passive ranging. Our goal is to improve the localization accuracy when no prior assumptions are made about the environment, with the objective of limiting the number of messages exchanged between the anchor and the sensors. In dense grid deployments of $N = 50$ tags or more, ISRiCH offers a reduction of the localization RMSE by a factor of three compared to SAL, while requiring only $2N$ messages.

Tests on grid networks show ISRiCH's ability to provide a network-wide correct association more than 90% of the time on larger grids, with an accuracy increasing with the number of tags. For example, with 49 tags, correct association when tags are $40\,\mathrm{cm}$ apart is found in 95.3% of our tests with ISRiCH while SAL never succeeds. Furthermore, the repetition of ISRiCH runs, together with an election algorithm, can ensure a perfect match in 2 to 5 runs for grids of $30\,\mathrm{cm}$ or more with as few as 25 tags.

In the remainder of this paper, we first review existing localization techniques and algorithms in Section II. Section III then describes the ISRiCH localization method, before evaluating its accuracy in Section V, using the experimental setup described in Section IV. We propose an extension of the ISRiCH algorithm to the position matching use case in Section VI. Finally, Section VII concludes this work.

II. Related work

As wireless sensor networks (WSN) became more popular, RF-based localization has seen many developments in recent years. Current research directions and challenges include ultra-low power localization [10], lightweight algorithms for distributed localization systems, and reduced infrastructure dependency [11]. In that context, Ultra-wideband (UWB) has emerged among available radio technologies as the best compromise between localization accuracy and the cost and energy constraints induced by WSN deployments. In this section, we review the main UWB localization methods in light of these challenges, from data acquisition to processing.

A. Multi-anchor localization techniques

A localization technique describes *what* information is measured to estimate a node's position and *how* it is collected. Despite their diversity, UWB localization techniques share some common concepts: In every localization scheme, we have a given number of *tags*, which are mobile or static nodes whose positions need to be determined, and one or more *anchors*, fixed nodes acting as reference points. A message exchange between anchors and tags is used to collect the necessary data to estimate the positions of the tags relative to the anchors. As introduced earlier, UWB localization can leverage either range- or angular-based methods that are described next.

1) Range-based methods: The most common category of UWB-based localization methods relies on the measurement of the propagation time of a frame. This duration, also known as time of flight, is then used to compute the distance between the emitter and receiver, given the speed of light. This procedure is called *ranging* and is made possible thanks to the specific large bandwidth spectrum of the UWB technology. A wide spectrum in the frequency domain relates to a very short impulse in the time domain, enabling the precise timestamping of the emission and reception of a frame by a radio. Localization algorithms vary depending on whether they directly use the time of flight of a frame transmitted between the tag and the anchors, as in Time of Arrival (ToA), or leverage the delayed reception of a frame by synchronized anchors, as in Time Difference of Arrival (TDoA).

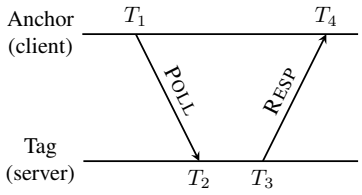


Fig. 1. Ranging exchange with TWR.

a) Time of Arrival (ToA): Here, the Two-Way Ranging (TWR) protocol is leveraged to derive the propagation time between anchor and tag as depicted in Fig. 1. Four timestamps are measured at emission (T_1 and T_3) or at reception (T_2 and T_4) of polling request and response messages. With T_2 and T_3 being carried in the response frame, the anchor (or *client* node) can derive its distance to the tag (or *server* node) from:

$$d = c.\frac{(T_4 - T_1) - (T_3 - T_2)}{2} \tag{1}$$

with c the speed of light. Note that anchor and tag can be set as *server* and *client* respectively if the tag needs to know its distance to the anchor. To localize a tag, one ranging procedure is required per fixed anchor.

The use of a two-way communication in TWR avoids the need for precise synchronization between tag and anchor, as the clock offsets are naturally compensated for in Eq. (1). Other ranging protocols exist that overcome limiting factors such as the relative clock drift between anchor and tag (SDS-TWR [4]), or the increased number of messages to perform the ranging with each anchor (N-TWR [3]). SDS-TWR adds 2 more messages in the tag-anchor ranging procedure to compensate for their relative clock drift, but at the cost of a reduced scalability and energy efficiency of the complete localization procedure. N-TWR factorizes ranging exchanges between a client node and several server nodes to reduce the total message count needed for localization. But relative clock drifts between client and server nodes can add imprecision since server nodes don't receive the frame at the same time.

Once the distance between a tag and several anchors is known with ranging, each anchor knows the distance d_i

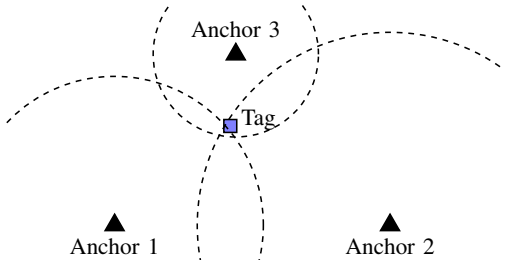


Fig. 2. Position estimation with TWR. Each dotted circle represents the estimated distance from an anchor to the tag with the tag at the intersection.

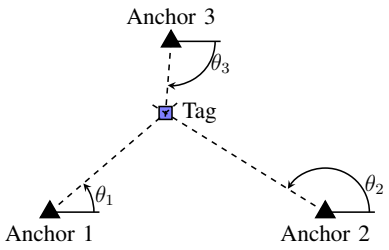


Fig. 3. Position estimation with AoA. Each dotted line represents the estimated direction from an anchor to the tag with the tag at the intersection

measured between itself and the tag and assumes that the tag lies on a circle (or a sphere in 3-dimensional (3D) localization) of radius $r = d_i$ around the anchor. With enough anchors (at least 3 in 2D and 4 for 3D), the intersection of the circles or spheres reveals the position of the tag, as illustrated in Fig. 2. Since ranging provides imprecise estimates, a perfect intersection of the circles often does not exist. An optimization method such as a least-squares resolution may then be used to get the most probable tag position.

With these methods, the tags thus have to be in range of at least 3 or 4 anchors, and preferably in direct line of sight to avoid drops in accuracy. In a permanent deployment as required by a tracking application in a warehouse for instance, leveraging multiple anchors for better localization is possible; however, for a temporary deployment of just a few hours or days, optimizing or simply over-provisioning the fixed infrastructure is economically prohibitive. Furthermore, although some solutions reduce the total number of messages [3], they do not scale to hundreds of nodes in a dense setup.

b) Time Difference of Arrival (TDoA): The TDoA localization technique differs from ToA by the nature of the distances that are measured. In this scheme, the tags emit only one message whose reception is timestamped by all the anchors. By comparing the timestamps received by a pair of anchors, the difference in distance between the tag and the two anchors can be derived [12]. Similarly to the ToA mechanisms, the combination of the distances from several anchor pairs reveals the position of the tag. This reduced number of messages, only one per tag, ensures a good scalability. However, contrary to TWR's bidirectional communication, the clock offsets are not naturally compensated: anchors need to be very precisely synchronized with a nanosecond accuracy to compare their timestamps. As such, the centimeter-level accuracy [13] and reduced number of messages come at the cost of a greater infrastructure dependency, making this technique impractical in temporary deployments.

2) Angular-based methods: These methods are built on the measurement of the direction of the tag as seen from an anchor. To do so, Angle of Arrival (AoA) anchors perform Phase Difference of Arrival (PDoA) [5] when receiving a signal from the tag. An anchor is equipped with an array of antennas that each measure the phase of a signal received from the tag. The direction of the tag is then extracted from the different phase values measured. Similarly to ranging, the position of the tag is found by combining measurements from several anchors, as shown in Fig. 3. AoA clearly simplifies the message exchange, with only one message emitted per tag, thus reducing energy consumption and network load. Furthermore, with PDoA-enabled devices becoming more easily available, more angular-based solutions have been proposed in recent years, reducing the accuracy gap between TWR and AoA [5].

However, several factors can impact the performance of AoA methods. First, the angular estimation tends to have a better accuracy at low angles, with higher uncertainty when the tag deviates from the orthogonal to the antenna plane [8]. The distance between anchor and tag is also of importance: a given angular error converts to greater positioning errors the further the tag is. Angular-based localization also comes at the cost of greater hardware complexity on the anchor side compared to range-based. Due to the angular measurements, the number of antennas, their nature, and the geometry of the antenna array have a major impact on the performance of the localization system [6], especially when working in 3D.

The ULoc method [7] is one such 3D-AoA scheme that, by using a custom-made 8-antenna array, results in an accuracy of $3.6\,\mathrm{cm}$ for a stationary tag. Nonetheless, the complexity of the anchor design, which requires strict synchronization between the 8 antennas, combined with the need for deploying several of them, hinders the widespread adoption of this system. Authors in [14] focus on evaluating and improving the hardware design of the anchors, using Qorvo DW3220 IC with only two antennas. Even though the localization is only 2D, they propose to use machine learning to reduce ambiguities such as the symmetry with regards to the antenna plane, expanding the capabilities of lower-cost 2-antenna anchors.

Similarly to TDoA methods, the angular-based paradigm reduces the total number of messages. The need for rolling out several anchors is still present, even if there is no need for a specific inter-anchor synchronization.

3) Performance of multi-anchor localization: With both angular- and ranged-based methods, the tag's position is computed by combining the measurements from multiple anchors. Because of the measurement uncertainties, the more anchors we have, the better the position estimation is after processing [8], reaching up to $5\,\mathrm{cm}$ for ranged-based localization [15] and $3.6\,\mathrm{cm}$ for angular-based localization [7] in optimal setups. As such, multi-anchor localization can solve the positioning of sensor on a dense grid of 20 or $30\,\mathrm{cm}$. Reaching this

performance however comes at the price of a need for *at least* 3 or 4 anchors, for respectively 2D or 3D localization. In fixed setups, this limitation is of little impact as the anchor infrastructure stays in place between uses. However, temporary deployments would benefit greatly from rolling out the least possible number of anchors. Next, we present the principle and performance of Single-Anchor Localization, compared to multi-anchor solutions.

B. Single-anchor localization

As the need for a multi-anchor infrastructure is a limiting factor in UWB localization [11], current research aims at designing infrastructure-light solutions. In Single-Anchor Localization (SAL), both ranging and angular information are leveraged at the same anchor to create an all-in-one system, also known as *hybrid localization*. SAL reduces deployment cost, opening the possibility for temporary or mobile localization devices. With both angular and ranging measurement, the position estimation is done by converting from relative polar coordinates to Cartesian using trigonometric functions.

The CELIDON [9] method uses the anchor mobility to provide a relative localization service. Anchors are mounted on the helmets of first responders in low visibility situations. Even if the 3D angular uncertainty is around $20°$ and depends on the bearer's orientation, its portability makes it a good basis for future relative ad-hoc localization systems. In ReLoki [6], authors use different types of antenna arrays to compare their performance when used jointly with TWR to compute the position of a tag. They show that for both 3D antenna arrays they tested, the final localization error reaches below $50\,\mathrm{cm}$ for low angle measurements.

So far, SAL trades reduced infrastructure costs for a reduced localization accuracy. In [8], the authors compare the performance of a SAL system with multi-anchor systems based on AoA and TWR in a 2D setup. Results show that SAL, with a Mean Absolute Error (MAE) of $48.4\,\mathrm{cm}$, performs better than a 4-anchor AoA-only scheme (without ranging) and is comparable with a 2-anchor TWR-only scheme (without angular measurements). Adding more anchors to the range-based system really outperforms SAL, with a MAE of $21.4\,\mathrm{cm}$ for 3 anchors and $18.7\,\mathrm{cm}$ for 4 anchors. Clearly, there is still room for improving accuracy in SAL solutions to become competitive with multi-anchor systems.

C. Collaborative localization

Collaborative localization expands on the more common anchor-tag scheme to collect and process information about the whole network rather than treating each tag independently. These collaborative methods can be used to reduce the dependence on an infrastructure [16], [17], increase the accuracy [18], [19], save energy [20], or limit the number of messages needed for the full network localization [19], [16]. In this section, we present different collaborative localization solutions, focused either on data collection with passive ranging, or on data processing with multidimensional scaling.

1) Multidimensional scaling: Multidimensional scaling, or MDS, is a data processing method that takes as input a *dissimilarity* metric between a set of elements in order to position them in a 2D or 3D space. Elements with low (resp. large) dissimilarity are positioned close to (resp. far from) one another, outputting a representation in the chosen number of dimensions. This is done through the use of the iterative SMACOF (Scaling by MAjorizing a COmplicated Function) algorithm [21] that minimizes a stress function.

MDS has been widely used in distributed localization solutions [18]. In this context, nodes perform pairwise ranging operations with each individual neighbor, or at least a large subset of them. The resulting distances are then used as the dissimilarity metric input for multidimensional scaling. In [22], authors propose a generalization of the classical MDS approach by adding fixed points of known position to the system: the anchors. Authors in [17] apply the same fixed points generalization as the previous work to perform the self-localization of a group of anchors where only a small subset originally know their position, thus expanding the infrastructure for the localization of tags. A key factor in this work is the use of *weighted* MDS to allow for out-of-range nodes: an unknown ranging is represented as a zero-weight to nullify its impact. While MDS can improve the accuracy of the localization system, the acquisition of the necessary data can be a limiting factor since for a network of N nodes, $\binom{N}{2}$ pairs must perform ranging.

2) Passive ranging: In passive ranging, a node aims at gaining positioning data by *overhearing* existing messages. As the total message count is reduced, passive ranging offers benefits both in terms of energy consumption and radio channel saturation. In a way, GNSS localization can be thought of as a special form of passive ranging, with the "tag" overhearing the satellites' messages without responding, although in this case the messages are precisely designed to help localize the tag. In a WSN context, we use the active ranging exchanges between other nodes to derive more data. Most of the current works use the TDoA localization scheme, so that passive nodes are fully passive, never sending any message. This can be done either with anchors overhearing anchor-tag ranging [19], tags overhearing anchor-anchor ranging [16], or tags overhearing anchor-tag ranging [20]. However, as these works rely on a large number of fixed anchors to be able to maintain good accuracy, their benefits in terms of energy saving and scalability do not easily extend to our low infrastructure needs.

III. ISRiCH

This section introduces ISRiCH, an infrastructure-light localization system designed for dense WSN deployments. We improve on the classical SAL scheme by refining the position estimations through the use of collaborative methods. Therefore, ISRiCH leverages *passive ranging* to compute the inter-tag distances without introducing more messages than SAL would. Inter-tag distances are then refined with a *drift compensation* mechanism, before processing them with MDS to derive the position of each tag.

A. Passive inter-tag ranging

Our central contribution is the design of a passive inter-tag ranging method that leverages the TWR messages a Single-Anchor Localization system uses to poll each tag. To do so, tags have to overhear RESP messages as depicted in Fig. 4 every time the anchor polls neighbor tags with TWR. In Fig. 4, when tag A sends its response to the anchor, it is overheard by neighbor tag B. Similarly, tag A listens to tag B's response to the anchor. For each overheard message, the tag records the time it has received it. Knowing the emission timestamps of RESP messages and the reception timestamps at tags A and B, we can derive the distance between A and B by adapting Eq. (1) to the appropriate timestamps:

$$d = c.\frac{(T^A_{B4} - T^A_{A3}) - (T^B_{B3} - T^B_{A4})}{2} \tag{2}$$

Here, T^X_{Mi} is the i-th timestamp collected by node X in the TWR procedure (cf. Fig. 1) on a RESP message sent by M.

Since the anchor is responsible for polling each tag, this passive ranging requires that every tag be in range of the anchor, a condition that is always satisfied in the considered use cases. Furthermore, to get a passive ranging between two tags, they need to be within range of each other. It is however not necessary for *all* the tags to be able to perform the ranging with *every other* tag. This situation will be addressed in Section III-C. If N tags are in line of sight of each other, overhearing allows us to derive at most $\binom{N}{2} = N(N-1)/2$ new distance measurements. For 10 tags, this amounts to $\binom{10}{2} = 45$ passive rangings.

The active and passive timestamps measured by the tags are collected and centralized at the anchor once every tag has been polled. Since tags are fixed, timestamp collection can be postponed to a later stage (for example, when the measured physical data is being collected). However, if localization information is needed earlier, tags are able to efficiently send timestamps to the anchor using concise serialization methods such as CBOR [23]. This communication overhead is further discussed in Section IV. With passive inter-tag ranging, we only need N TWR operations instead of $N + \binom{N}{2}$. It of course offers benefits on radio channel congestion, but it also reduces the duration needed to gather all timestamps, and thus the duration sensors have their radio on, preserving energy.

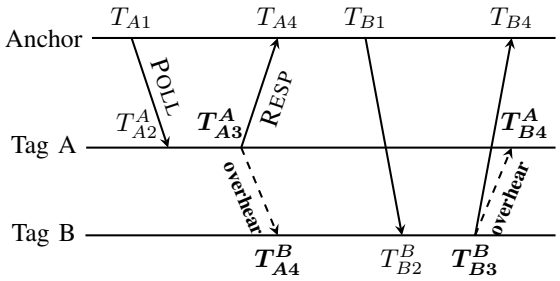


Fig. 4. Passive ranging procedure using overhearing with two tags.

B. Clock drift compensation

Due to the lower quality of embedded oscillators, internal clocks of tags are imperfect, resulting in varying clock frequencies. As such, we must remain cautious when using timestamps measured at different nodes in the same formula, as the time base for each one is different. For two nodes A and B, the relationship between the hardware time t^B of node B and the hardware time t^A of node A is modeled with a linear approximation: $t^B = \alpha_{B,A}t^A + \theta_{B,A}$. Here, t^A is the reference time base and $\alpha_{B,A}$ is the relative drift of the frequency of B to the frequency of A. While the offset $\theta_{B,A}$ is naturally compensated for in the TWR formula, the relative drift remains an issue. Thus, by identifying the time base used to measure each timestamp in Eq. (1), the distance obtained from a ranging between A and B is:

$$d = c.\frac{(T^A_4 - T^A_1) - (T^B_3 - T^B_2)}{2} \tag{3}$$

This equation can be expressed in the time base of node A:

$$d = c.\frac{(T^A_4 - T^A_1) - \alpha_{B,A}(T^A_3 - T^A_2)}{2} \tag{4}$$

The drift thus leads to a distance estimation error of:

$$error_d = \frac{c}{2}(1 - \alpha_{B,A})(T^A_3 - T^A_2) \tag{5}$$

From this result, we know that the ranging error due to the drift is greater if the delay between T^A_3 and T^A_2 is longer. With an active anchor-tag TWR, this delay can be minimized to a few milli- or microseconds. But our passive ranging scheme causes this delay to be much larger. Indeed, as illustrated in Fig. 4, tag A has to wait for tag B to be polled to retrieve the overheard timestamp T^A_{B4}. Thus, the longest overheard ranging delay is observed for the case where A is the tag polled first and B the tag polled last. For a setting with N tags, this longest overheard delay causes an imprecise distance that is not acceptable. For example, assuming $N = 100$ tags, a ± 2 PPM (Parts Per Million) clock drift, and a TWR operation lasting $10\,\mathrm{ms}$, Eq. (5) gives us a passive ranging error between the first and last tags to be polled that could grow up to $600\,\mathrm{m}$. As we aim at positioning sensors that are centimeters apart, such a level of error is absolutely detrimental. To overcome this major problem, it is either possible to minimize the inter-ranging delay or to find a way to compensate for relative clock drifts. Next, we propose to combine both solutions by adding a drift compensation stage that avoids introducing additional delay with the transmission of additional messages.

First, we note that classical drift compensation using CFO measurements [24], as is available on most recent UWB chips, cannot be applied here. The observations of such a system are time-varying by nature, and, if they are enough for active ranging, the delays in our system imply the need for a more smoothed-out estimation. We thus choose to compute the relative drift of each tag to the anchor by using the already acquired timestamp measurements. Deriving the drift of clock A relative to reference clock R is done by measuring the same duration Δ on both clocks. The drift is given by Δ^A/Δ^R.

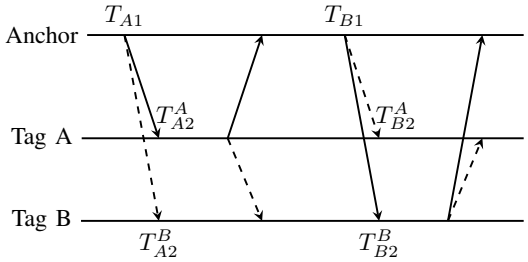


Fig. 5. Passive timestamping for the clock drift compensation mechanism.

Practically, we need two messages emitted by the anchor and received by tag A to compute this drift. When looking at the current timestamps recorded by ISRiCH in Fig. 4, there is one timestamp missing to measure the relative drift of a tag with respect to the anchor clock. Since it is detrimental to add any new message in the full ISRiCH polling sequence, we continue building on the passive timestamping idea to get the missing timestamp. As shown in Fig. 5, each polling message sent by the anchor can be overheard by all tags. Thus, each tag can calculate Δ, with timestamps originating from both its direct poll by the anchor and the overheard one. In order to get the best drift estimate for a given tag x, we record at x the timestamps of the $N-1$ overheard polls and perform the following linear regression:

$$\hat{\alpha}_{x,anchor} = \frac{\sum_{i=1}^{N}(T_{i1}-\bar{T_1})(T^x_{i2}-\bar{T^x_2})}{\sum_{i=1}^{N}(T_{i1}-\bar{T_1})^2} \tag{6}$$

with $\bar{T_1} = \frac{1}{N}\sum_{i=1}^{N} T_{i1}$ and $\bar{T^x_2} = \frac{1}{N}\sum_{i=1}^{N} T^x_{i2}$. We assume in this equation that each tag is numbered between 1 and N. From the drift of each tag, timestamps can be expressed in the anchor time base before inter-tag distances calculation.

The performance of our drift compensation method is evaluated experimentally by comparing the localization error of a classical active TWR between tags A and B and our passive drift-compensated ranging mechanism with respect to different artificially added delays between the ranging of tag A and that of tag B by the anchor. Results in Fig. 6 show that the drift-compensated passive ranging offers similar performance to an active TWR even with up to 10 seconds delay. It paves the way for great scalability, since hundreds of tags could be polled during this duration.

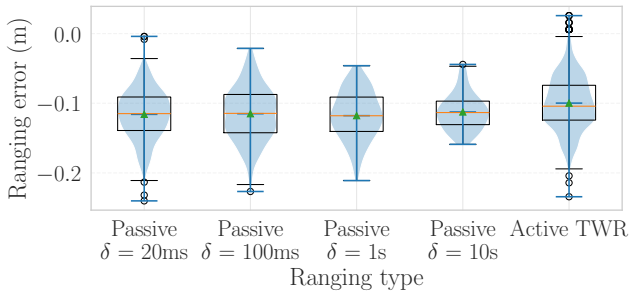


Fig. 6. Distribution of the localization error between active and drift-compensated passive ranging operations, with 300 measurements per scenario.

TABLE I
SIMULATION PARAMETERS

Parameter	Value
Angular error (°)	$\mathcal{N}(0, 3.3^2)$
Ranging error (m)	$\mathcal{N}(0, 0.25^2)$

C. Processing with MDS

Once the data acquisition has been done, we can represent the estimated distances in a matrix:

$$\boldsymbol{D} = \begin{pmatrix} 0 & d_{anchor,i} & d_{anchor,j} & \cdots & d_{anchor,N} \\ d_{anchor,i} & 0 & d_{i,j} & \cdots & d_{i,N} \\ \vdots & & & & \\ d_{anchor,N} & d_{i,N} & d_{j,N} & \cdots & 0 \end{pmatrix}$$

This matrix $\boldsymbol{D}$ can then be used as the dissimilarity input of multidimensional scaling. As done in [17], we use a weighted MDS and set the weights of missing distances to 0 to ignore them, improving the reliability of our system if inter-tag communication is not possible in the whole network. To avoid being trapped in a local minimum with MDS [18], we initialize optimization with position estimates computed with the regular SAL procedure, using the distance and angular measurements obtained by the anchor. The angular element of this initialization also offers benefits in preemptively resolving any symmetrical or rotational ambiguities, while the inclusion of the anchor, whose position is by definition known, resolves the translation ambiguity.

The computational complexity of the MDS is dominated by the SMACOF [21] algorithm and is $O(N^2)$ with N the number of tags. Though this computation can be made offline by the anchor, some caution is necessary when dealing with large networks, as detailed next.

IV. Experimental Setup

The performance of our system is evaluated using testbed experiments and simulations. ISRiCH has been deployed on the SLICES-FR IoT-LAB testbed [25] using Qorvo DWM1001-DEV Ultra-wideband boards. Since the testbed does not offer AoA-capable nodes, we artificially add noisy angular measurements. Simulated angles are generated with a randomness factor based on the AoA performance results in [5], [8]. Moreover, a simulator has been developed[1] to further diversify the test setups and evaluate the scalability of ISRiCH. The choice of an in-house simulator is motivated by the need for modeling the greater influence of clock drift.

Simulations have been calibrated with experimental results to be as faithful as possible to real-world conditions. The simulator calibration parameters are given in Table I. Both angular and ranging errors follow normal distributions with zero mean, assuming that no global bias is present. The standard deviation of these distributions is chosen to fit practical evaluations and confidence intervals [5], considering the case where only one measurement is taken per value to reduce the number of messages, rather than repeating the ranging to compute the

[1] https://gitlab.irit.fr/rmess/bperroux/isrich-sim

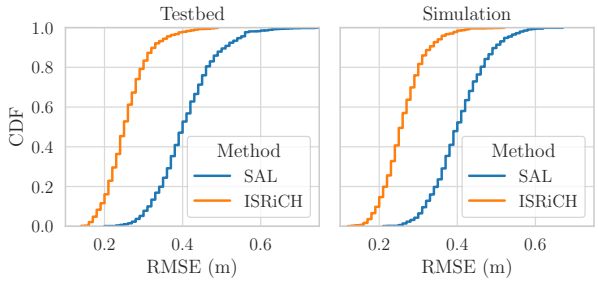


Fig. 7. Comparison of the RMSE between testbed data and simulation using the same positions.

average distance. Evaluation of our simulation parameters with respect to our testbed is presented in Section V, along with an analysis of their effect on SAL's and ISRiCH's performance.

Furthermore, practical implementation aspects have been evaluated on our testbed. As our passive ranging mechanism implies the timestamping of more messages than active ranging, memory usage and data collection must be dimensioned accordingly. Each ranging entry in memory is composed of the polled tag ID and the timestamps for the polling and responses. On DW1000 chips, these timestamps are stored on 40 bits. Considering a tag ID on 2 bytes, the memory needs are 12 bytes per tag. For a fully-connected network of N tags, storing all overheard timestamps totals $12(N-1)$ bytes per tag. With one DW1000 maximum frame size of 1023 bytes [26], a tag can upload ranging data of $N = 80$ neighbors without compression, keeping data collection a light operation.

Once ranging data is collected, the question of the scalability of the MDS computation remains. We have executed the MDS algorithm on a Raspberry Pi 3B that could be used as the anchor controller. Observations confirm the $O(N^2)$ evolution of the time complexity, reaching an average execution time of $11.5\,\mathrm{s}$ for $N = 100$. Since MDS is not meant to be executed in real-time, as would be needed for tracking applications for example, this execution time is acceptable for our static tags.

V. Localization accuracy

We first aim at evaluating the localization accuracy of ISRiCH. To do so, we perform 1000 repetitions of ISRiCH and SAL on FIT IoT-Lab. For each localization method, we evaluate the localization accuracy using the root mean square error (RMSE) metric: $RMSE = \sqrt{\frac{1}{N}\sum_i^N e_i^2}$. Fig. 7 represents the empirical cumulative distribution function (CDF) of the RMSE for SAL and ISRiCH on the testbed. We observe that half of the time, ISRiCH results in a RMSE of less than $0.25\,\mathrm{m}$, with a 95% threshold of $0.36\,\mathrm{m}$. Comparatively, the median RMSE of SAL is $0.40\,\mathrm{m}$ with a 95% threshold of $0.55\,\mathrm{m}$. Compared to a mean absolute error (MAE) metric, RMSE reveals outliers thanks to the squared error. As such, this difference in RMSE not only shows the better global accuracy of ISRiCH but also the reduced risk for large individual errors.

To diversify the scenarios tested, the next results are produced using our in-house simulator introduced in Section IV. Before using the simulator, we have calibrated it to assert that our simulations match the experiments done earlier. Fig. 7 compares SAL and ISRiCH performance obtained on the testbed and in simulation. We observe very similar performance, with SAL anchor performance close to the one in previous works [8]. Simulations will help us expand investigated scenarios.

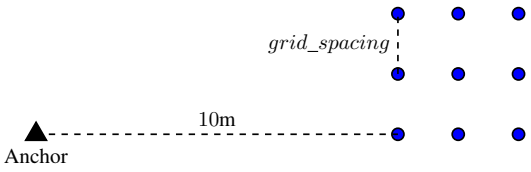


Fig. 8. Experimental setup for the grid scenarios, here with 9 tags

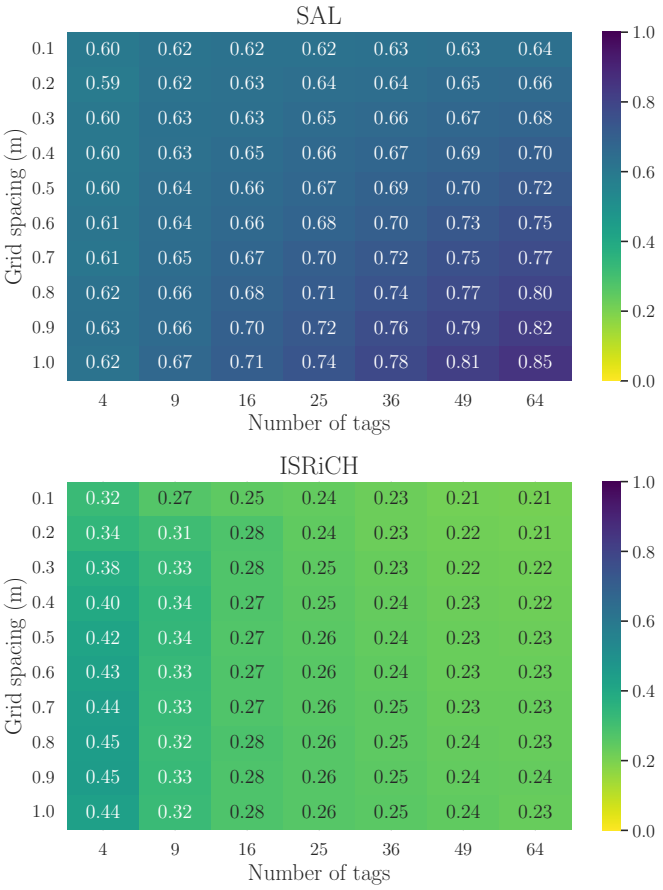
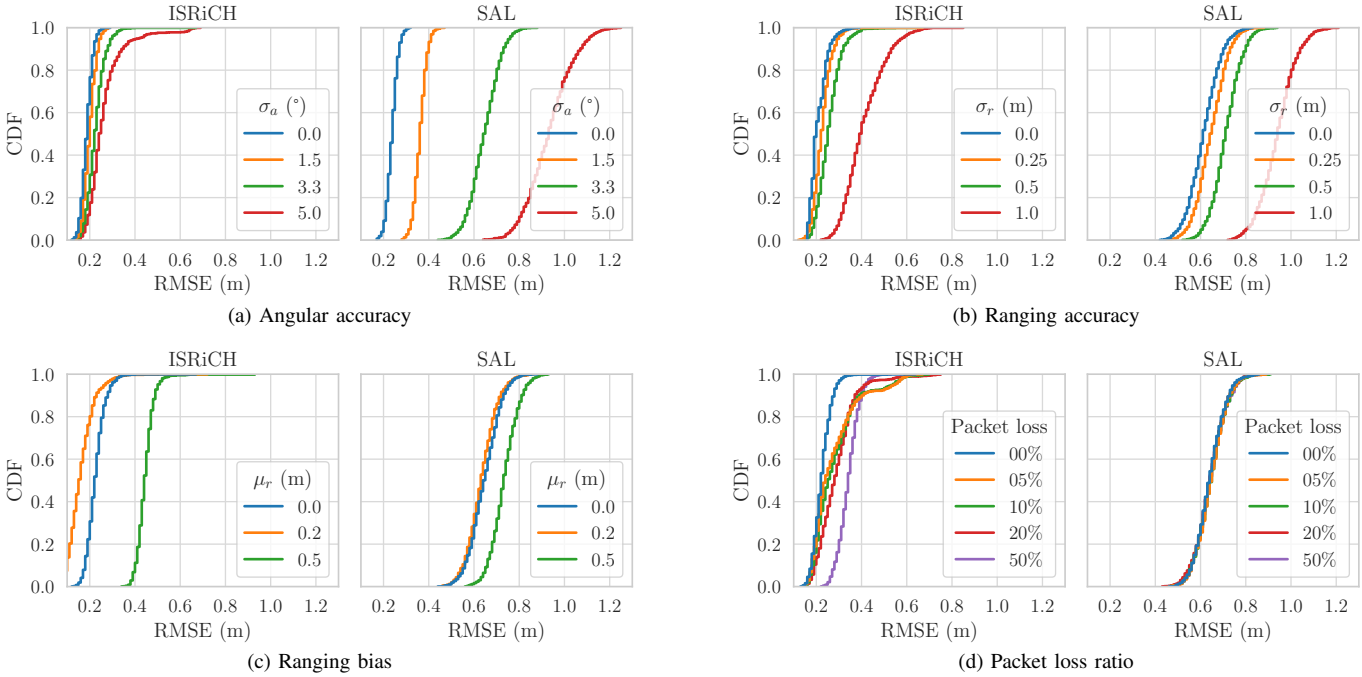


Fig. 9. Mean RMSE of SAL and ISRiCH on grid scenarios (1000 repetitions)

To mimic the use case of modal analysis on a structure, we test scenarios where a given number of tags is positioned on a regular grid. The anchor is located at $10\,\mathrm{m}$ from the nearest tag, as represented in Fig. 8. The *grid_spacing* parameter varies from $0.1\,\mathrm{m}$ to $1\,\mathrm{m}$. It is thus small, but not necessarily negligible compared to the anchor-tag distances. We perform simulations for a varying number of tags and *grid_spacing* values, setting 1000 repetitions per scenario.

Average RMSE results are given in Fig. 9, showing the clear superiority of ISRiCH compared to SAL. As network and grid size vary, the two solutions exhibit different accuracy trends. For SAL, the best performance is obtained for a reduced number of tags and small grid spacings. This can be explained



(a) Angular accuracy

(b) Ranging accuracy

(c) Ranging bias

(d) Packet loss ratio

Fig. 10. CDF of the RMSE of ISRiCH and SAL for a grid with $grid_spacing = 0.2m$ with $N = 36$, while varying different simulation parameters.

by simple trigonometry: for a given angular error, localization error is proportional to the anchor-tag distance. Since SAL relies heavily on angular measurements, increasing N and $grid_spacing$ increases the distance between anchor and tags and thus degrades localization accuracy. On the contrary, ISRiCH's accuracy improves with the number of tags. This is exactly what we expected since positioning benefits from $\binom{N}{2}$ passive ranging distances providing additional information. However, we observe that improvements slow down starting at $N = 25$ tags. From this observation, it calls for clustered implementations in the future that would reduce memory and radio usage needs and add another layer of scalability.

Additionally, we evaluate the impact of the anchor's angular accuracy on SAL and ISRiCH performance. Fig. 10a compares the RMSE for both methods when varying the standard deviation of the angular error generation, σ_a (angular error follows a normal law $\mathcal{N}(0, \sigma_a^2)$, cf. Section IV). Results show that while SAL is highly dependent on the angular accuracy to position the tags, the influence of this parameter is limited with ISRiCH. The resistance of ISRiCH to angular errors stems from the inter-tag ranging that mitigates excessive errors in MDS. As such, ISRiCH could allow for the use of lower-precision, and thus lower-cost, anchors while still maintaining good localization accuracy.

A similar evaluation has been done on both the standard deviation σ_r (Fig. 10b) and the bias μ_r (Fig. 10c) of the ranging error parameter of the simulation following the normal law $\mathcal{N}(\mu_r, \sigma_r^2)$. Indeed, while our simulator was calibrated with both testbed data and a literature review, challenging environments may cause increased errors due to NLOS conditions. As an example, deployment inside an aircraft cabin may suffer from signal reflections, artificially lengthening the measured distance and causing a positive bias. Results show a similar impact of σ_r on both localization methods, indicating neither a competitive advantage nor a flaw of ISRiCH in this respect, while a positive bias μ_r hinders ISRiCH more.

Finally, we investigated the effect of another risk of challenging environments: packet loss. As ISRiCH needs a bidirectional overhearing to compute inter-tag distances, this point of concern must not be ignored. While SAL is naturally resistant, as packet loss between anchor and tags is detected and causes retransmissions, ISRiCH's passive ranging suffers more from this effect. However, results in Fig. 10d still show a good resilience, as even with a lower number of passive rangings, ISRiCH strongly improves on classical SAL, as seen in Fig. 9.

VI. Matching Problem

A. Goal and method

When deploying sensors for mechanical wave propagation analysis, the sensors are generally placed at predefined points or following a specific pattern. For the modal analysis of an aircraft wing, accelerometers have to be placed on a grid spanning the structure. As such, we have a list of candidate positions. To avoid the manual mapping of sensor identities to candidate positions, we aim to automatically and uniquely associate the tag's identity with one of the candidate locations on the grid. To do so, we leverage a localization protocol (SAL or ISRiCH) that derives, for each tag, an estimated position. Then, our problem boils down to assigning the candidate location to the estimated ones, uniquely.

To perform the matching, we rely on the *Kuhn-Munkres (KM) matching algorithm* [27]. This algorithm finds the association $m = (x_1, \ldots, x_i, \ldots, x_N)$ with x_i the identifier of the tag in candidate position i, enforcing $x_i \neq x_j \quad \forall (i, j), i \neq j$.

In our implementation, the KM algorithm minimizes the sum, over all tags, of the squared Euclidean distance between the tag's position estimate $p(x_i)$ and its assigned candidate position $c(i)$. Formally, we have:

$$m_{opt} = \arg\min_{m \in M} \sum_{i \in [1,N]} \text{dist}\left(c(i), p(x_i)\right)^2 \tag{7}$$

Compared to other methods such as *stable matching* [28], the KM algorithm is resistant to biases carried over from bad angular measurements as it looks for a global optimum rather than favoring strong associations at the cost of a few very bad ones (this is the benefit of summing the squared distances and not the distances directly). The KM algorithm has a complexity of $O(N^3)$. It is executed on the anchor. With our basic Raspberry Pi 3B, the average execution time of this algorithm for $N = 100$ tags is 13 s. Together with the ISRiCH localization, it takes less than 25 s to obtain a match.

B. Results

To evaluate the performance of the matching algorithm, we define as a *perfect match* the situation where all tags have been correctly identified. The percentage of perfect matches over 1000 runs of each simulated configuration is shown in Fig. 11. While the SAL system can provide perfect matches for a low number of tags on large grids, its ability to do so strongly decreases in more challenging networks. On the other hand, ISRiCH, while still struggling with the smallest grid spacings, can almost ensure perfect matching in more spread-out setups or with a large number of tags. Additionally, even when the grid spacing is reduced to only 20 cm, rolling out our algorithm with at least 36 tags yields a perfect match more than half of the time. Both of these observations can be linked to the performance in RMSE: the larger the RMSE relative to the grid spacing, the harder it is to obtain a good match.

To enforce the computation of a perfect match, we evaluate the potential of repeating the full ISRiCH algorithm, from data acquisition to matching, until a match is selected using a *match election* strategy. This election consists in comparing successively computed $m^1_{opt}, m^2_{opt}, \ldots, m^R_{opt}$ matches to find similarities between them. Since a match is a quite large vector and matching errors being localized, finding twice the same match is a good first hint of having found a correct association. To evaluate the performance of this election strategy, we perform successive runs until this criterion is met and compare this elected match with the perfect match.

Performance is evaluated with our simulator for 25-tag grids of varying grid spacing. Fig. 12 shows the number of ISRiCH runs before a match election is conclusive. In all scenarios tested, our procedure resulted in the election of the perfect match. For grid spacing values of 30 cm or more, a match is found in 2 runs of ISRiCH 50% of the time. Furthermore, while a spacing of 30 cm needs a maximum of 8 runs, for grid spacings of 70 cm or more, two runs are enough more than 99% of the time. The more constrained case, with only 20 cm in between tags, is however harder to conclusively match.

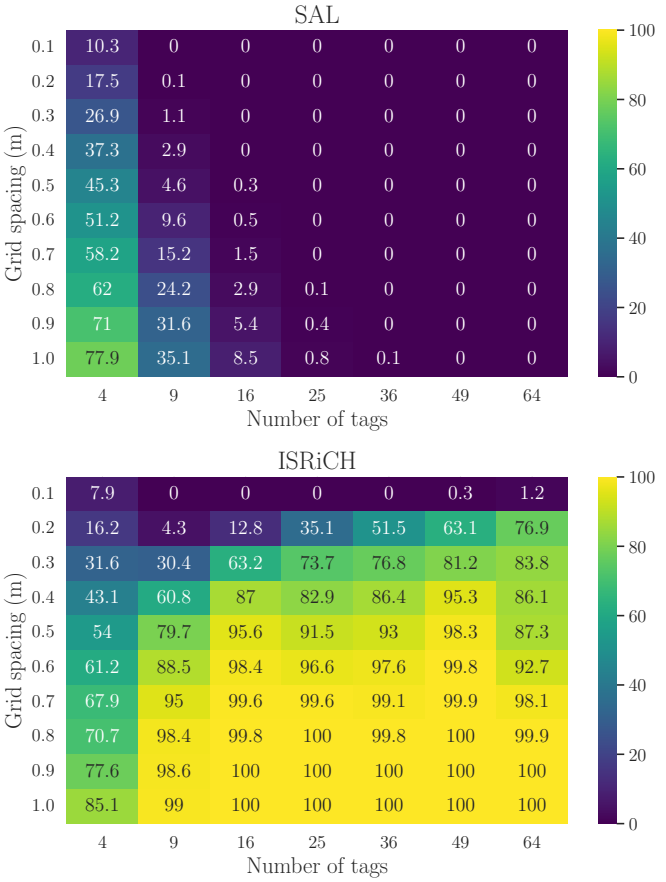


Fig. 11. Percentage of perfect matches on grid scenarios, with 1000 repetitions

The required number of runs is linked to the ratio of perfect matches observed in Fig. 11. Assuming that each ISRiCH run is an independent Bernoulli trial with a success probability being the observed perfect match ratio for the scenario of interest, denoted p_m, the total number of runs until obtaining twice the same match follows a negative binomial distribution. Then, the probability of needing exactly n runs is given by:

$$P(X = n) = (n-1)p_m^2(1-p_m)^{n-2}, \qquad n \geq 2 \tag{8}$$

The application of Eq. (8) to the case with $N = 25$ tags and $grid_spacing = 0.2$ gives us a probability of success in 2 runs of 0.12 (using $p_m = 0.351$), as observed in Fig. 12. Similarly, the same formula can be leveraged to determine that ISRiCH, for a grid of 64 tags and 20 cm spacing, succeeds in a 2-run election more than 50% of the time ($p_m = 0.769$), with at most 8 runs.

We can conclude that the matching election enables us to ensure the choice of a correct matching with only one anchor. However, for the dense setups, this mechanism can be costly as it implies acquiring new ranging data and computing the MDS and matching algorithms multiple times.

VII. Conclusion & Future work

In this paper we have presented ISRiCH, a localization protocol for temporarily deploying a dense WSN. With only one anchor, ISRiCH improves state-of-the-art Single-Anchor

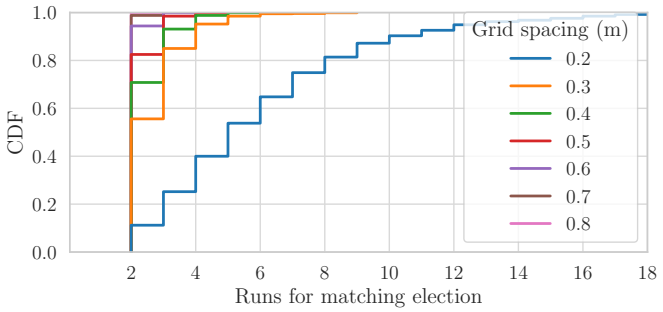


Fig. 12. Empirical cumulative distribution function (CDF) of the number of runs before election of a perfect matching, on a grid of 25 tags.

Localization systems. Our contribution is to introduce passive ranging and timestamping to get inter-tag ranging and perform network-wide drift compensation. Inter-tag ranging data is processed with multidimensional scaling to position the sensors relative to each other, the anchor location being used to find absolute positions. With both simulations and experimentation on the SLICES-FR IoT-LAB testbed [25], we show that this holistic approach improves the accuracy by a factor of 3, reaching a median RMSE of below 50 cm more than 90% of the time. In matching scenarios, where the estimated positions of the tags must be associated with a set of candidate sensor positions, ISRiCH strongly outperforms SAL systems by providing perfect matches more than 90% of the time on large grids. Furthermore, we show that the repetition of the algorithm with a matching election mechanism finds correct matches in less than 6 runs more than 99% of the time for grids of 30 cm spacings.

Future work will extend ISRiCH to 3D localization with the use of novel 3D-PDoA enabled anchors to encompass more complex use cases. Practical deployments of the extended 3D-ISRiCH protocol are envisioned to demonstrate its capacity in providing efficient modal or acoustic analysis thanks to localization and matching.

References

[1] M. Cobos, F. Antonacci, A. Alexandridis, A. Mouchtaris, and B. Lee, "A Survey of Sound Source Localization Methods in Wireless Acoustic Sensor Networks," *Wireless Communications and Mobile Computing*, vol. 2017, pp. 1–24, 2017.

[2] S. Yiu, M. Dashti, H. Claussen, and F. Perez-Cruz, "Wireless RSSI fingerprinting localization," *Signal Processing*, vol. 131, pp. 235–244, 2017.

[3] F. Despaux, A. van den Bossche, K. Jaffrès-Runser, and T. Val, "N-TWR: An accurate time-of-flight-based N-ary ranging protocol for Ultra-Wide band," *Ad Hoc Networks*, vol. 79, pp. 1–19, 2018.

[4] IEEE Computer Society, "IEEE Std 802.15.4a," IEEE Computer Society, Standard 802.15.4a-2007, Aug. 2007.

[5] G. K. J. Fischer, T. Schaechtle, A. Gabbrielli, J. Bordoy, I. Häring, F. Höflinger, and S. J. Rupitsch, "A Systematic Survey and Comparative Analysis of Angular-Based Indoor Localization and Positioning Technologies," *IEEE Communications Surveys & Tutorials*, vol. 28, 2026.

[6] J. Prince Mathew and C. Nowzari, "ReLoki: A Light-Weight Relative Localization System Based on UWB Antenna Arrays," *Sensors*, vol. 24, no. 1616, p. 5407, Jan. 2024.

[7] M. Zhao, T. Chang, A. Arun, R. Ayyalasomayajula, C. Zhang, and D. Bharadia, "ULoc: Low-Power, Scalable and cm-Accurate UWB-Tag Localization and Tracking for Indoor Applications," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 3, Sep. 2021.

[8] B. Van Herbruggen, S. Luchie, R. Wilssens, and E. De Poorter, "Single Anchor Localization by Combining UWB Angle-of-Arrival and Two-Way-Ranging: an Experimental Evaluation of the DW3000," in *2024 International Conference on Localization and GNSS (ICL-GNSS)*. IEEE, Jun. 2024, p. 1–7.

[9] J. Tiemann, O. Fuhr, and C. Wietfeld, "CELIDON: Supporting First Responders through 3D AOA-based UWB Ad-Hoc Localization," in *2020 16th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. IEEE, Oct. 2020, p. 20–25.

[10] M. Rostami and K. Sundaresan, "Enabling high accuracy pervasive tracking with ultra low power UWB tags," in *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking (MobiCom)*. ACM, 2022, p. 459–472.

[11] T. Savić, X. Vilajosana, and T. Watteyne, "Constrained Localization: A Survey," *IEEE Access*, vol. 10, p. 49297–49321, 2022.

[12] J. Xu, M. Ma, and C. L. Law, "Position Estimation Using UWB TDOA Measurements," in *2006 IEEE International Conference on Ultra-Wideband*. IEEE, 2006, pp. 605–610.

[13] P. Grasso, M. S. Innocente, J. J. Tai, O. Haas, and A. M. Dizqah, "Analysis and Accuracy Improvement of UWB-TDoA-Based Indoor Positioning System," *Sensors*, vol. 22, no. 23, p. 9136, 2022.

[14] T. Margiani, S. Cortesi, M. Keller, C. Vogt, T. Polonelli, and M. Magno, "Angle of Arrival and Centimeter Distance Estimation on a Smart UWB Sensor Node," *IEEE Transactions on Instrumentation and Measurement*, vol. 72, p. 1–10, 2023.

[15] M. Elsanhoury, P. Mäkelä, J. Koljonen, P. Välisuo, A. Shamsuzzoha, T. Mantere, M. Elmusrati, and H. Kuusniemi, "Precision Positioning for Smart Logistics Using Ultra-Wideband Technology-Based Indoor Navigation: A Review," *IEEE Access*, vol. 10, pp. 44 413–44 445, 2022.

[16] H. Chen and A. Dhekne, "PnPLoc: UWB Based Plug & Play Indoor Localization," in *2022 IEEE 12th International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, Sep. 2022, p. 1–8.

[17] P. Corbalán, G. P. Picco, M. Coors, and V. Jain, "Self-Localization of Ultra-Wideband Anchors: From Theory to Practice," *IEEE Access*, vol. 11, p. 29711–29725, 2023.

[18] N. Saeed, H. Nam, T. Y. Al-Naffouri, and M.-S. Alouini, "A State-of-the-Art Survey on Multidimensional Scaling-Based Localization Techniques," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, p. 3565–3583, 2019.

[19] T. Laadung, S. Ulp, M. M. Alam, and Y. L. Moullec, "Novel Active-Passive Two-Way Ranging Protocols for UWB Positioning Systems," *IEEE Sensors Journal*, vol. 22, no. 6, p. 5223–5237, Mar. 2022.

[20] S. Cortesi, L. Schulthess, D. Plozza, C. Vogt, and M. Magno, "Eco-WakeLoc: An Energy-Neutral and Cooperative UWB Real-Time Locating System," 2026.

[21] J. de Leeuw, "Application of convex analysis to multidimensional scaling," in *Recent developments in statistics*, O. Barndorff-Nielsen, J. L. Jensen, and W. S. Kendall, Eds. North-Holland Publishing Company, 1977, pp. 133–145.

[22] C. Di Franco, E. Bini, M. Marinoni, and G. C. Buttazzo, "Multidimensional scaling localization with anchors," in *2017 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*. IEEE, Apr. 2017, p. 49–54.

[23] C. Bormann and P. E. Hoffman, "Concise Binary Object Representation (CBOR)," RFC 8949, Dec. 2020.

[24] J. Krška and V. Navrátil, "Utilization of Carrier-Frequency Offset Measurements in UWB TDoA Positioning with Receiving Tag," *Sensors*, vol. 23, no. 5, p. 2595, 2023.

[25] C. Adjih, E. Baccelli, E. Fleury, G. Harter, N. Mitton, T. Noel, R. Pissard-Gibollet, F. Saint-Marcel, G. Schreiner, J. Vandaele, and T. Watteyne, "FIT IoT-LAB: A Large Scale Open Experimental IoT Testbed," in *IEEE World Forum on Internet of Things (IEEE WF-IoT)*. IEEE, Dec. 2015, pp. 459–464.

[26] Qorvo, Inc., *DW1000 Ultra-Wideband Transceiver Datasheet*, Qorvo, 2017, version 2.21. [Online]. Available: https://www.qorvo.com/products/d/da007946

[27] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[28] D. Gale and L. S. Shapley, "College admissions and the stability of marriage," *The American Mathematical Monthly*, vol. 69, no. 1, pp. 9–15, 1962.