

This item is the archived peer-reviewed author-version of:

Improving the accuracy and robustness of ultra-wideband localization through sensor fusion and outlier detection

Reference:

Dw ek Nathan, Birem Merw an, Geebelen Kurt, Hostens Erik, Mishra Anurodh, Steckel Jan, Yudanto Risang.- Improving the accuracy and robustness of ultra-w ideband localization through sensor fusion and outlier detection
IEEE robotics and automation letters - ISSN 2377-3766 - 5:1(2020), p. 32-39
Full text (Publisher's DOI): <https://doi.org/10.1109/LRA.2019.2943821>
To cite this reference: <https://hdl.handle.net/10067/1635610151162165141>

Improving the Accuracy and Robustness of Ultra-wideband Localization through Sensor Fusion and Outlier Detection*

Nathan Dwek^{1,3}, Merwan Birem¹, Kurt Geebelen¹, Erik Hostens¹,
Anurodh Mishra¹, Jan Steckel², and Risang Yudanto¹

Abstract—This article presents sensor fusion techniques for ultra-wideband-based localization to achieve sufficient accuracy and robustness for the control of AGVs in an industrial environment. We propose two outlier detection methods in combination with an EKF, and present experimental validation where 10 cm accuracy is achieved even in difficult NLOS conditions.

Index Terms—Localization, Sensor Fusion, Outlier Detection, Range Sensing, Ultra-wideband

I. INTRODUCTION

As research into indoor localization is receiving attention [1], [2], ultra-wideband (UWB) based localization has emerged as one candidate solution, with the performance and availability of commercial off the shelf hardware as well as turnkey solutions growing steadily [3], [4]. The principle of operation is to install anchors (also sometimes denoted as beacons in the literature) at known locations in the environment and to measure through the exchange of UWB signals the distance and/or the angle between those anchors and an antenna installed on the person or object to be localized (referred to as the tag). From those ranges or angles and the known anchors positions, the tag can then be localized by performing multilateration (for ranges) or triangulation (for angles) [5].

UWB-based localization has however not yet found its exact niche. Since it is infrastructure-based, setup cost is high. This makes it inadequate for lower-cost applications where the target accuracy is in the order of meters, an order of magnitude worse than what can be achieved using UWB [6]. Furthermore, it can only track persons or objects that are equipped with a tag, which must include a UWB-specific transducer. For the tracking of crowds for example, smartphone-based localization techniques using WiFi Received Signal Strength (RSS) or fingerprinting are more popular solutions [1].

For use-cases where the higher requirements on accuracy may justify the infrastructure cost, UWB-based localization has not proven reliable enough. This is especially true in industrial environments where metallic obstructions and

sources of electro-magnetic interferences (EMI) are present, resulting in the measurement of ranges failing or being biased and inaccurate. Such events are generically referred to as UWB outliers in this article. This has driven past research to focus on detecting and handling Non-Line-of-Sight (NLOS) conditions and improving the overall accuracy and reliability of the localization in realistic conditions [7]–[14].

In this article, we present sensor fusion techniques to tackle those issues and specifically the detection of UWB outliers, with a focus on the localization of industrial vehicles. In section II, we further detail the problem statement and present the proposed approach. This leads in section III to the general design of the sensor fusion. Then, two novel outlier detection techniques are presented in section IV. Those techniques are experimentally validated together with the general sensor fusion in section V, where the setup is presented, datasets are made available [15], and performance metrics are shown and compared to previous results. We conclude in section VI where we also look at further improvements to our work.

II. PROPOSED APPROACH

We present in this work sensor fusion techniques for the localization of vehicles in an industrial environment, based on UWB ranging data. The goal is to reach accuracy, refresh rate and robustness all sufficient for the control of automated guided vehicles (AGVs).

Typical UWB polling rates, in the order of tens of Hz for affordable large scale deployments with many tags, will not be sufficient to reach this goal. Moreover, due to EMI and NLOS conditions, sufficient UWB data might not be available at all times, and some ranges might have to be discarded as outliers. For those two reasons, sensor fusion is needed of UWB ranging with one or several faster and/or more reliable sensors. As the industrial vehicles we consider are typically equipped with wheel encoders, we choose to use encoder data as an additional measurement. This is a standard scheme where dead-reckoning, which ensures continuity of the estimated position, is used in conjunction with a source of absolute position to compensate for the integration drift [7].

As the vehicles considered in this work are expected to drive on relatively flat surfaces, the focus is set on 2D localization in the horizontal plane, with some prior knowledge on the vertical position being available. Full 6 degrees of freedom pose estimation is not required, but the yaw angle will be estimated, both because it is necessary for the rest

*This work has been carried out within the framework of Flanders Make's ICON project LOCATOR (LOCalization system for Accurate Tracking and navigation for autonomous OpeRation, HBC.2016.0469) funded by the agency Flanders Innovation & Entrepreneurship (VLAIO) and Flanders Make. Flanders Make is the Flemish strategic research centre for the manufacturing industry.

¹Flanders Make Strategic Research Centre, Belgium

²FTI CoSys Lab, University of Antwerp, Belgium

³nathan.dwek@flandersmake.be

of the sensor fusion, and because it is a required feedback signal for the control of an AGV.

To reach the target accuracy, the focus is set on improving sensor fusion techniques, and not on designing custom UWB hardware or sensors. In section V the proposed techniques are validated using two-way ranging (TWR) [16] with commercial off-the-shelf hardware. TWR was chosen as it does not require very accurate clock synchronization between anchors, unlike time-of-arrival (ToA) or time-difference-of-arrival (TDoA) based schemes [16]. It is expected that the developed techniques translate to other UWB hardware and that they are sufficiently generic such that minimal effort is required to adapt them to ranging schemes which do not output time-of-flight (ToF) measurements such as ToA and TDoA. Angle-of-arrival (AoA) based schemes are not directly considered in this article, but the principles behind the proposed outlier detection methods still apply.

In this work, the main source of improvement is the development of better outlier detection techniques for the UWB range measurements. This topic has received attention in the past already [7]–[11], [14], and remains a challenge for the deployment of UWB in industrial environments. One of the outlier detection techniques presented in this article relies on channel power diagnostics being available, but we also propose a technique that does not require them, and that as a consequence should be applicable in all situations and for other types of range sensors.

III. GENERAL SENSOR FUSION DESIGN

An extended Kalman Filter (EKF) is used to perform sensor fusion of the UWB ranges and of the encoder data, based on the vehicle model. It is not the goal of this article to present a detailed derivation of such an EKF, as this topic is very well covered in literature [7], [8], [17]. Some highlights and design choices are however given below. First, the traditional notations are introduced.

The EKF estimates the state x of a vehicle based on measurements y . The vehicle is modeled as a dynamical system by (1) and (2), where f describes the dynamics of the system and h describes how the system is observed. u is the system input and q and r are sources of noise.

$$\dot{x} = f(x, u, q) \quad (1)$$

$$y = h(x, r) \quad (2)$$

For the purposes of the EKF, this model must be discretized in time and linearized at every time-step around the current values of x , u and y into (3) and (4), which now deal with small-signal variables.

$$\tilde{x}[t+1] = A[t]\tilde{x}[t] + B[t]\tilde{u}[t] + \tilde{q} \quad (3)$$

$$\tilde{y}[t] = H[t]\tilde{x}[t] + \tilde{r} \quad (4)$$

Here, \tilde{q} and \tilde{r} are zero-mean Gaussian random variables (we denote such variables simply as Gaussian in the rest of this article for the sake of conciseness), with covariance matrices Q and R . The covariance matrix of the uncertainty

on the state estimation is denoted P . Finally, the innovation s is the prediction error as given in (5).

$$s = y_{meas} - y_{pred} \quad (5)$$

It is also a Gaussian random vector, with its covariance S given by (7).

$$s \sim \mathcal{N}(\bar{0}, S) \quad (6)$$

$$S = HPH^T + R \quad (7)$$

Similarly to [7], the EKF used in this work has as variable time-step Δt in order to accommodate measurements at different sample rates, and handles a varying number of measurements being received at once. At each time-step, numerical integration of the continuous time vehicle dynamics (1) is used to perform the prediction step. The predicted state $x(t + \Delta t)$ is computed, along with $x(t + \frac{\Delta t}{2})$, which is used to derive the A matrix for that time-step. Both the A and H matrices are computed by performing symbolic automatic derivation using CasADi [18], and H is derived taking into account the exact position of the tag the vehicle. This results in the UWB ranges actually contributing to the orientation estimation.

The vehicle model also allows to correctly tune the different sources of noise in the EKF representation. For example, the sideways speed noise is limited by how much the wheels can slip on the ground and is lower than the forward and yaw acceleration which are simply unknown system inputs.

The vertical coordinate z is part of the estimated state. It is modeled as constant, with variations only due to a source of noise on the vertical speed. The standard deviation of that source of noise is set very low since, as mentioned earlier, very good prior knowledge of the z position is available.

The encoder measurement model includes a slowly-varying bias, as indicated in (8) and (9) for $i = 1, \dots, n_{encoders}$.

$$v_{meas}^i = v^i + \delta^i + r_v^i \quad (8)$$

$$\delta^i[t+1] = \delta^i[t] + w^i \quad (9)$$

Here, v^i denotes the speed measured by a given encoder and the bias δ^i is modeled as a slow random walk. r_v^i are independent and identically distributed (i.i.d.) Gaussian noises, and so are w^i .

Multilateration is not done as a separate pre-processing step, but the EKF works directly with UWB ranges. Reference [7] lists the advantages of this more elegant approach. It allows to always maximally use the available UWB information, even if not enough ranges (less than respectively 3 or 4 in the 2D or 3D case) are available to perform multilateration. Crucially to this work, it also allows to perform outlier detection at the level of individual UWB ranges rather than at the level of the position computed through multilateration of all the ranges, outliers or not. We add to these arguments of [7] that working directly with ranges has the advantage that the EKF will then implicitly handle the geometrical dilution of precision (GDOP) resulting from the selected anchors and current vehicle position; only the ranging measurement noise

is required for each anchors. Those sources of noise are assumed i.i.d.

At each time-step where UWB ranges are available, outlier detection is performed before the update step, in order to select only the reliable measurements which will contribute to that update step.

IV. OUTLIER DETECTION

Two outlier detection methods are proposed, one is based purely on the estimation of S , the other relies on channel power diagnostics.

A. Based on the Estimated Covariance of the Innovation

Hypothesis testing has already been proposed for the detection of outliers in the context of UWB localization [7], [8]. The idea is to determine whether the prediction error s (possibly multidimensional) can be attributed with sufficient probability to its supposed distribution $\mathcal{N}(0, S)$.

Reference [7] proposes to look at the marginal distribution of every element (i.e. sensor or anchor) of the prediction error, and to discard individual measurements which fall outside of their confidence interval (e.g. 3σ). This has the advantage of treating every UWB range individually, but it does not take cross-correlations between those ranges into account. The UWB ranges are heavily correlated since they measure the distances between the different anchors and a single tag on the vehicle. Checking that the correlations are respected is important as it enforces that the ranges agree on a unique position of the tag.

Reference [8] proposes to use the Mahalanobis distance to check whether the entire measurement vector can be attributed to the supposed distribution. This is a rigorous way to take into account the cross-correlations between the UWB ranges. However, it does not allow to identify which individual range or ranges must be discarded, and all ranges must be valid at once for a measurement vector to be kept. This makes it inadequate if frequent drops of only some of the anchors are expected.

We propose a method that allows to detect individual outlying range or ranges while taking into account the cross-correlations between them. The principle is to compute for each element of the prediction error s its distribution conditioned on all the other elements $P(s_i | s_{j \neq i})$. In this case where $s \sim \mathcal{N}(0, S)$, this univariate conditional distribution is also Gaussian, with mean and variance given by (10) and (11) from [19, Section 3.4].

$$\mu_i = S_{ij} S_{jj}^{-1} s_{j \neq i} \quad (\text{Correlation}) \quad (10)$$

$$\sigma_i^2 = S_{ii} - S_{ij} S_{jj}^{-1} S_{ij}^T \quad (\text{Narrowing}) \quad (11)$$

Here, for N simultaneous measurements, S is size $N \times N$, S_{jj} is S without its i th row and i th column (size $(N-1) \times (N-1)$) and S_{ij} is the i th row of S without its i th element (size $1 \times (N-1)$).

Intuitively, using those conditional probabilities enforces that measurements agree with each other, but also allows a larger disagreement with the prediction if there is indeed consensus amongst the measurements. There is however

one problem that if a single strong outlier occurs, all the other measurements, although valid, become very unlikely conditioned on that single outlier, which leads to false-positives through contamination.

To remedy this, we propose algorithm 1 which proceeds

Algorithm 1 Iterative S -based outlier detection

Require: Outlier threshold $n > 0$

repeat

for every non-outlying element s_i of s **do**

 Compute μ_i and σ_i of $P(s_i | s_{j \neq i})$ with (10), (11)

 Build i th confidence interval $[\mu_i - n\sigma_i, \mu_i + n\sigma_i]$

end for

if one or several s_i out of their confidence interval **then**

 Find the s_{i_M} furthest outside of its interval

 Remove only s_{i_M} from s

 Remove only the i_M th column and row from S

end if

until no outlier found or no measurement remaining

return s and S to be used by the EKF update

iteratively and eliminates only the strongest outlier at each iteration. This means that if a single strong outlier occurs, it will be eliminated first and will not contaminate the conditional distributions of the other measurements during the next iterations.

B. Based on Channel Power Diagnostics

Identifying NLOS conditions based on features of the complex impulse response (CIR) of the UWB channel has often been considered [9]–[14]. References [12], [13] investigate the choice of those features. Reference [12] proposes a classifier to discriminate between LOS and NLOS conditions, while reference [13] investigates the direct correlation between those features and the ranging error.

A similar but slightly simpler approach, recommended by Decawave themselves [9], is to use the ratio between the energy in the timing pulse (First Peak Power Level or FPPL) and the overall Received Signal Strength (RSS) as a criterion between NLOS and Line-of-Sight (LOS) conditions. Those two diagnostics already allow for good detection, as energy and maximum amplitude of the received signal are significantly more attenuated in NLOS conditions [13]. This simplicity makes this approach widespread, especially as those two diagnostics are provided out of the box by the very popular Decawave transducers.

This last approach can however be improved since FPPL and RSS are correlated, which means that NLOS conditions will cause a dip both in the RSS and in the FPPL. This is pictured in Fig. 1 which shows the channel diagnostics along with the true ranging error as NLOS conditions occur around 120s. Fig. 1 indeed shows a strong correlation between events in the RSS - FPPL delta (in dB) and events in the ranging error, but it also shows how the RSS is affected as well by the NLOS conditions. Furthermore, Fig. 1 shows that the RSS measurement fluctuates more than the FPPL measurement. This is understandable, as the RSS is the

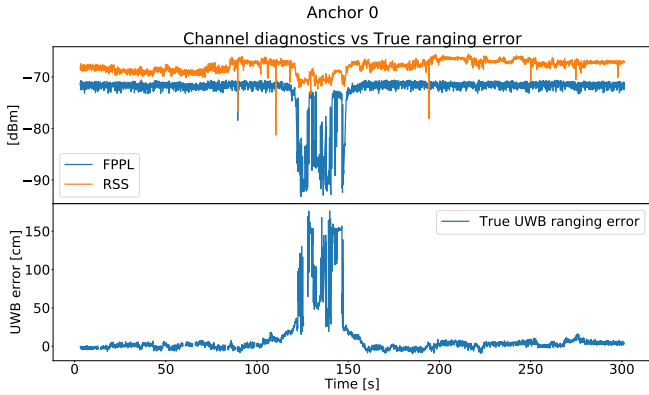


Fig. 1: Effect of NLOS conditions on channel diagnostics and ranging error. Note how the NLOS conditions at 120 s affect both the FPPL measurement and the ranging error, but also the RSS measurement. Note also how the RSS measurement is noisier than the FPPL measurement. In this particular instance, a quantitative correlation could be established between FPPL and ranging error, but in most cases the link between FPPL and ranging error stops at events happening simultaneously.

result of the integration of the signal power over the whole transmission, whereas the FPPL only measures the energy in the timing pulse, when signal-to-noise ratio is maximal. We conclude from those observations that a better indicator could be obtained by working solely with the FPPL and comparing it with an expected value based on previous measurements. This avoids comparing with the RSS, which trends similarly in the presence of NLOS conditions, and allows to only work with the less noisy of the two measurements.

We propose to use the Friis transmission equation (12) from [20] to model the FPPL measured at the receiver, and to then perform hypothesis testing on the measurement against the predicted value.

$$\begin{aligned}
 P_{Rx} = & \underbrace{P_{Tx} + G_{Tx} + G_{Rx}}_{\text{Quasi-static parameters}} \\
 & + \underbrace{10n \log \left(\frac{\lambda}{4\pi r} \right)}_{\text{Propagation}} \\
 & + \text{Noise and other unexpected contributions}
 \end{aligned} \quad (12)$$

In this particular case, P_{Rx} is the FPPL measured at the receiver, P_{Tx} the FPPL at the transmitter and G_{Rx} and G_{Tx} are antenna gains which include directivity and losses. λ is the central wavelength of the UWB channel and r is the distance between the transmitter (the anchor or the tag) and the receiver (the tag or the anchor). The path loss exponent n is set to $n = 2$ (free space propagation) as the goal is to model an ideal LOS channel to then perform hypothesis testing. If the channel is indeed purely LOS, the model fits with the last term reducing to only simple Gaussian noise. Otherwise, the last term must include additional transmission and reflections losses and/or n has to change value to account for multipath propagation [20].

If the antennas are sufficiently omnidirectional, $P_{Tx} + G_{Tx} + G_{Rx}$ can be considered as a near constant contribution to the overall budget with variation due only to temperature and ageing. From there, the proposed method is to run an inner univariate KF independently for each anchor-tag UWB channel to estimate $P_{Tx} + G_{Tx} + G_{Rx}$ of that channel (modelled as a slow random walk) based on FPPL measurements. From this estimation of $P_{Tx} + G_{Tx} + G_{Rx}$, the FPPL is predicted using equation (12) where r is computed based on the estimated position of the robot.

As each inner channel KF is univariate and independent, detecting an outlier can be as simple as checking that the measured FPPL falls in the confidence interval around the predicted FPPL. Measuring an FPPL outside the confidence interval means that the channel does not fit the ideal LOS model, a condition for the corresponding range to be reliable, with sufficient probability. As a consequence, when an FPPL measurement is determined to be an outlier by the associated channel KF, the update step of that UWB channel KF is not done and the range provided by the corresponding anchor is also discarded by the localization EKF.

Depending on the use-case and the available measurements, one of the two proposed outlier detection techniques can be used or both in combination. FPPL-based outlier detection is less computationally intensive than S -based outlier detection and is completely independent from the choice of sensor fusion technique as it only relies on the position estimation and not on the estimation of S . It has however the downsides that it requires the FPPL measurement to be available, that it can only detect types of outliers where there is a correlation between FPPL and ranging error, and that it does not exploit cross-correlations in S to enforce agreement between the selected ranges. In subsection V-B, an experimental comparison is shown of the performance of the two proposed techniques and of their combination.

When both techniques are combined, FPPL-based outlier detection should be performed before S -based outlier detection as it is less computationally intensive. Discarding ranges early makes the matrices in equations (10) and (11) smaller which especially speeds up the inversions.

V. EXPERIMENTAL VALIDATION

The experimental validation of the outlier detection and of the overall sensor fusion is presented in this section.

A. Setup and Datasets

The proposed techniques are used to track a Pioneer P3-DX robot that is driving through a roughly $8 \text{ m} \times 10 \text{ m}$ test area. As pictured in Fig. 2 and Fig. 3, an ArUco marker [21] on the robot is tracked by 8 cameras covering the test area to provide mm-accurate position ground-truth (GT). 8 Decawave TREK1000 [22] devices are used as anchors, with their spatial configuration varying from dataset to dataset. The tag placed on the robot is another Decawave TREK1000. The UWB devices are configured to perform double-sided TWR [16] at 12.5 Hz. ROS [23] is used to control the robot and acquire the encoder, UWB and GT data.




Fig. 2: The P3-DX robot tracked in the experiments



Fig. 3: Experimental setup. Depending on the experiment, the anchors are mounted on the fixed metallic trusses or on tripods further away from the testing area. The cameras providing optical ground-truth are mounted on the trusses and calibrated using ArUco markers at known positions on the floor.

Several datasets were acquired, with varying trajectories, anchor placements and obstructions. Those datasets [15] are provided alongside this work under a CC-BY 4.0 license, and raw data can also be made available on request.

To establish baseline performance in the ideal case, simple trajectories were executed with no obstructions and good anchor placements. Those datasets are pictured in Fig. 4, with the positions of the antennas indicated by the labels 0 to 7. The antennas are hung at a height of around 4 m.

A slightly more challenging dataset was acquired where the robot drove right under some anchors, as depicted in Fig. 5, which resulted in outliers due to the toroidal directivity of the antennas. The result of simple multilateration without outlier detection is also shown in green on Fig. 5 to highlight the parts of the trajectory where outliers occurred. The lower right part of the trajectory shows that anchor 1, provides faulty ranges as the robot drives close to right under it, resulting in simple multilateration giving poor accuracy in that area. The same phenomenon happens, although with a less dramatic effect, around anchors 0 and 5.

Finally, to emulate realistic conditions with complex mixed obstructions, a dataset was acquired with a large vehicle placed in the middle of the test area, as depicted in Fig. 6. The vehicle in the middle of the test area was preventing or disrupting ranging to certain anchors in func-




Fig. 4: Ideal LOS datasets: Ground Truth (GT) trajectory, and trajectory estimated using the proposed sensor fusion techniques (labeled as KF). Labels 0 to 7 on the figure denote the positions of the anchors.

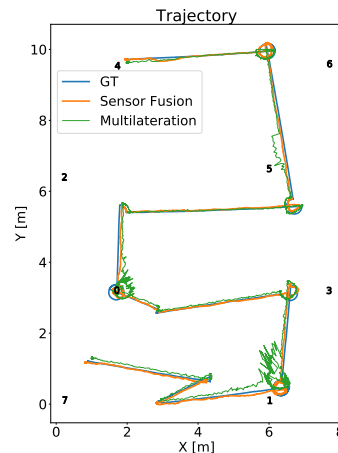


Fig. 5: “Stop & Go” dataset: Ground Truth (GT) trajectory, and trajectories estimated using the proposed sensor fusion techniques and simple multilateration. Labels 0 to 7 on the figure denote the positions of the anchors. Note how the multilateration estimation is affected by UWB outliers as the robot drives under or near anchors 0, 1 and 5. This is due to the toroidal directivity of the anchors.

tion of the position of the robot, due to refraction through glass windows, edge diffraction, transmission through thin metallic sheets and/or complete reflection on thicker metallic elements depending on which part(s) of the vehicle the UWB wave was interacting with.

Again, the result of simple multilateration (in green on Fig. 6) visually shows how much UWB ranging was disrupted by NLOS conditions, and how these translate into inaccurate position estimation if not handled properly.

The proposed sensor fusion techniques were implemented in Python using Scipy [24] and tested on those datasets. The results of those tests are presented below.

B. Performance Metrics

As proposed by [25], we use the cumulative distribution function (CDF) of the 2D localization error to quantify the tracking accuracy.

1) *Baseline Performance in the Ideal LOS Case:* First the performance of the sensor fusion is shown in the purely LOS

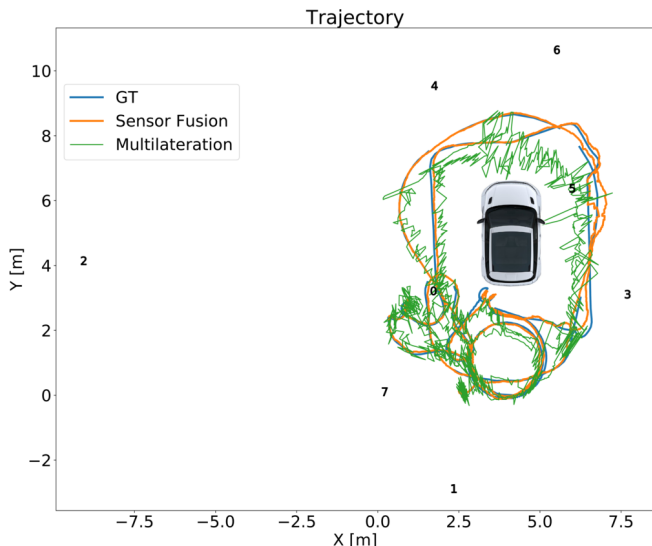


Fig. 6: “Evoque” Dataset: Ground Truth (GT) trajectory, and trajectories estimated using the proposed sensor fusion techniques and simple multilateration. Labels 0 to 7 on the figure denote the positions of the anchors.

cases. Fig. 7 shows the accuracy for the three datasets that




Fig. 7: 2D localization error CDF for the LOS cases

did not contain any outliers, and Table I gives a summary of the corresponding accuracy metrics.

TABLE I: Accuracy metrics for the LOS cases

Dataset	2D localization error		
	Median	Third quartile	90th percentile
“Straight Y”	5.6 cm	7.2 cm	8.4 cm
“Circles 1”	4.5 cm	7.1 cm	9.8 cm
“Circles 2”	4.3 cm	5.6 cm	6.5 cm

The obtained localization performance is on par with or better than what has been previously reported on UWB tracking for this ideal case [2], [14], [17], [25]. Those previous results have been obtained with similar anchor densities i.e. about 8 anchors covering a roughly 100 m^2 area. Note however that we have not found state of the art results that could be compared in a rigorous manner and that use encoders in combination with UWB ranges, which

should arguably give better performance than e.g. UWB and Inertial Measurement Unit (IMU). Note also that contestants did not have access to encoder data at the 2016 IPIN robot tracking (track 4) competition [2], which remains a reference regarding the tracking of indoor vehicles.

Similar performance has also been obtained with datasets from three industrial partners, which further validates our results. Each of these datasets was acquired using a different vehicle and different sensors, at different scales, and in two cases with custom UWB hardware. One was acquired inside a plant of a leading metal manufacturer, one in a large outside container terminal and one at a farming automation specialist. In two cases, anchor density was significantly lower, with 8 anchors covering areas up to 2500 m^2 .

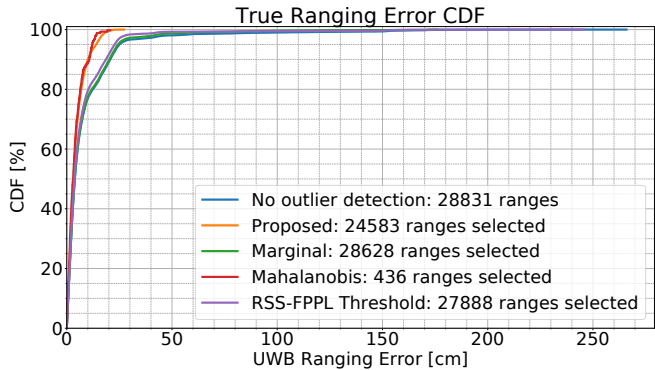
2) *Performance in Realistic Cases:* With the basic sensor fusion validated, the focus is now set on benchmarking the outlier detection techniques and overall accuracy in the more realistic “Stop & Go” and “Evoque” datasets.

To evaluate outlier detection performance, the CDF of the absolute measurement error of the selected ranges is shown per outlier detection method, using the known GT positions to compute the reference ranges. This is depicted in Fig. 8 for the “Stop & Go” dataset and for the “Evoque” dataset.


Fig. 8 shows the improvement obtained with the proposed outlier detection compared to the existing methods presented in section IV. In the “Stop & Go” dataset, using only the marginal distributions does not sufficiently eliminate outliers, and neither does using a simple threshold on the RSS - FPPL delta. As expected, using the Mahalanobis distance as a criterion results in the elimination of many valid ranges. While it eliminates all measurements with an absolute error above 30 cm, it selects only 436 ranges, compared to the 24583 selected by the proposed method, which achieves a very similar CDF and eliminates all measurements with an absolute error above 40 cm.

Similarly, in the “Evoque” dataset, the Mahalanobis distance criterion results in the elimination of all but one measurements, which can be understood as less than 5% of the measurement vectors did not contain an element with an absolute error below 20 cm. Outlier detection based on the marginal distributions approaches the proposed method in this particular case but is still slightly inferior. In general, it appears that from use-case to use-case, one of the existing outlier detection criteria might approach the proposed method but the proposed method remains robust across all the test-cases, including the ones carried out at industrial partner’s sites and with their hardware.

This improved accuracy of the ranges selected to be used by the localization EKF translates into an overall improved localization performance, as depicted in Fig. 9, which compares the localization accuracy of the EKF with the proposed outlier detection method to the accuracy of the same EKF with existing outlier detection methods and to the accuracy of simple multilateration. The accuracy of the EKF with Mahalanobis outlier detection also indicates the quality of the encoder measurements. As nearly all ranges are eliminated, the EKF is essentially performing dead-reckoning.



(a) “Stop & Go” dataset



(b) “Evoque” dataset


Fig. 8: CDF of the absolute UWB ranging error for different outlier detection methods. “Marginal”: measurements further than 3σ away from their expected value are eliminated. “Mahalanobis”: innovation vectors with a Mahalanobis norm higher than 3 are eliminated. “RSS-FPPL Threshold”: ranges where the channel has RSS - FPPL > 10 dB are eliminated, as recommended by Decawave [9]. The proposed outlier detection method allows to limit the true ranging error from sometimes over 1 m to always under 40 cm.

Again, the improvement obtained by applying the techniques proposed in this article is clearly visible. These techniques achieve localization with third quartile of the 2D localization error at or below 10 cm even in the harshest case. Third quartile error below 10 cm was also achieved with the coverage of $50\text{ m} \times 50\text{ m}$ test-area of a running factory.

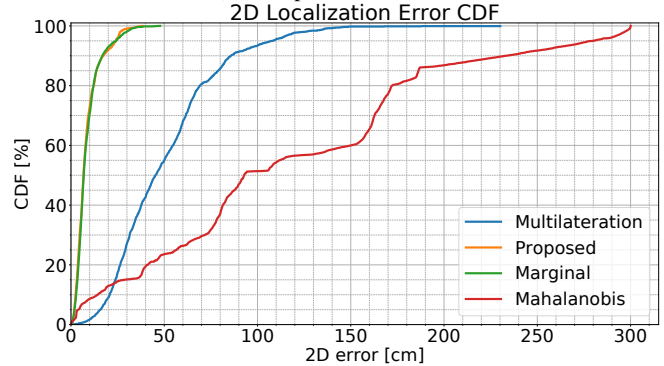
VI. CONCLUSIONS AND FUTURE WORK

In this article we have shown how UWB can be used in conjunction with dead-reckoning in order to achieve tracking with 10 cm third quartile 2D localization error. This was experimentally validated even when complex disruptions of the UWB channel were introduced to mimic a harsh industrial environment. These results are made possible by the application of standard sensor fusion techniques as well as the development of improved outlier detection methods more specific to UWB based localization.

In our “Evoque” dataset, where more than 30% of the measured ranges are outliers with true ranging error from 20 cm up to more than 3.5 m, we have shown experimentally



(a) “Stop & Go” dataset



(b) “Evoque” dataset

Fig. 9: CDF of the 2D localization error for different outlier detection methods, and for simple multilateration

how the proposed outlier detection methods successfully manage to select a subset of 70% of these measurements with errors no greater than 50 cm. This leads to a dramatic improvement of the resulting localization accuracy and we have shown how in the harshest conditions sensor fusion and outlier detection allowed to bring the third quartile of the 2D localization error from over 3 m to 10 cm.

As our approach revolves around improving the processing of the measurement data rather than its acquisition and resulting quality, it is very generic. We have proven this by applying our techniques on datasets collected by three industrial partners in addition to our own, with different vehicles, UWB hardware and overall conditions.

We believe that these results can be improved upon going forward. First, the UWB hardware used during the acquisition of the datasets presented in this article could be upgraded or better modeled. The current antennas [26] do not have perfectly omnidirectional group delay but are modeled as such. There is also reason to believe that the UWB ranging error could be further modeled as a function of the vehicle position. Previous work has shown the multilateration localization error to be seemingly repeatable [14], [17].

Second, the approach proposed in this article could be made more powerful and generic by also taking advantage of data from an IMU. This would require full 6 degrees of freedom pose estimation to correctly estimate the reference frame in which accelerations and rotations are measured as

well as some IMU-specific techniques for the compensation of bias and drift, as has been done before [7], [14]. The advantages would be a possibly improved accuracy, and more importantly redundancy and/or the ability to build a robust, independent localization package that would not rely on vehicle encoder data being available.

Finally, the proposed outlier detection techniques are “hard” in the sense that a measurement is either used with its usual weight, or not used at all. There have been attempts [14] at “soft” outlier detection techniques that weigh the incoming measurements based on the estimated condition of the UWB channel. Somewhat similarly, reference [13] proposes to correct the measured ranges based on features in the CIR of the UWB channel instead of classifying measurements and discarding the NLOS ones. Estimating the parameters of a more accurate but also more complex UWB range measurement model based on the information contained in the CIR and possibly information from other sources such as a map of the environment or channel history would be the most comprehensive solution. However, this remains a (highly dimensional) challenge, and would probably require using a particle filter to handle the non-Gaussian measurement model, which is not always practical on embedded targets.

To make large-scale deployments simpler and UWB a stronger candidate localization solution overall, research should also investigate how to estimate the position of the anchors in the reference frame of interest without resorting to expensive field surveys. This can be made possible by formulating the estimation of the whole trajectory and of the anchor positions from the UWB ranges as a single optimization problem. If other sensors are used in this estimation, it becomes a Simultaneous Localization and Mapping (SLAM) problem, a topic that is also receiving attention. In either case, the outlier detection methods presented in this article remain of interest.

REFERENCES

- [1] D. Lymberopoulos and J. Liu, “The Microsoft Indoor Localization Competition: Experiences and Lessons Learned,” *IEEE Signal Processing Magazine*, vol. 34, no. 5, pp. 125–140, Sep. 2017.
- [2] M. C. Pérez-Rubio, C. Losada-Gutiérrez, F. Espinosa, J. Macias-Guarasa, J. Tiemann, F. Eckermann, C. Wietfeld, M. Katkov, S. Huba, J. Ureña, J. M. Villadangos, D. Gualda, E. Díaz, R. Nieto, E. Santiso, P. del Portillo, and M. Martínez, “A realistic evaluation of indoor robot position tracking systems: The IPIN 2016 competition experience,” *Measurement*, vol. 135, pp. 151–162, Mar. 2019. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0263224118310728>
- [3] A. R. Jimenez Ruiz and F. Seco Granja, “Comparing Ubisense, BeSpooon, and DecaWave UWB Location Systems: Indoor Performance Analysis,” *IEEE Transactions on Instrumentation and Measurement*, vol. 66, no. 8, pp. 2106–2117, Aug. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7891540/>
- [4] “Pozyx: Accurate positioning,” <https://pozyx.io>, accessed: 2019-03-05.
- [5] A. Alarifi, A. Al-Salman, M. Alsaleh, A. Alnafessah, S. Al-Hadhrani, M. A. Al-Amr, and H. S. Al-Khalifa, “Ultra Wideband Indoor Positioning Technologies: Analysis and Recent Advances,” *Sensors (Basel, Switzerland)*, vol. 16, no. 5, May 2016. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4883398/>
- [6] M. Vossiek, L. Wiebking, P. Gulden, J. Weighardt, and C. Hoffmann, “Wireless local positioning - concepts, solutions, applications,” in *Radio and Wireless Conference, 2003. RAWCON '03. Proceedings*, Aug. 2003, pp. 219–224.
- [7] J. D. Hol, F. Dijkstra, H. Luinge, and T. B. Schon, “Tightly coupled UWB/IMU pose estimation,” in *2009 IEEE International Conference on Ultra-Wideband*, Sep. 2009, pp. 688–692.
- [8] B. Hepp, T. Nägeli, and O. Hilliges, “Omni-directional person tracking on a flying robot using occlusion-robust ultra-wideband signals,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2016, pp. 189–194.
- [9] “Dw1000 metrics for estimation of non line of sight operating conditions,” https://www.decawave.com/wp-content/uploads/2018/10/APS006_Part-3-DW1000-Diagnostics-for-NLOS-Channels_v1.1.pdf, Decawave, Tech. Rep. APS013 Part 3 Application Note, 2016.
- [10] J. Khodjaev, Y. Park, and A. S. Malik, “Survey of NLOS identification and error mitigation problems in UWB-based positioning algorithms for dense environments,” *annals des télécommunications*, vol. 65, no. 5-6, pp. 301–311, Jun. 2010. [Online]. Available: <https://link.springer.com/article/10.1007/s12243-009-0124-z>
- [11] D. Neiryck, M. O’Duinn, and C. McElroy, “Characterisation of the NLOS Performance of an IEEE 802.15.4a Receiver,” Decawave, Tech. Rep., 2015.
- [12] S. Marano, W. M. Gifford, H. Wymeersch, and M. Z. Win, “NLOS identification and mitigation for localization based on UWB experimental data,” *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 7, pp. 1026–1035, Sep. 2010.
- [13] H. Wymeersch, S. Marano, W. M. Gifford, and M. Z. Win, “A Machine Learning Approach to Ranging Error Mitigation for UWB Localization,” *IEEE Transactions on Communications*, vol. 60, no. 6, pp. 1719–1728, Jun. 2012.
- [14] J. Tiemann, L. Koring, P. Gorczak, and C. Wietfeld, “Improving the Robustness of Control-Grade Ultra-Wideband Localization,” *IFAC-PapersOnLine*, vol. 51, no. 10, pp. 103–108, Jan. 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2405896318305640>
- [15] N. Dwek, “Locator Lommel datasets,” Apr. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.2621414>
- [16] “The implementation of two-way ranging with the dw1000,” https://www.decawave.com/wp-content/uploads/2018/08/aps013_dw1000_and_two_way_ranging_v2.2.pdf, Decawave, Tech. Rep. APS013 Application Note, 2015.
- [17] J. Tiemann, F. Eckermann, and C. Wietfeld, “ATLAS - an open-source TDOA-based Ultra-wideband localization system,” in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Oct. 2016, pp. 1–6.
- [18] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, “CasADi – A software framework for nonlinear optimization and optimal control,” *Mathematical Programming Computation*, In Press, 2018.
- [19] M. L. Eaton, *Multivariate statistics: a vector space approach*, ser. Wiley series in probability and mathematical statistics. New York: Wiley, 1983.
- [20] R. C. Johnson and H. Jasik, Eds., *Antenna engineering handbook*, 2nd ed. New York, NY: McGraw-Hill, 1984, oCLC: 10489292.
- [21] S. Garrido-Jurado, R. Muñoz-Salinas, F. Madrid-Cuevas, and M. Marín-Jiménez, “Automatic generation and detection of highly reliable fiducial markers under occlusion,” *Pattern Recognition*, vol. 47, no. 6, p. 2280–2292, Jun 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.patcog.2014.01.005>
- [22] “Decawave TREK1000: Two-Way-Ranging (TWR) RTLS IC Evaluation Kit,” <https://www.decawave.com/product/trek1000-evaluation-kit/>, accessed: 2019-03-11.
- [23] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, “Ros: an open-source robot operating system,” in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA) Workshop on Open Source Robotics*, Kobe, Japan, May 2009.
- [24] E. Jones, T. Oliphant, P. Peterson *et al.*, “SciPy: Open source scientific tools for Python,” <http://www.scipy.org/>, 2001–, [Online; accessed 2019-03-14].
- [25] F. Potortù, S. Park, A. R. Jiménez Ruiz, P. Barsocchi, M. Girolami, A. Crivello, S. Y. Lee, J. H. Lim, J. Torres-Sospedra, F. Seco, R. Montoliu, G. M. Mendoza-Silva, M. D. C. Pérez Rubio, C. Losada-Gutiérrez, F. Espinosa, and J. Macias-Guarasa, “Comparing the Performance of Indoor Localization Systems through the EvAAL Framework,” *Sensors (Basel, Switzerland)*, vol. 17, no. 10, Oct. 2017.
- [26] “DecaWave reference planar monopole antenna,” Decawave, Tech. Rep., 2014.