Robot Localization via Odometry-assisted Ultra-wideband Ranging with Stochastic Guarantees

Valerio Magnago[1], Pablo Corbalán[1], Gian Pietro Picco[1], Luigi Palopoli[1], Daniele Fontanelli[2]

***Abstract*— We consider the problem of accurate and high-rate self-localization for a mobile robot. We *adaptively* combine the speed information acquired by proprioceptive sensors with *intermittent* positioning samples acquired via ultra-wideband (UWB) radios. These are triggered only if and when needed to reduce the positioning uncertainty, itself modeled by a probabilistic cost function. Our formulation is agnostic w.r.t. the source of uncertainty and enables an intuitive specification of user navigation requirements along with stochastic guarantees on the system operation. Experimental results in simulation and with a real platform show that our approach *i)* meets these guarantees in practice *ii)* achieves the same accuracy of a fixed periodic sampling but with significantly higher scalability and lower energy consumption *iii)* is resilient to errors in UWB estimates, enabling the use of low-accuracy ranging schemes which further improve these two performance metrics.**

I. Introduction

The accurate estimation of the position of a robot is key for its correct navigation [1]. In outdoor scenarios, absolute localization is provided by GPS, often complemented by a magnetometer-based compass. However, indoor and other GPS-denied scenarios demand alternative localization technologies and approaches. Several have been developed to date, with different tradeoffs concerning, e.g., accuracy, infrastructure requirements and cost, and energy consumption.

Why UWB in Robotics? Ultra-wideband (UWB) radios are rapidly gaining traction, thanks to a new generation of smaller and cheaper transceivers, e.g., the DecaWave DW1000 [2], offering decimeter-level accuracy. In comparison with other techniques, e.g., LiDAR [3] or vision on natural [4] and artificial [5], [6] landmarks, UWB enables both distance estimation (ranging) *and* communication among devices within the *same* radio chip. This is a major asset in system design, especially in robotics applications where concerns about weight, form factor, and complexity dominate, e.g., drones and space exploration. Other radio-based technologies, for example WiFi [7], offer the same asset but with an order of magnitude decrease in accuracy and significantly higher energy consumption.

Motivation. Nevertheless, the control of a robot requires the acquisition of positioning information at relatively high frequency, typically in the range 10–50 Hz. In this respect, the application of UWB to robotics faces two key challenges. First, positioning entails several message exchanges with nearby anchors. The commonly-used two-way ranging (TWR) in the IEEE 802.15.4 standard [8] requires 2 packets per estimate in its simplest form (single-sided TWR, SS-TWR); 4 packets are often used (double-sided TWR, DS-TWR) to improve accuracy. A positioning estimate from N anchors therefore requires $2N$ or $4N$ messages, respectively. As these occur on the shared wireless medium, the *scalability* to scenarios with multiple robots is intrinsically at odds with the achievable update rate, i.e., tracking accuracy [9]. Further, although UWB has relatively low *energy consumption* w.r.t. WiFi, it still consumes hundreds of mW, often prohibitive for small and lightweight robots where the energy budget matters [10]. Based on these considerations, reducing the UWB sampling rate without compromising accuracy is imperative for the practical application of this localization technology.

Approach. This paper tackles both challenges by combining the *speed* information acquired by low-cost, low-accuracy proprioceptive sensors, e.g., tachometers, IMU and/or odometers, with the *distance* information obtained by high-cost, high-accuracy UWB devices. Ego-motion estimators are widely used in robotics and can usually be executed at high frequency, but unfortunately suffer from drift phenomena induced by dead-reckoning.

The key idea of this paper is to use UWB measurements *intermittently and adaptively*, by triggering them only if and when needed to keep the positioning uncertainty under control. We model uncertainty directly by estimating the a-posteriori probability density function (PDF) based on an Unscented Kalman Filter (UKF). This enables us to provide *stochastic guarantees* and to empower the robot user with the ability to *intuitively specify application requirements* in terms of an accuracy threshold combined with a confidence level, as in "a positioning error $\leq$20 cm in $\geq$90% of the samples". After concisely illustrating the assumptions and problem formulation (Section II), we outline our intermittent localization approach (Section III). We then provide a detailed model for a specific robotic setup and UWB system (Section IV). The latter is used to evaluate the performance of our approach via simulation (Section V) and then validate these results with real-world experiments (Section VI) exploiting the ground truth offered by a high-precision optical system.

Related Work and Contribution. The combination of ego-motion estimators and localization sensors is relatively common in robotics [11] and typically achieved by using the localization sensor, UWB in our case, at a rate lower than needed by control, leveraging the ego-motion estimates in be-

[1]Authors are with the Dipartimento di Ingegneria e Scienza dell'Informazione (DISI), University of Trento, Italy first.last@unitn.it.

[2]Author is with the Dipartimento di Ingegneria Industriale (DII), University of Trento, Italy daniele.fontanelli@unitn.it.

The research in this paper has received funding from the European Union via the EIT Digital AWARD and D-TWIN projects, and from the Italian government via the NG-UWB project (MIUR PRIN 2017).

tween samples. However, selecting the appropriate sampling period is tricky if robots do not move at constant velocity. The sampling rate must be dimensioned for the worst-case, yielding an inefficient average performance; further, the system may still fail to meet the required guarantees if the worst-case is incorrectly estimated. In contrast, our approach *adapts* the sampling rate to the current uncertainty, *without requiring a priori estimation and regardless of its nature*. By triggering UWB samples only when needed we achieve the same accuracy at a fraction of the cost (<50% in many cases, Section V–VI). Moreover, as the sampling rate is not fixed and bound a priori, we can cope with unexpected changes due to environmental factors, e.g., the presence of uneven or slippery terrain, or the malfunctioning of the odometer.

To the best of our knowledge, the only adaptive combination of odometry and UWB is presented in [12], where UWB samples are triggered based on the maximum eigenvalue of the covariance matrix estimated via an Extended Kalman Filter (EKF). In comparison, our approach offers significant advantages: *i)* formulating user requirements in terms of error threshold and confidence level rather than maximum eigenvalues is remarkably more intuitive, simplifying the robot configuration and operation, and *ii)* the technique in [12] is unnecessarily conservative, triggering UWB samples when not required and ultimately significantly reducing the potential improvements, as we show quantitatively in Section V. Moreover, the work in [12] uses DS-TWR to improve the accuracy of UWB samples and hence of the overall localization. Instead, we show experimentally that in our approach *uncertainty is dominated by the odometry error* (Section VI). This is a key finding, enabling the use of the lower-accuracy SS-TWR scheme—further slashing in half the bandwidth and energy demands—without detriment to localization accuracy.

The paper ends with brief concluding remarks and an outlook on future work (Section VII).

II. Assumptions and Problem Formulation

We begin by discussing our assumptions about the robot model and by formalizing the problem we address.

A. Model

The fixed, right-handed reference frame for platform localization is referred to as $\langle W \rangle = \{O_w, X_w, Y_w, Z_w\}$, as shown in Fig. 1. The robotic vehicle is regarded as a rigid body $\mathcal{B}$ moving on the plane $X_w \times Y_w$. Let t_s denote the sampling period, assumed common for all onboard sensors. The generalized coordinates of the robot at time kt_s are denoted by $p_k = [x_k, y_k, \theta_k]^T$, with (x_k, y_k) being the coordinates of the origin of frame $\langle B \rangle = \{O_b, X_b, Y_b, Z_b\}$ attached to the robot, and with θ_k representing the angle between X_b and X_w. The kinematic model of a generic drift–less, input–affine robot can be written as:

$$\begin{cases} p_{k+1} = p_k + G_k(p_k, u_k + \epsilon_k) \\ z_k = h(p_k) + \eta_k \end{cases} \quad (1)$$

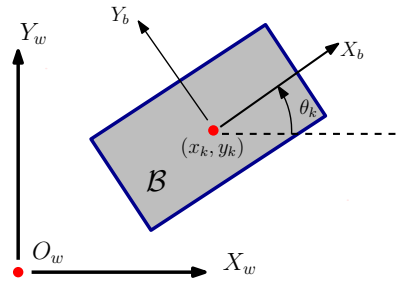

Fig. 1. Platform model represented as a rigid body $\mathcal{B}$ moving on the $X_w \times Y_w$ plane with an attached reference frame $\langle B \rangle$.

where p_k is the robot state, u_k is the input vector, assumed piecewise constant between $(k-1)t_s$ and kt_s, ϵ_k is the additive zero-mean uncertainty term affecting the input quantities and $G_k(\cdot)$ is the input vector function. At any sampling period, the robot samples its internal sensors (e.g., encoders, IMU), but the estimate of the state is affected by the dead-reckoning effect modeled in the term ϵ_k. To prevent the drift of the estimate, the robot is also allowed to collect a measurement of a quantity z_k from an exteroceptive sensor, the UWB localization system in our case, every t_s time units. The quantity z_k is related to the state through the non-linear function $h(\cdot)$. We assume that *i)* the measurement z_k is affected by a noise term η_k *ii)* $h(p_k)$ ensures non-linear observability of the robot state p_k in $\langle W \rangle$, and *iii)* each measurement of z_k has a cost $c(z_k)$. For UWB localization, the cost term accounts for the communication bandwidth used for every measurement, which reduces the number of robots that can use the system at the same time, and for the energy consumed in transmitting and receiving packets, which reduces the robot autonomy.

B. Problem Formulation

Our problem can be formalized in the following terms. Given a robot with kinematic model (1), produce an estimation $\hat{p}_k$ for the state such that: *i)* the cost $c(z_k)$ is minimized, and *ii)* the uncertainty is upper-bounded by a threshold Λ_k with a confidence ψ_k. The accuracy is modeled by a non-negative performance function $H_k(\cdot)$ of the localization error $e_k = \hat{p}_k - p_k$. The performance function $H_k(\cdot)$ is generally multi-dimensional: $H_k : \mathbb{R}^2 \times \mathbb{S} \to \mathbb{R}^n$, where $\mathbb{R}^2$ is the dimension of the state space, $\mathbb{S}$ the space of angles in the plane, and n the number of constraints required. The confidence is defined as a probability, i.e., $\psi_k \in [0, 1]$. Hence, the accuracy requirement is analytically expressed as:

$$\Pr\left[H_k(e_k) < \Lambda_k\right] \geq \psi_k, \forall k. \quad (2)$$

The performance function $H_k(e_k)$ is time-varying, as it depends on changing conditions of the application scenario. For instance, if the robot is moving along a corridor oriented along X_w, it is reasonable to upper-bound the error along Y_w, to avoid collisions with the walls, and the orientation θ_k, to ensure that the robot moves along the chosen direction. A

possible performance function could be:

$$H_k(e_k) = \begin{bmatrix} |\hat{y}_k - y_k| \\ |\hat{\theta}_k - \theta_k| \end{bmatrix} \leq \begin{bmatrix} \lambda_{y_k} \\ \lambda_{\theta_k} \end{bmatrix} = \Lambda_k, \tag{3}$$

with confidence ψ_k. In other cases, we may want to limit the position error on the plane, e.g., to safely enter a room through the door. We could model this by:

$$H_k(e_k) = \sqrt{(\hat{x}_k - x_k)^2 + (\hat{y}_k - y_k)^2} \leq \lambda_k = \Lambda_k. \tag{4}$$

The possible choices of Λ_k and ψ_k are evidently limited by the intensity of the noise η_k as per (1).

III. Approach

Let $f_k(e)$ denote the PDF of the estimation error $e_k = \hat{p}_k - p_k$ at timestamp kt_s. We can rewrite (2) as the Riemann integral:

$$\int_{\mathcal{E}_k} f_k(e)de > \psi_k, \forall k, \tag{5}$$

over the state subspace:

$$\mathcal{E}_k = \left\{ e \in \mathbb{R}^2 \times \mathbb{S} : H_k(e) < \Lambda_k \right\}. \tag{6}$$

Importantly, (5) can be used as a trigger event for the measurement of $z_k = h(p_k)$. As long as $\int_{\mathcal{E}_k} f_k(e)de > \psi_k$, the degradation is not sufficient to violate the accuracy constraint, so no external measurement is needed. However, in the absence of external measurements the dead-reckoning effect "inflates" the probability distribution of the error, with $f_k(e) \to 0$ and, since $\mathcal{E}_k$ is a closed set, $\int_{\mathcal{E}_k} f_k(e)de \to 0$. Thereby, after a sufficient number of steps the threshold ψ_k is approached. When this event occurs, the z_k measurement is triggered to reset the error to acceptable levels. This policy evidently ensures that the accuracy constraint is never violated and limits remarkably the number of expensive queries to the UWB localization system.

A different solution proposed in the literature [12] considers the maximum eigenvalue of the covariance ellipsoid as a trigger for the sampling event of z_k. A first serious problem with this approach is that it is difficult to establish a link between the switching logic, the application requirements and the environment constraints, which in our approach is quite naturally captured by the triple $(H_k(e_k), \lambda_k, \phi_k)$. Further, as shown in the next section, the use of the ellipsoid could lead to an overly conservative triggering of measurements.

IV. A Detailed Model for a Practical Case

We now apply our approach to a case of practical interest, which is the subject of our evaluation via simulation (Section V) and real-world validation (Section VI).

For the robot kinematics, we assume a unicycle model. The robot planar coordinates (x_k, y_k) correspond to the midpoint of the traction wheels, assumed to be the origin O_b of the body frame $\langle B \rangle$. The Y_b axis points through the left wheel, while the X_b axis is oriented in the forward direction, as shown in Fig. 2 for the robot used for the experiments. Assuming that relative encoders are attached to each wheel shaft, measuring respectively the left and right wheel angular displacements δ_k^l and δ_k^r in the time interval $[kt_s, (k+1)t_s]$,

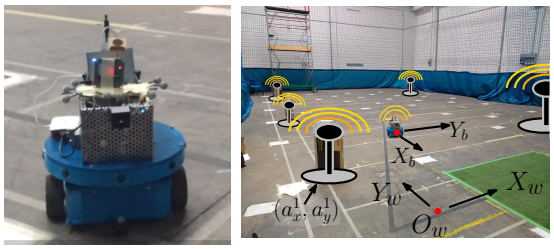

Fig. 2. Unicycle-like vehicle adopted in the experiment and representation of the reference system. UWB anchors are depicted as well.

we can compute the linear motion of the left or right wheel as $L_k = \frac{\phi_l}{2}\delta_k^l$ and $R_k = \frac{\phi_r}{2}\delta_k^r$, being ϕ_l and ϕ_r the wheels radii. With reference to the generic model (1) and assuming for simplicity that in t_s the speed of the wheels is constant, we can safely assume that the inputs are $u_k = [L_k, R_k]^T$. As a consequence, the non-linear input vector is given by

$$G_k = \begin{bmatrix} r_k \left(\sin\left(\theta_{k-1}\right) - \sin\left(\theta_{k-1} + \gamma_k\right)\right) \\ r_k \left(\cos\left(\theta_{k-1} + \gamma_k\right) - \cos\left(\theta_{k-1}\right)\right) \\ \gamma_k \end{bmatrix},$$

where $r_k = \frac{R_k + L_k}{2} + \frac{\epsilon_{R_k} + \epsilon_{L_k}}{2}$ and $\gamma_k = \frac{R_k - L_k}{b} + \frac{\epsilon_{R_k} - \epsilon_{L_k}}{b}$. As reported in [13], the additive white Gaussian noise covariance matrix E_k of $\epsilon_k = [\epsilon_{L_k}, \epsilon_{R_k}]^T$ can be computed as:

$$E_k = \begin{bmatrix} k_L^2 |L_k| & 0 \\ 0 & k_R^2 |R_k| \end{bmatrix}, \tag{7}$$

with ϵ_{L_k} and ϵ_{R_k} being uncorrelated.

A. Measurement System

As a global localization system, we use UWB sensors. This system is capable of carrying out ranging measurements from a set of fixed anchors. Let us assume to have m anchors, each with known plane position (a_x^i, a_y^i), with $i = 1, \ldots, m$; an actual deployment of UWB anchors is shown in Fig. 2. Therefore, the measure z_k is composed by the distance measure coming from m different anchors, i.e., $z_k = [z_k^1, z_k^2, ..., z_k^m]$, where

$$z_k^i = \sqrt{\left(x_k - a_x^i\right)^2 + \left(y_k - a_y^i\right)^2}.$$

The zero-mean, white Gaussian noise η_k affecting each measurement is described through the covariance matrix:

$$N_k = \text{diag}\left\{\nu_k^1, \nu_k^2, \ldots, \nu_k^n\right\}, \tag{8}$$

where $\nu_k^i \in \mathbb{R}_{\geq 0}$ is the variance of the i-th ranging measure and the noise terms of the anchors are assumed uncorrelated.

B. Stochastic Guarantees

One widely adopted solution for collecting both a localization estimate $\hat{p}_k$ and the associated pdf of the error $f_k(e)$ is to resort to Bayesian filters, which are able to coherently fuse the prior on the odometry with the triggered measurement from the UWB sensing system. Due to the non-linear nature of the model and of the measures, we select the Unscented Kalman Filter (UKF) [14]. This enables us to capture the posterior mean $\hat{p}_k$ and covariance Σ_k up to

TABLE I

ANCHOR POSITIONS EXPRESSED IN $\langle W \rangle$.

	Anchor i				
	1	2	3	4	5
$\mathbf{a}^{\mathbf{i}}_{\mathbf{x}}$ **[m]**	1.76	-7.2	-7.2	1.9	-1.9
$\mathbf{a}^{\mathbf{i}}_{\mathbf{y}}$ **[m]**	-5.3	-5.3	1.9	1.9	5.3

the 3^{rd} order of the Taylor series expansion for the non-linearities, modeled as a Gaussian. Without loss of generality, we adopt as performance function $H_k(e)$ the position error defined in (4), hence deriving for (6) the following

$$\mathcal{E}_k = \left\{ e \in \mathbb{R}^2 \times \mathbb{S} : \sqrt{(\hat{x}_k - x_k)^2 + (\hat{y}_k - y_k)^2} < \lambda \right\}, \tag{9}$$

where λ is assumed here to be constant. $\mathcal{E}_k$ is a cylinder contained in the state space $\mathbb{R}^2 \times \mathbb{S}$, and it can be equivalently represented with $\mathcal{A}_k \times \mathbb{S}$, being $\mathcal{A}_k$ the base of the cylinder in the $X_w \times Y_w$ plane. (5) can be written as:

$$\int_{\mathcal{E}_k} \frac{1}{\sqrt{2\pi d_{p_k}}} e^{-\frac{(p-\hat{p}_k)^T \Sigma_k^{-1} (p-\hat{p}_k)}{2}} dp, \tag{10}$$

where $d_{p_k} = \det(\Sigma_k)$. By marginalizing along θ and defining as Σ_k^{xy} the top 2×2 matrix of Σ_k, (10) can be simplified as:

$$\frac{1}{\sqrt{2\pi d_{xy_k}}} \int_{\mathcal{A}_k} e^{-\frac{(p^{xy}-\hat{p}_k^{xy})^T \Sigma_k^{xy-1} (p^{xy}-p_k^{xy})}{2}} dp^{xy}, \tag{11}$$

where p^{xy} and $\hat{p}_k^{xy}$ are the entries related to the cartesian positions of p and $\hat{p}_k$, respectively, while $d_{xy_k} = \det(\Sigma_k^{xy})$. Expressing the previous integral in cylindric coordinates where $v(\psi) = [\cos(\psi), \sin(\psi)]^T$ is the unit vector indicating the integration direction, (11) can be rewritten as

$$\frac{1}{\sqrt{2\pi d_{xy_k}}} \int_0^{2\pi} \int_0^{\lambda} \rho e^{-\frac{\rho v(\psi)^T \Sigma_k^{xy-1} \rho v(\psi)}{2}} d\rho \, d\psi. \tag{12}$$

The integral in ρ can be solved analytically leading to:

$$\frac{1}{\sqrt{2\pi d_{xy_k}}} \int_0^{2\pi} \frac{e^{\lambda^2 C(\psi)} - 1}{2C(\psi)} d\psi \tag{13}$$

where $C(\psi) = -v(\psi)^T \Sigma_k^{xy-1} v(\psi)/2$. Very efficient ad-hoc numeric solutions can be designed for this integral within the desired level of accuracy.

V. SIMULATION-BASED EVALUATION

To evaluate the proposed approach, we developed a simulator by assuming that the model noise of the odometer ϵ_k and the UWB measurement noises η_k are generated by white, zero-mean Gaussian stochastic processes, and by imposing parameters mirroring the actual experimental platform described in Section VI. The covariance matrix of ϵ_k in (7) is computed with $k_L = k_R = 0.004$ m$^{1/2}$, while for the covariance matrix of the UWB noise η_k in (8) we have $\nu_k^i = 10^{-2}$ m^2, $\forall i = 1, \ldots, m$, with $m = 5$. The anchors are deployed in the environment as reported in Table I.

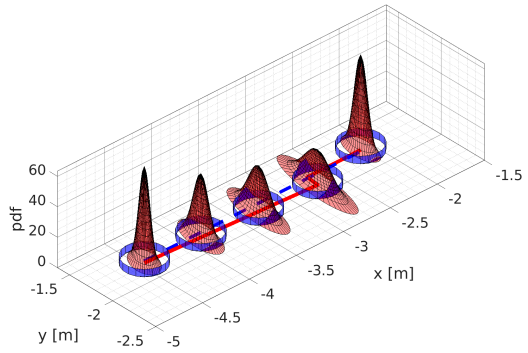

Fig. 3. Time evolution of the marginal PDF of $f_k(e)$ for a robot moving along a straight line parallel to X_w. Both, the actual (dashed line) and estimated (solid line) trajectories are reported. At the fourth depicted position, the integral in (13) exceeds the given confidence ψ due to dead-reckoning, hence the UWB system is triggered and the PDF narrows down (fifth depicted PDF).

We first impose a straight line trajectory originating from position $(-4.5, -1.8)$ m and moving along increasing coordinates of the X_w axis (Fig. 3) with $\lambda = 0.2$ m for (4) and $\psi = 0.75$ for (5), both constant in time. In Fig. 3 we see that along the actual motion (dashed line) the PDF of the error marginalized w.r.t. $\mathbb{S}$ increases due to dead-reckoning, and the estimated trajectory (solid line) therefore deviates from the actual path. In the same figure, the base of the cylinders $\mathcal{A}_k$ are depicted as well (blue circles). When the robot reaches the fourth depicted position, the integral in (13) goes below the confidence ψ; a reading from the UWB system is therefore triggered, causing the PDF to narrow down (last depicted PDF in Fig. 3).

Fig. 4 shows the cumulative distribution function (CDF) of the localization error as a function of the user-defined confidence ψ and threshold λ (Fig. 4a) as well as the impact of the dead-reckoning ϵ_k and UWB ranging uncertainty η_k (Fig. 4b). Note that the CDFs reported are computed by collecting the error values at their maximum right before triggering an UWB measurement, i.e., in the worst-case scenario. Our technique easily adapts to changing performance metrics (ψ, λ) or increasing the odometry or ranging uncertainty (ϵ_k, η_k), tightly satisfying the user-defined bound. Increasing the uncertainty, however, has a remarkable effect on the frequency at which the UWB system is triggered (Table II). The UWB sampling frequency is dominated by the odometry uncertainty, while the UWB ranging uncertainty plays a significantly minor role. This allows us to use more energy- and time-efficient (although less accurate) ranging schemes like SS-TWR instead of DS-TWR as in [12], saving energy and increasing scalability w.r.t. the state of the art.

Besides the discussed intrinsic robustness of our approach, we now show how the algorithm effectively reduces the communication bandwidth utilization w.r.t. a periodic sampling using a trajectory with time-varying velocity. The rationale is to further highlight the adaptability of the solution at hand: when the trajectories are regular and the velocity is almost

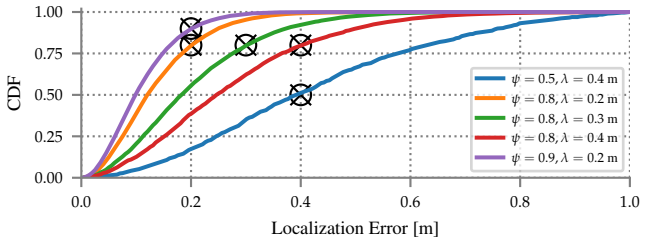

(a) Different (ψ, λ)

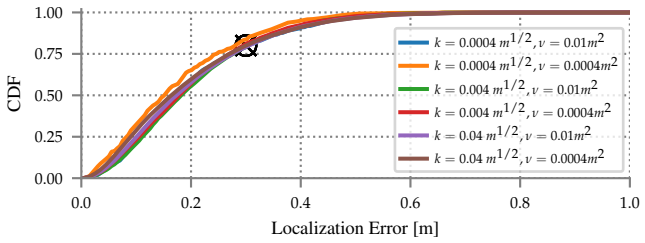

(b) Different (k, ν)

Fig. 4. Localization accuracy under different (ψ, λ, k, ν). Our solution adapts the CDFs (top) to the user-defined threshold λ and confidence ψ, highlighted with crossed circles on each plot. Changing the encoder or ranging uncertainty for $\lambda = 0.3$ m and $\psi = 0.8$ (bottom) barely affects the resulting localization error.

TABLE II

FREQUENCY OF UWB TRIGGERING AS A FUNCTION OF THE UNCERTAINTY ϵ_k AND η_k.

	$\mathbf{k_L = k_R [m^{1/2}]}$		
$\nu[\mathbf{m^2}]$	**0.0004**	**0.004**	**0.04**
0.0004	0.02 Hz	0.07 Hz	0.52 Hz
0.01	0.02 Hz	0.08 Hz	0.6 Hz

constant (see odometer model (7) and its effect in Table II), a low-frequency periodic sampling is sufficient; when instead the trajectories become more challenging, a high-frequency sampling is needed. Therefore, a periodic sampling approach should be fine-tuned on the worst-case scenario, which leads to an overuse of the shared UWB system. The adaptability of the proposed solution, instead, automatically enables the use of the UWB system only when strictly needed. The example we use in our analysis is taken from an industrial warehouse application and assumes that the robot moves on a planned path with a bound of 0.5 m/s^2 in the forward and lateral accelerations, a maximum speed of 1.5 m/s, and by starting and stopping with zero velocity. The resulting velocity profile (Fig. 5) is composed of a region in which the robot is moving straight at full speed in the middle of two turning sections. The obtained pseudo-frequency distribution of UWB measurements, i.e., the inverse of the local sampling periods, is depicted in Fig. 6. The frequency distribution is in the interval $[0.6, 1.5]$ Hz with a mean value of 1.08 Hz. Lower frequencies are relative to the slow-paced region of the velocity profile, whereas high frequencies are due to the fast-paced region (grey area in Fig. 5). Fig. 7 shows the worst-case localization error CDFs for this scenario obtained

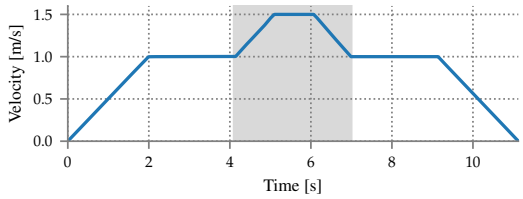

Fig. 5. Velocity profile.

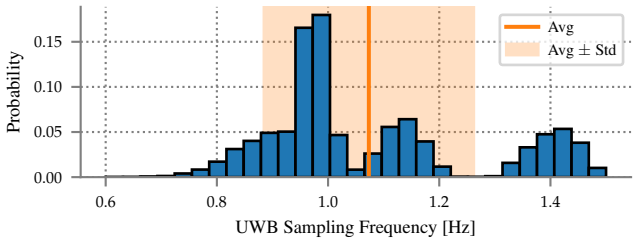

Fig. 6. UWB sampling pseudo-frequencies distribution.

with our adaptive mechanism w.r.t. periodic UWB sampling at fixed rates. The periodic rates are respectively selected based on the minimum, mean, and maximum frequency in Fig. 6. When the robot moves at 1 m/s (left), all the CDFs comply with the user-defined bound ($\psi = 0.8, \lambda = 0.2$ m). However, increasing the velocity to 1.5 m/s (right) makes the periodic approaches with $f_s < 1.5$ Hz unable to cope with the fast speed of the robot. Our technique, instead, adapts its sampling rate as needed and satisfies the requirements irrespective of the actual working conditions.

Finally, we compare our proposed technique against the one in [12], which is based on the eigenstructure of the localization error covariance. In particular, we take the greater axis of the ellipse which includes the ψ probability to scale the greater eigenvalue of the (x, y) position uncertainty matrix obtained with Principal Component Analysis. The conservativeness of the methods relying on the eigenvalues is reported in Fig. 8, where the CDFs of the two methods are compared. It is evident that our method follows the requirements more tightly in both tested configurations. In the first case, $\lambda = 0.25$ m and $\psi = 0.8$, our technique requires a mean UWB triggering frequency of 0.16 Hz instead of 0.25 Hz; in the second case, $\lambda = 0.2$ m and $\psi = 0.9$, the proposed adaptive system frequency is 0.38 Hz, while the eigenvalue technique yields 0.50 Hz. Therefore, our technique can potentially reduce the required UWB sampling rate by a 24–36% factor w.r.t. state-of-the-art techniques based on the eigenvalue.

VI. EXPERIMENTAL RESULTS

To validate our technique, we ran an experimental campaign in an environment equipped with 14 OptiTrack cameras [15] covering a 12×8 m^2 area and providing mm-level localization at 125 Hz. The unicycle-like wheeled robot used for the experiments and the testing arena equipped with 5 UWB anchors are reported in Fig. 2. The robot had a UWB tag in its center position and measured its

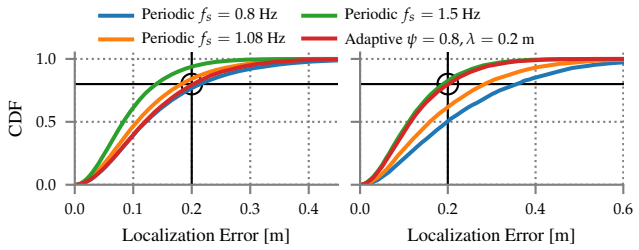

Fig. 7. Comparison between periodic and adaptive UWB sampling with the robot moving at 1 m/s (left) and 1.5 m/s (right). Our adaptive mechanism satisfies the user requirements despite changes in the robot velocity.

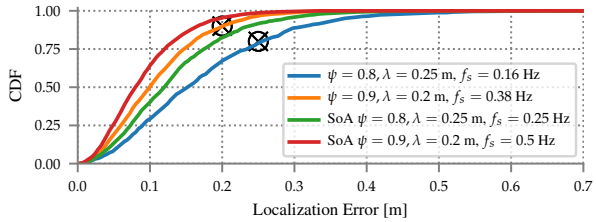

Fig. 8. Comparison with the state of the art (SoA) for different (ψ, λ). SoA techniques based on the eigenvalue are conservative, over provisioning UWB measurements, increasing consumption, and decreasing scalability.

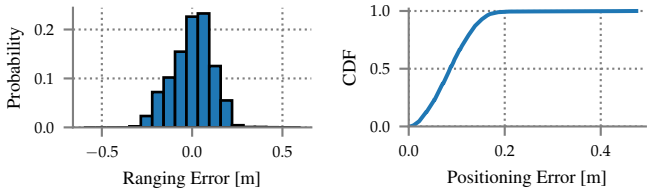

Fig. 9. UWB ranging (left) and positioning (right) error.

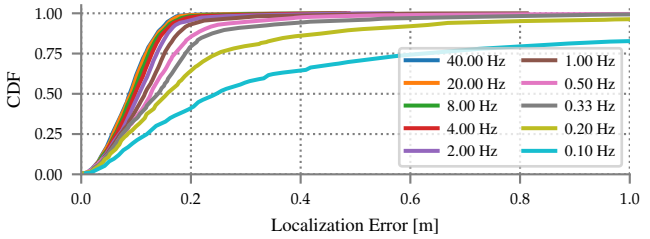

Fig. 10. Localization error with periodic UWB sampling at different frequencies.

ego-motion by wheel encoders sampled at a frequency of 50 Hz. The UWB tag performs SS-TWR in a round-robin fashion with all $m = 5$ anchors at a ranging frequency of 200 Hz, providing a maximum positioning rate of 40 Hz. SS-TWR involves a two-way message exchange between the tag and an anchor. The firmware is implemented atop Contiki OS [16] for the DecaWave EVB1000 platform [17], equipped with an STM32F105 MCU, the DW1000 transceiver, and a PCB antenna. UWB and odometry data are sent via WiFi to a laptop, which also stores the ground truth data acquired from the OptiTrack system. We recorded positioning information from 6 different generic trajectories, yielding ∼13000 location samples overall.

As a baseline, we first present the positioning results along the predefined trajectories when the UWB system is activated periodically and in isolation, i.e., without fusing UWB ranging information with the wheel encoders data. Fig. 9 shows the characterization of the ranging precision of UWB (left) and the positioning results (right) when measurements are collected at the maximum positioning frequency of 40 Hz. The standard deviation of the UWB ranging error is $\sigma = 11$ cm with a maximum error of 48 cm (Fig. 9, left). The ranging measurements are used to retrieve the position using a non-linear least squares solver, yielding the positioning error, computed with (4), shown in the right side of Fig. 9. The mean positioning error is $\mu_p = 9$ cm with a standard deviation of $\sigma_p = 5$ cm and a 90^{th} percentile error of 14 cm.

Fig. 10 shows the localization error when fusing odometry and UWB ranging data with the UKF, considering different periodic UWB sampling frequencies f_s. With $f_s = 40$ Hz the accuracy obtained is the same as the positioning error using only UWB (Fig. 9, right). Decreasing $f_s = 1$ Hz, the UKF yields a mean localization error $\mu = 12$ cm with a standard deviation of $\sigma = 8$ cm and a 90^{th} percentile error of 18 cm, slightly increasing the error w.r.t. the UWB-only solution, but reducing f_s by 97%. This, in turn, reduces the energy consumption, while increasing the scalability of the system. Unfortunately, the *right* sampling frequency that satisfies a user-defined bound (ψ, λ) and fulfills the worst-case scenario for robot navigation can only be speculated; recall the analysis of the previous section, summarized in Fig. 7.

Sample trajectories of the experiments are reported in Fig. 11, while the resulting PDFs are reported in Fig. 12, proving the applicability of the approach in an actual situation. Finally, Fig. 13 reports the mean frequencies and their standard deviations as a function of the imposed performance indices. Two final remarks are in order. First, the curve at periodic 1 Hz sampling in Fig. 10 has a maximum error $\lambda = 20$ cm with a probability of $\psi = 90\%$. With the same configuration of (λ, ψ) we obtain a frequency reduced to almost one half, as shown in the second orange bar on the left side of Fig. 13; this further confirms our claims in an experimental scenario. Second, the results of our comparison against [12], shown in Fig. 8, hold also in our experimental setup; however, they are here not reported here due to space limitations.

VII. Conclusions

We proposed a localization approach for mobile robots that relies on fusing the relative information coming from the encoders with absolute measurements coming from an external UWB infrastructure. External measurements are triggered adaptively only if and when the estimated positioning accuracy fails to meet the user-specified requirements. Our mathematical framework directly estimates uncertainty via the a-posteriori PDF and enables stochastic guarantees on system performance. Simulation and real-world experiments

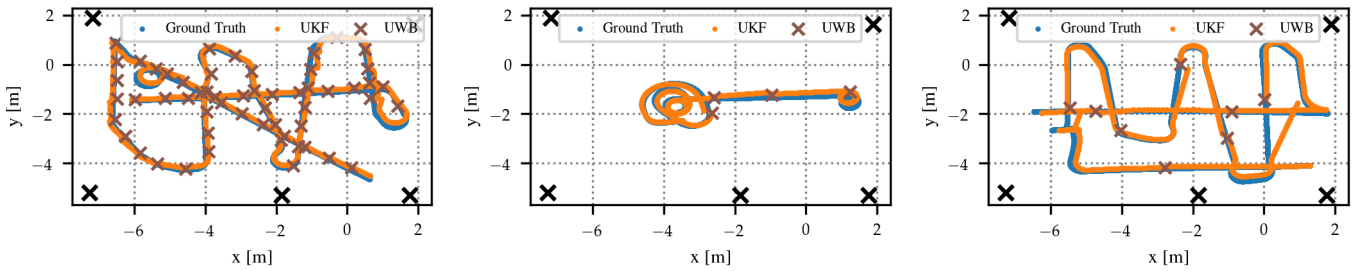

(a) ($\psi = 0.9, \lambda = 20$ cm) (b) ($\psi = 0.9, \lambda = 50$ cm) (c) ($\psi = 0.9, \lambda = 100$ cm)

Fig. 11. Localization tracking across three trajectories with $\psi = 0.9$ and different threshold λ. Each black cross represents a UWB anchor. The UKF output (orange) follows accurately the ground truth measurements (blue). As we increase λ, the number of UWB measurements (brown) needed decreases.

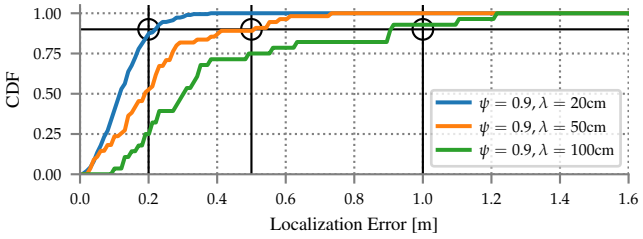

Fig. 12. Localization error with dynamic UWB sampling and $\psi = 0.9$ for different thresholds λ. The only measurements considered are taken just before triggering UWB ranging (i.e., worst-case scenario).

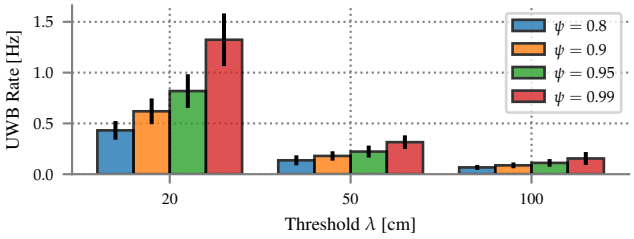

Fig. 13. UWB localization rate with $m = 5$ ranging exchanges for different confidence intervals ψ and thresholds λ over all trajectories.

confirm the soundness of our technique and its effectiveness in reducing the UWB sampling rate w.r.t. a fixed periodic solution, with direct benefits in terms of scalability and energy consumption. Further, our system allows for low-accuracy ranging schemes, which further improve these metrics.

We are building on these results along several dimensions. From a theoretical point of view, we are seeking a formalization of the conditions under which our technique yields an optimal sampling rate. From a system point of view, we are quantitatively investigating the scalability our technique can achieve in scenarios with multiple robots and its dependency on their motion patterns. Finally, from an application point of view, we are experimenting with various definitions of H_k geared towards specific classes of navigation tasks, concretely demonstrating the flexibility of our approach.

References

[1] R. Siegwart and I. R. Nourbakhsh, *Introduction to Autonomous Mobile Robots*. Scituate, MA, USA: Bradford Company, 2004.

[2] DecaWave, "DW1000 Data Sheet," 2016.

[3] M. J. Gallant and J. A. Marshall, "Two-dimensional axis mapping using LiDAR," *IEEE Trans. on Robotics*, vol. 32, no. 1, pp. 150–160, Feb. 2016.

[4] B. Dzodzo, L. Han, X. Chen, H. Qian, and Y. Xu, "Realtime 2D code based localization for indoor robot navigation," in *Proc. IEEE Int. Conference on Robotics and Biomimetics (ROBIO)*, Shenzhen, China, Dec. 2013, pp. 486–492.

[5] F. Zenatti, D. Fontanelli, L. Palopoli, D. Macii, and P. Nazemzadeh, "Optimal Placement of Passive Sensors for Robot Localisation," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and System (IROS)*. Daejeon, South Korea: IEEE/RSJ, Oct. 2016, pp. 4586–4593.

[6] P. Nazemzadeh, D. Fontanelli, D. Macii, and L. Palopoli, "Indoor Localization of Mobile Robots through QR Code Detection and Dead Reckoning Data Fusion," *IEEE/ASME Transactions on Mechatronics*, vol. 22, no. 6, pp. 2588–2599, Dec. 2017.

[7] P. Chen, Y. B. Xu, L. Chen, and Z. A. Deng, "Survey of WLAN fingerprinting positioning system," *Applied Mechanics and Materials*, vol. 380, pp. 2499–2505, Aug. 2013.

[8] "IEEE 802.15.4-2015, Standard for Low-Rate Wireless Networks."

[9] T. Chang and A. Mehta, "Optimal scheduling for resource-constrained multirobot cooperative localization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1552–1559, July 2018.

[10] M. Lahijanian, M. Svorenova, A. A. Morye, B. Yeomans, D. Rao, I. Posner, P. Newman, H. Kress-Gazit, and M. Kwiatkowska, "Resource-performance tradeoff analysis for mobile robots," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1840–1847, July 2018.

[11] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.

[12] E. Fresk, K. Ödmark, and G. Nikolakopoulos, "Ultra wideband enabled inertial odometry for generic localization," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 11 465–11 472, 2017.

[13] K. S. Chong and L. Kleeman, "Accurate odometry and error modelling for a mobile robot," in *Proceedings of International Conference on Robotics and Automation*, vol. 4, April 1997, pp. 2783–2788 vol.4.

[14] E. A. Wan and R. Van Der Merwe, "The unscented kalman filter for nonlinear estimation," in *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)*, Oct 2000, pp. 153–158.

[15] Optitrack, 2019. [Online]. Available: https://optitrack.com/motion-capture-robotics/

[16] P. Corbalán, T. Istomin, and G. P. Picco, "Poster: Enabling Contiki on Ultra-wideband Radios," in *Proc. of EWSN*, 2018.

[17] DecaWave, "DecaWave ScenSor EVB1000 Evaluation Board," 2013.