

A Real-Time Multi-Antenna SAR-based method for 3D localization of RFID Tags by a Moving Robot

Anastasios Tzitzis, Aristidis Chatzistefanou Raptopoulos, Traianos Yioultsis, *Member, IEEE*,
Antonios G. Dimitriou, *Senior Member, IEEE*

Abstract—In this paper, we present a prototype method for 3D localization of UHF RFID tags by a moving robot. The proposed method represents an extension of Phase ReLock, [24], in 3D space. Phase measurements are collected by at least two antennas, thus a multi-antenna synthetic aperture is created. We propose a new optimization problem, which involves measurements from all different antennas. Thanks to phase-unwrapping, the new multi-antenna optimization problem preserves its convex-like properties. Consequently, it is solved rapidly by standard optimization techniques. Furthermore, we introduce a confidence metric that can identify measured data, which deteriorate the accuracy of the estimations. These can be removed from the measurements' dataset. We have conducted extensive measurements, employing four antennas on top of a SLAM-enabled robot; i.e. the robot is able to create a map of the unknown environment and continuously estimate its pose inside the map. The proposed method outperforms prior art with respect to accuracy and computation time. It achieves mean 3D error less than 20cm with an estimation-time of only 0.17s per tag on an average laptop.

Index Terms—RFID, 3D Localization, Nonlinear Optimization, Phase Unwrapping, Performance Evaluation, Robotics, SLAM.

I. INTRODUCTION

RFID technology attracts growing interest due to the wide variety of related applications. In contrast to traditional optical technologies (camera, laser, barcode), RFID technology does not require visual contact, while its low cost and ability of high reading-rates makes its deployment beneficial in the fields of logistics, retail supply chain, healthcare, agriculture, smart houses, access control systems, libraries, etc. Concurrently, the proper utilization of the physical quantities measured by a commercial RFID reader (i.e. the phase and power of the RFID tag's signal) has drawn increasing attention in research, [1] - [6], in order to develop and improve localization methods that can accurately estimate the tag's location and essentially locate the "tagged" objects in the environment.

Localization in three dimensions usually requires the installation of multiple RFID readers and antennas. Adaptive Power Multilateration [7] proposes the deployment of a minimum of 4 antennas and the dynamical adjustment of their transmit power to estimate the antennas-to-tag distances. 3D-Batl [8] deploys again 4 antennas and exploits the power

path loss model to compute the antennas-to-tag and inter-tag distances. BackPos [9] is based on hyperbolic positioning; phase-difference measurements from each antenna-pair creates a hyperbolic locus of possible tag locations.

Accurate 3D localization demands for an increased number of installed antennas. However, the cost of a fixed network of multiple readers and antennas can be prohibitive and such solutions may not be applicable, especially in large areas. An alternative is to exploit moving antennas that can collect measurements from multiple locations. The motion of the antenna generates a virtual array, also known as synthetic aperture (SAR), an approach already known from radar applications.

The technique of SAR was initially applied in RFID technology by the holographic method [10] and later by its variation [11]. In [10], phase measurements taken along the antenna's trajectory are exploited to create a holographic image that quantifies the expectation a given point of the image is the actual tag's location. Eventually, it solves a Maximum-Likelihood (ML) problem, where the peak of the (non convex-type) image is found through an exhaustive search on all image points. 3DLoc [12] and 3DinSAR [13] perform holographic imaging in same manner, while they utilize a single antenna moving in two perpendicular directions in order to accomplish localization in three dimensions.

Such methods require antenna adjustment and control by a handling system, at the expense of time and effort since the presented process requires manual operation. On the contrary, an RFID-equipped robot that can navigate inside the environment allows for automatic inventorying and localization of all surrounding objects, without human intervention. Deploying a moving robot, though, raises the necessity of continuously tracing its pose while in motion, since SAR-based algorithms request for knowledge of the antenna's coordinates. Localizing a robot that moves along random trajectories is definitely a task that can be quite complex and challenging.

Antenna(s) mounted on robotic platforms are exploited in [14] - [16] in order to estimate the tag's location based on a ML approach; the solution of the localization problem corresponds to the maximum of a proper matching function, which is non convex. Therefore, the global maximum is found by exhaustive search over an assigned grid of possible tag locations. The trajectory of the robot in [14] is accurately obtained by a camera-system installed on the ceiling of the area, which continuously monitors the robot's motion, while in [16], a different solution is presented, which fuses odometry data with measurements from RFID tags placed at known positions. The performance of such grid-based methods strictly

Manuscript received December 7, 2020

This research has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH – CREATE – INNOVATE (project code:T1EDK-03032).

All authors are with the School of Electrical and Computer Engineering, Aristotle University of Thessaloniki, Greece, e-mail: antodimi@auth.gr.

depends on the grid resolution. The number of required calculations is proportional to the number of the grid points; hence, the computational cost for 3D problems and large tag populations is prohibitive for real-time applications.

In such framework, PSO-SAR [17] applies particle swarm optimization to the non-convex cost function, in order to reduce [14]'s computational cost, caused by the grid's exhaustive search. The algorithm moves a swarm of particles toward the global solution of the problem, bypassing local maxima. In [18], an extended Kalman filter is presented to online localize the tags with respect to the employed robot. Kalman filters have the ability to account for the variance of the phase measurements' Gaussian error, while relative localization is achieved quite fast, in comparison with grid-based methods, due to the filter's low complexity. However, both methods have the disadvantage of dependency on the algorithm's parameters, which necessitate preliminary investigation for the optimal values.

In addition to the robot-location-aware methods, the "fingerprinting" methods do not require any knowledge of the antenna's position(s). They employ tags at known positions as reference landmarks and evaluate the resemblance of the measured quantities between reference and target tags; such resemblance is quantified by a proper formula including several design parameters. LANDMARC [19] compares the RSSI values and estimates the location of the target tag as the middle point of the K nearest reference tags. Vire [20] proposes an improved method based on LANDMARC, by introducing simulated tags, in addition to a few real reference tags. In similar manner, the fingerprint technique can be applied to measurements collected along synthetic apertures (i.e. by a moving robot), as demonstrated in [21] and [22], while the performance is online evaluated, by treating the reference tags as target ones. Pinit [23] also employs a moving robot to collect measurements and compares the multipath profile between target and reference tags to accomplish accurate localization in scenarios of severe multipath.

Although the fingerprinting-based methods have low complexity and are considered real-time, their performance is quite sensitive to the algorithm's design parameters (e.g. number of neighbors considered, etc), requiring an analysis to find and use the optimal values. Most importantly, such methods depend significantly on the number and density of the employed reference tags; an increase of the reference tag population may increase the achieved accuracy but it slows down the estimation process by increasing the computational cost. Especially for three-dimensional localization, a dense 3D reference grid is necessary, while measuring the ground truth of the reference tags demands for excessive preliminary preparation.

Recently, we proposed a novel SAR-based localization method, called "Phase ReLock" [24], which exploits phase measurements obtained by a single antenna mounted on top of a prototype robotic vehicle. By performing phase-unwrapping to the measured data and creating a new unwrapped phase model, a convex-type optimization problem is built, which is able to be rapidly solved by standard nonlinear optimization algorithms. This property makes "Phase ReLock" indepen-

dent of a calculation's grid and prevents from performing a computationally expensive search, in contrast to other SAR methods which are grid-based. Experimental results prove that tag-estimations are derived orders of magnitude faster in comparison with state-of-the-art grid-based methods, while slightly better accuracy is delivered.

In [25] and [26], it was demonstrated that localization in three dimensions can be feasible by a single antenna, when non-straight synthetic apertures are exploited; i.e. the robot moves in curved trajectories. In [25] particularly, a simulation-analysis was implemented to evaluate the effect of the robot's trajectory on the achieved accuracy and reliability. However, even if the 3D location of the tag is estimated in less than a second, the delivered accuracy can not be considered quite satisfactory for realistic inventorying applications (mean localization error of about $0.5m$).

Additionally, the required estimation of the robot's trace is accomplished through SLAM (Simultaneous Localization and Mapping). In addition to RFID equipment, the robot carries a set of sensors (laser and cameras), which enable it to autonomously navigate in any a-priori unknown environment, create a 2D/3D map of it, and continuously estimate its trace in the created map by fusing odometry and laser data. Thus, the complex task of the robot is to identify and pinpoint the surrounding objects into the created map [27]. However, a trajectory derived by any SLAM algorithm suffers from various errors, induced among others by deformations of the created map. Such errors are expected to directly affect the accuracy of any RFID localization method applied.

In this paper, we present the extension of Phase ReLock to multi-antenna synthetic apertures. Initially, the phase curve obtained by each available antenna is unwrapped. In contrast to [25], we create a new optimization problem, including measurements from multiple antennas and searching for the actual coordinates of the tag in 3D space. The new multi-antenna optimization problem, preserves its convex-like property. As a result, not only it can be solved rapidly by standard optimization techniques, exactly as its equivalent 2D problem, but it is also robust and independent of algorithm's design parameters. Furthermore, we introduce a confidence metric that can identify less accurate estimations, which are then treated accordingly. An experimental campaign is conducted to evaluate the method's performance under different environmental conditions and algorithm's parameters and compare it against other localization methods.

The paper is organised as follows: Section II presents the localization method "Phase ReLock 3D". Section III introduces a confidence metric that evaluates the performance of the method, while Section IV gives the experimental results and comparisons. Finally, Section V concludes our findings.

II. PROBLEM FORMULATION

Consider the objective of localizing a static RFID tag by a moving robot that carries the required RFID equipment; i.e. reader(s) and antenna(s). The position of a tag is denoted as $A_{tag} = [x_{tag}, y_{tag}, z_{tag}]$. The antenna's location at time-index t , is denoted as $A_t = [x_t, y_t, z_t]$, while the phase of

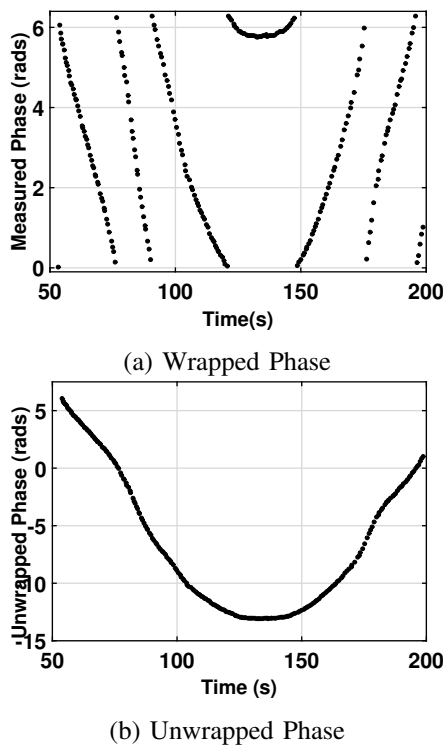


Fig. 1: Phase Unwrapping eliminates the 2π jumps and produces a continuous curve with support in $(-\infty, +\infty)$.

the backscattered signal from the tag to the antenna recorded at the same time t is denoted as θ_t . The measured phase is associated with two times the Euclidean distance between the antenna and the tag, $\|A_{tag} - A_t\|_2$, in order to account for the round trip propagation from a monostatic RFID reader. The theoretical phase-distance model is:

$$\phi_t = \left(\frac{2\pi}{\lambda} 2 \|A_{tag} - A_t\|_2 + \phi_0 \right) \bmod 2\pi, \quad t \in [1, T] \quad (1)$$

where ϕ_t is the expected phase that would have been measured if the tag's and antenna's locations were A_{tag} and A_t , respectively; notice θ_t is the measured phase sample. T is the total number of recorded measurements. In (1), λ stands for the wavelength of the carrier frequency, while ϕ_0 is an additional phase-rotation, irrelevant to the wave propagation. It corresponds to a phase offset induced by the tag's and the reader's electronics, cable effects, antennas' leakage, etc. However, for the same RFID setup, i.e. same antenna and tag pair, ϕ_0 is common for all measurements, but unknown.

The final output of (1) is a remainder to 2π , to indicate that the phase takes values only in intervals of 2π , as shown in Fig. 1 (a). Phase is a periodic function that repeats for every $\lambda/2$ change of antenna-to-tag distance. This kind of periodicity introduces an ambiguity; phase cannot directly reflect the true value of distance between tag and antenna, but just a series of possible distance values, which differ from each other by increments of half wavelengths, $\frac{\lambda}{2}$. Considering the above, an alternative phase-distance model can be given by

$$\phi_t = \frac{2\pi}{\lambda} 2 \|A_{tag} - A_t\|_2 + \underbrace{\phi_0 + 2k_t\pi}_{c_t}, \quad t \in [1, T] \quad (2)$$





Fig. 2: The motion of robot leads to the creation of a multi-antenna synthetic aperture, composed by M single-antenna apertures at different height.

where $2k_t\pi$ with $k_t \in \mathbb{Z}$, represents the aforementioned cycle ambiguity, which is not common for all measurements, as the presence of subscript t can imply.

In order to deal with this issue, phase unwrapping can be deployed. By properly adding multiples of 2π to each measured sample, the phase curve is reconstructed, such that it obtains a continuous form, free of jumps every 2π . The unwrapped phase measurements, denoted as $\check{\theta}_t$, have now support in $(-\infty, +\infty)$, as shown in Fig. 1 (b). In accordance to the reconstructed measured phase, the unwrapped phase-distance model becomes:

$$\check{\phi}_t = \frac{2\pi}{\lambda} 2 \|A_{tag} - A_t\|_2 + \underbrace{\phi_0 + 2k_t\pi}_c, \quad t \in [1, T], \quad (3)$$

where $\check{\phi}_t$ is the expected unwrapped phase that would have been measured if the tag's and antenna's locations were A_{tag} and A_t , respectively; notice $\check{\theta}_t$ is the unwrapped measured value. In (3), $2k_t\pi$ represents again the cycle ambiguity. However, this ambiguity term is now common for all unwrapped phase samples, in contrast to the case of wrapped data. Since both terms ϕ_0 and $2k_t\pi$ are common for all collected data of same antenna-tag pair, their summation c , is also a common constant, but still unknown.

A. Multi-Antenna Synthetic Aperture

An antenna mounted on a moving robot allows tag-readings from multiple locations, generating a synthetic aperture, also known as virtual antenna array. When M antennas are employed at different heights on the robot, M synthetic apertures are formed, composing a multi-antenna synthetic aperture, as depicted in Fig. 2 (a). For the t^{th} out of T_j aperture locations of the j^{th} out of M antennas, we define the observation sequence $\mathbf{O}_t^{(j)} = (\mathbf{A}_t^{(j)}, \check{\theta}_t^{(j)}, \lambda^{(j)})$, which holds information about the coordinates of the antenna, the unwrapped phase value and the carrier wavelength, see Fig. 2 (b). The absence of index t in λ implies that the operating frequency of each antenna j is fixed for the entire inquiry process, but is not necessarily the same for different antennas. By applying the observation sequences $\mathbf{O}_t^{(j)}$, $j \in [1, M]$ and $t \in [1, T_j]$, to the theoretical phase model (3), a system of equations $S = \{S^{(1)}, \dots, S^{(j)}, \dots, S^{(M)}\}$ is built:

$$S = \left\{ \begin{array}{l} S^{(1)} \left\{ \begin{array}{l} \check{\theta}_1^{(1)} = \left(\frac{4\pi}{\lambda^{(1)}} \|\mathbf{A}_{tag} - \mathbf{A}_1^{(1)}\|_2 + c^{(1)} \right) \\ \check{\theta}_2^{(1)} = \left(\frac{4\pi}{\lambda^{(1)}} \|\mathbf{A}_{tag} - \mathbf{A}_2^{(1)}\|_2 + c^{(1)} \right) \\ \vdots \\ \check{\theta}_{T_1}^{(1)} = \left(\frac{4\pi}{\lambda^{(1)}} \|\mathbf{A}_{tag} - \mathbf{A}_{T_1}^{(1)}\|_2 + c^{(1)} \right) \\ \vdots \\ \check{\theta}_1^{(j)} = \left(\frac{4\pi}{\lambda^{(j)}} \|\mathbf{A}_{tag} - \mathbf{A}_1^{(j)}\|_2 + c^{(j)} \right) \\ \vdots \\ \check{\theta}_t^{(j)} = \left(\frac{4\pi}{\lambda^{(j)}} \|\mathbf{A}_{tag} - \mathbf{A}_t^{(j)}\|_2 + c^{(j)} \right) \\ \vdots \\ \check{\theta}_{T_j}^{(j)} = \left(\frac{4\pi}{\lambda^{(j)}} \|\mathbf{A}_{tag} - \mathbf{A}_{T_j}^{(j)}\|_2 + c^{(j)} \right) \\ \vdots \\ S^{(M)} \left\{ \begin{array}{l} \check{\theta}_1^{(M)} = \left(\frac{4\pi}{\lambda^{(M)}} \|\mathbf{A}_{tag} - \mathbf{A}_1^{(M)}\|_2 + c^{(M)} \right) \\ \check{\theta}_2^{(M)} = \left(\frac{4\pi}{\lambda^{(M)}} \|\mathbf{A}_{tag} - \mathbf{A}_2^{(M)}\|_2 + c^{(M)} \right) \\ \vdots \\ \check{\theta}_{T_M}^{(M)} = \left(\frac{4\pi}{\lambda^{(M)}} \|\mathbf{A}_{tag} - \mathbf{A}_{T_M}^{(M)}\|_2 + c^{(M)} \right) \end{array} \right. \end{array} \right. \end{array} \right.$$

Sub-system $S^{(j)}$ corresponds to the synthetic aperture of the j^{th} antenna and $c^{(j)}$ is the common offset applied to all unwrapped measurements of j^{th} antenna. The unknown parameters of each sub-system $S^{(j)}$ are 4: the tag's coordinates $\mathbf{A}_{tag} = [x_{tag}, y_{tag}, z_{tag}]$ and the phase offset $c^{(j)}$. Extending to augmented system S , the latter consists of $\sum_1^M T_j$ equa-

tions and $3 + M$ unknown parameters represented by vector $\mathbf{p} = (\mathbf{A}_{tag}, c^{(1)}, \dots, c^{(M)})$.

System S represents a data-fit problem, according to which, the set of unwrapped phase measurements $\check{\theta}_t$, should match to the set of theoretical/expected values $\check{\phi}_t$, derived by (3). Since S is an overdetermined system, the solution is obtained in a least square sense and the optimum parameters represented by vector $\mathbf{p}^{opt} = (\mathbf{A}_{tag}, c^{(1)}, \dots, c^{(M)})^{opt}$ are sought to minimize the sum of squared differences between the measured and expected values. Particularly,

$$\mathbf{p}^{opt} = \arg \min_{\mathbf{p}} \check{F}(\mathbf{p}) \quad (4)$$

where matching function $\check{F}(\cdot)$ follows:

$$\check{F}(\mathbf{p}) = \sum_{j=1}^M \left[\sum_{t=1}^{T_j} \left(\check{\theta}_t^{(j)} - \frac{4\pi}{\lambda^{(j)}} \|\mathbf{A}_{tag} - \mathbf{A}_t^{(j)}\|_2 - c^{(j)} \right)^2 \right] \quad (5)$$

B. Nonlinear Optimization

Objective function \check{F} in (5) is nonlinear and cannot be solved in closed form by simple matrix techniques. Minimization across the parameter space is accomplished through nonlinear optimization that exploits iterative algorithms. Starting from an initial selection of the unknown parameters, denoted as \mathbf{p}_0 , each iteration adjusts the solution, such that \check{F} in (5) repeatedly decreases and the fit is improved. This procedure stops when some convergence criteria are met; i.e. a (local, in general) minimum is reached. Seeking the optimum values \mathbf{p}^{opt} though, implies seeking the global minimum of \check{F} and not just a local one. Therefore, in order for any of the standard optimization algorithms to be used, the involved cost function should exhibit one and only global minimum, in other words be convex. Fig. 3 represents the 3D illustration of \check{F} along the directions x_{tag} , y_{tag} , z_{tag} ; the three 2D slices of the main plot pass through the tag's estimated position. The smooth change of coloring implies convex-type surfaces. Additionally, the subplot at the upper left corner of Fig. 3 depicts the shape of \check{F} on the $x_{tag} - c^{(1)}$ parameter-plane, indicating convexity of the cost function with respect to the ambiguity term $c^{(1)}$. \check{F} maintains the property of convexity along $c^{(j)}$ for $j \in [2, M]$, as well. It is also worth noting, that the shape of \check{F} for the $y_{tag} - c^{(1)}$ and $z_{tag} - c^{(1)}$ planes is also similar. To better demonstrate the cost-function, each slice of the main plot is also shown separately in Fig. 4, where \check{F} 's smooth and convex-type surface can be observed. The involved objective function does not suffer from local extrema and deployment of nonlinear optimization algorithms is feasible.

Depending on the way the minimization-problem is solved at each iteration, there are various optimization algorithms. Herein, the method adopted is based on the trust-region strategy [28] - [30]. Let vector $\mathbf{p}_k = (\mathbf{A}_{tag}, c^{(1)}, \dots, c^{(M)})_k$ denote the estimation of the unknown parameters at the k^{th} iteration. According to the trust-region approach, the original cost function \check{F} in (5) is approximated near \mathbf{p}_k with a model

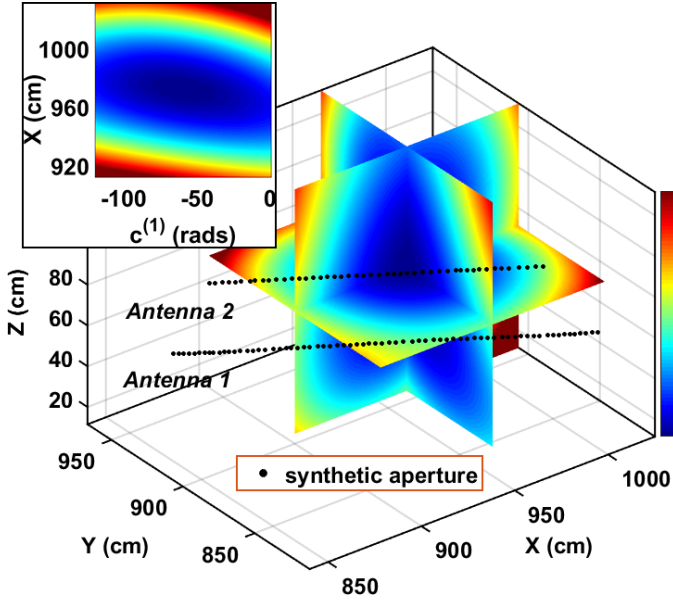


Fig. 3: Cost function \check{F} in (5) shown in 2D slices on the $(x_{tag} - y_{tag})$, $(y_{tag} - z_{tag})$, $(z_{tag} - x_{tag})$ and $(x_{tag} - c^{(1)})$ planes. The smooth change of coloring implies convex-type surfaces. The black dots represent the synthetic aperture of each employed antenna.

m_k , which is normally the second-order Taylor expansion of the former:

$$\begin{aligned} m_k(s) &= \check{F}(\mathbf{p}_k + s) \\ &= \check{F}(\mathbf{p}_k) + \nabla \check{F}(\mathbf{p}_k)^T s + \frac{1}{2} s^T \nabla^2 \check{F}(\mathbf{p}_k) s \end{aligned} \quad (6)$$

Model $m_k(s)$ is then minimized, while its minimum is sought within a trust-region of radius Δ_k , centered at \mathbf{p}_k :

$$\mathbf{s}_k = \arg \min_s m_k(s), \quad s. t. \quad \|\mathbf{D}_k \mathbf{s}_k\|_2 \leq \Delta_k \quad (7)$$

where \mathbf{s}_k is called trial step and Δ_k essentially determines the maximum allowed value of \mathbf{s}_k . Furthermore, since the problem's parameters may take values that have widely different scales, a spherical trust region may not be efficient. \mathbf{D}_k corresponds to a scaling matrix that ensures vector $\mathbf{D}_k \mathbf{s}_k$ has entries of same order of magnitude. Step \mathbf{s}_k is only the minimum of model m_k and its performance should be evaluated with respect to \check{F} . We define the agreement between the actual function \check{F} and the approximate model m_k as the ratio between their corresponding reductions:

$$\rho_k = \frac{F(\mathbf{p}_k) - F(\mathbf{p}_k + \mathbf{s}_k)}{m_k(0) - m_k(\mathbf{s}_k)}. \quad (8)$$

Ratio ρ_k is exploited to decide whether trial step \mathbf{s}_k is accepted for next iteration and adjust the trust region's radius Δ_{k+1} . Since the aim of each iteration is to decrease \check{F} , step \mathbf{s}_k is accepted as long as ρ_k takes positive values. Rejecting a step corresponds to a wasted iteration/computation time and hence, a general strategy is to reject it only when it leads to

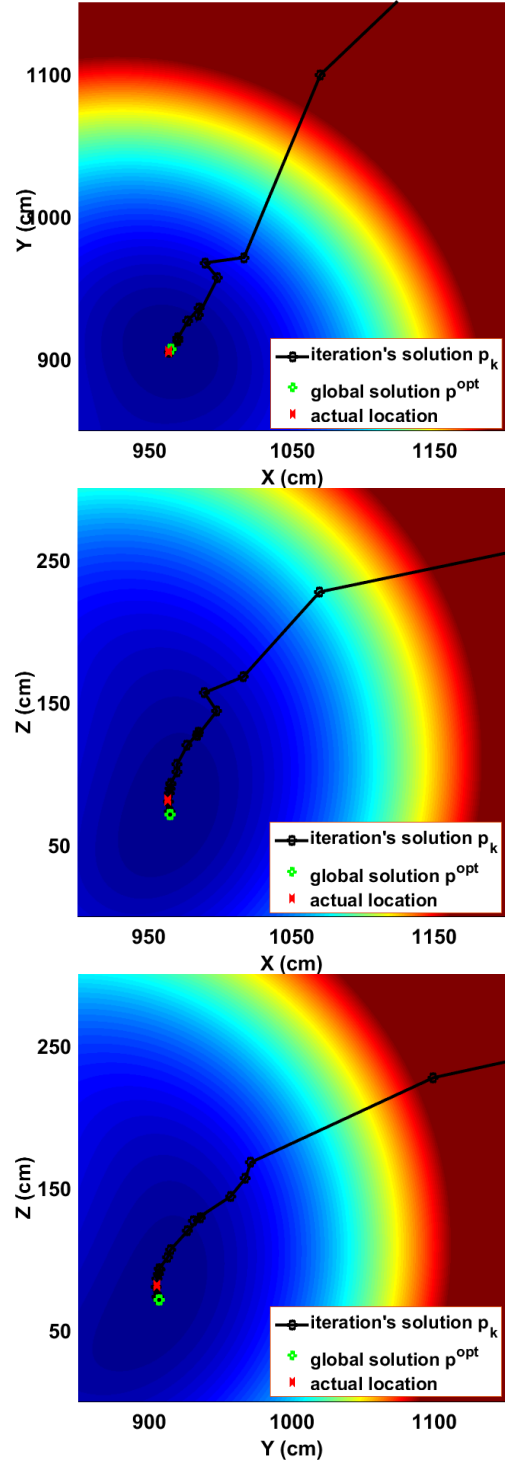


Fig. 4: The x-y, x-z and y-z plane of objective function \check{F} shown in Fig. 3. Additionally, the progress of the optimization algorithm is shown, which iteratively converges to the global minimum. Black circles represent the estimation at each iteration, green circle the accepted solution and red cross the actual tag location.

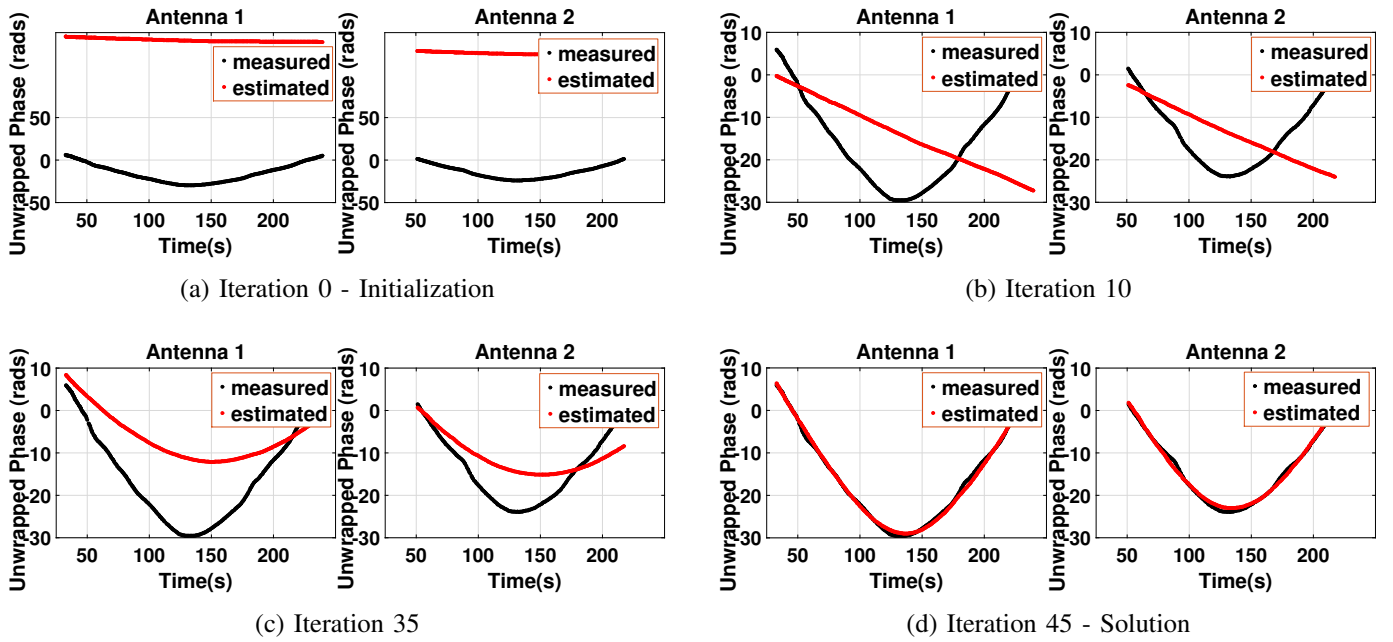


Fig. 5: Comparison between the measured unwrapped phase curves collected by the antennas and the estimated phase curves produced by (3) for the parameter values $(A_{tag}, c^{(1)}, c^{(2)})$ estimated at different algorithm's iterations. The estimated curves produced by the initial values have a bad fit to the measured ones (a), however as the values are refined, the fit is improved (b)-(c), when eventually the optimum values deliver the best match.

an undesired increase of \check{F} , i.e. a negative ρ_k . The parameter vector of the next iteration is hence updated, such that

$$p_{k+1} = \begin{cases} p_k + s_k, & \text{if } \rho_k \geq 0 \\ p_k, & \text{if } \rho_k < 0 \end{cases} \quad (9)$$

Large values of ρ (close to 1) indicate that model m_k is an accurate approximation of \check{F} and hence, the trust region can be enlarged to allow potential larger and ambitious steps. On the contrary, small values of ρ_k imply that there was a great disagreement between m_k and \check{F} and the trust region should be shrunk in order to restrict the maximum allowed value of s_k . The region's radius is updated as follows:

$$\Delta_{k+1} = \begin{cases} \eta \Delta_k, & \text{if } \rho_k \geq c_1 \\ \Delta_k, & \text{if } c_1 > \rho_k \geq c_2 \\ \Delta_k / \eta, & \text{if } c_2 > \rho_k \end{cases} \quad (10)$$

where c_1, c_2 and η are non negative quantities for evaluating the suitability of the approximate model and adjusting the trust region's radius for next iteration; typical values are $c_1 = 0.75, c_2 = 0.25, \eta = 2$.

The iterative procedure shall be repeated until some solution p_k is accepted as the global solution of the original optimization problem. Convergence to the minimum is identified by evaluating the gradient of \check{F} at p_k ; the algorithm stops when

$$\|\nabla \check{F}(p_k)\|_2 \leq \epsilon, \quad (11)$$

where ϵ is a small non negative value that represents the termination tolerance. Alternative termination criteria are i) no further decrease of the amplitude of cost function \check{F} , ii) no

change of the parameter values p_k for consecutive iterations, iii) a maximum number of iterations.

Fig. 4 demonstrates the progress of converging to the global minimum of \check{F} (5) by applying the trust region algorithm. The initial values of the parameters were $p_0 = (A_{tag}(cm), c^{(1)}(rad), c^{(2)}(rad))_0 = (2000, 2000, 500, 0, 0)$. Starting by these parameter values, the algorithm adjusts them at each iteration (depicted by black circles), until it eventually reaches the global minimum of \check{F} that corresponds to the optimum values $p^{opt} = (965.1, 906.7, 71.4, -54.3, -58.8)$. The optimum solution is depicted with a green circle, while the true tag's coordinates are represented by the red cross.

Additionally, Fig 5 compares the measured (unwrapped) phase curves obtained by the two apertures shown in Fig. 3, with the estimated curves produced by (3) for the parameters estimated at different algorithm's iterations. Initially, there is a great disagreement between the measured and estimated data, see Fig. 5 (a), since the initial values p_0 are far from the actual solution. However, as the algorithm gets closer to the global minimum of \check{F} through the iterative process, the produced estimated curves tend to resemble the measured ones, see Fig. 5 (b)-(c). Eventually, the algorithm converges to the global minimum and the optimum parameter values p^{opt} deliver the best match, see Fig. 5 (d).

Summarizing, "Phase ReLock 3D" solves the 3D localization problem by initially unwrapping the phase measurements obtained from a multi-antenna synthetic aperture. Then, it crafts and solves an optimization problem by deploying the trust-region optimization algorithm.




Fig. 6: 3D illustration of \tilde{F} in (5), when different number of antennas are used to locate same tag.

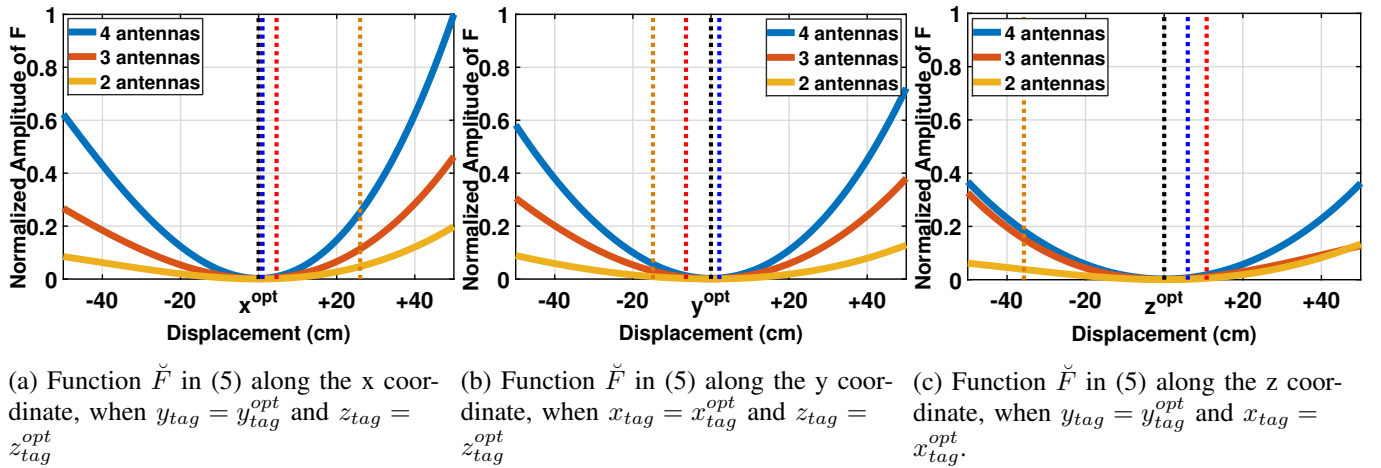


Fig. 7: The curves of \tilde{F} along x_{tag} , y_{tag} and z_{tag} direction, for different number of employed antennas. The curves have been centered such that their minimums coincide. The x-axis of each plot represents the displacement from the estimated parameter-value. The colored vertical dashed lines correspond to the actual coordinate of the tag with respect to the estimated coordinate represented by the black dashed line passing through the minimum of each curve.

III. PERFORMANCE EVALUATION

After deriving the best-fit values of the unknown parameters, $\mathbf{p}^{opt} = (\mathbf{A}_{tag}, c^{(1)}, \dots, c^{(M)})^{opt}$, the performance of the optimization should be evaluated in order to obtain an intuition of the expected localization error. Various factors can deteriorate the efficiency of any localization method, such as multipath/noise, inadequate collection of measurements, inadequate number of employed antennas, errors induced by the phase-unwrapping algorithm, improper robot trajectory, etc. Either case will lead to a poor estimation.

This effect will be reflected in the variance of the estimated parameters; i.e. a common measure in optimization that quantifies the (un)certainly of the estimated solution based on the quality of the fit. In fact, the variance represents how much the fit will worsen, as the parameters take values away from the optimum ones. Small variance indicates that only values, which tend to be quite close to the best-fit values, can produce an equivalently good fit as the latter. On the contrary, high variance implies that a fit, almost as good as the best one, can be achieved even for values that are quite spread out from the optimum values. Thus, the smaller the value of variance is,

the better and more confidently the solution is identified.

Asking how much the fit worsens as the parameters take values away from the optimum, is essentially asking how much the amplitude of cost function \tilde{F} in (5) increases. When the function's curve is quite steep towards the minimum, values away from it lead to a highly increased amplitude and a much worse fit, whilst for nearly flat curves, \tilde{F} 's amplitude increases slowly for values away from the optimum. As a result, quite steep curves are likely to produce reliable estimations, whilst for nearly flat curves, the ambiguity of the estimation is increased.

Therefore, the *local curvature around the global minimum* will be explored in order to evaluate the confidence of the estimation and consequently, the expected localization error. The information of curvature along each parameter-direction is represented by the respective second partial derivative of \tilde{F} in (5). Since the variance is inversely related to the curvature of \tilde{F} , its value for each parameter in $\mathbf{p} = (\mathbf{A}_{tag}, c^{(1)}, \dots, c^{(M)})$ is given by:

$$\sigma_{\mathbf{p}(i)}^2 = C(i, i), \quad i \in [1, 3 + M] \quad (12)$$

where $C(i, i)$ is the i^{th} diagonal element of the covariance matrix C , given by

$$C = \sigma_r^2 H^{-1} \quad (13)$$

where σ_r^2 is the variance of the problem's residuals and H stands for the $(3 + M) \times (3 + M)$ Hessian matrix of \check{F} in (5), computed for the optimum values \mathbf{p}^{opt} :

$$H = \nabla^2 \check{F}(\mathbf{p}^{opt}). \quad (14)$$

When the variance of each parameter is available, the corresponding confidence intervals can be computed. The confidence interval (CI) is essentially a range of values around the estimated one, which contain the true value with a certain degree of confidence; typically the 95% interval is exploited, i.e. true value lies within the calculated range with a 95% probability. The narrower the length of confidence interval, the more confident the estimation. Hereinafter, CI will refer to the 95% confidence interval.

Figs 6 (a), (b) and (c) show the cost function \check{F} (5), when measurements from 2, 3 or 4 antennas, respectively, are used in order to locate the same tag. By observing the colors in Fig. 6, one can notice that for values of A_{tag} away from the optimum, the amplitude of \check{F} increases more abruptly as the number of employed antennas increases. Hence, the confidence of the estimation is higher for the case of 4 employed antennas, since the produced curve is much steeper around the minimum.

This effect is also addressed in Fig. 7, which depicts the produced curves of \check{F} along each coordinate-axis x_{tag} , y_{tag} and z_{tag} , when 2, 3 and 4 antennas are employed. In order to acquire a better comparison of the curvatures, all curves have been centered such that their minimums coincide, while x-axis of each plot represents the displacement from the estimated coordinate x^{opt} , y^{opt} and z^{opt} . When measurements by 4 antennas are exploited, the produced curves are much steeper around the minimum, leading to a more reliable and accurate identification of the solution. Furthermore, the vertical dashed lines indicate the localization error along each coordinate. Since x-axis of each plot represents the displacement from the estimated coordinate, the colored lines represent the difference between the latter and the tag's true coordinate; e.g. in 7 (a) the yellow dashed line passing through the value of $+25cm$ implies that the localization error along the x coordinate is $25cm$. The closer each colored line lies to the optimum value (displayed by the black dashed line passing through the minimum), the lower the localization error. One can notice that the error decreases for each coordinate, as the number of antennas increases.

However, it should be noted that the curvature of the cost function, and hence the confidence of the estimation, does not always increase, by increasing the number of involved antennas/measurements. Sometimes, data from an antenna may be corrupted by noise or multipath. In such cases, the corresponding reliability decreases when the specific data are accounted by the cost function. Thanks to the proposed criterion, i.e. the length of confidence interval CI , such data may be identified and excluded from the estimation-process. Such experimental results will be presented in section IV-D.




Fig. 8: Photo of robot "Frida" from a recent international exhibition.



(a) Two banners placed side-by-side



(b) Two banners placed back-to-back

Fig. 9: Photos during the experiments.

IV. EXPERIMENTAL CAMPAIGN

We have constructed a prototype robot, "Frida" (see Fig. 8), which was employed for the experimental campaign. Frida carries 2 Impinj R420 readers and 4 circularly polarized UHF RFID antennas per side, installed at different heights on the robot; in the following experiments, only one reader and 4 antennas were activated, illuminating one half-space as the robot moves. The heights of the antennas were $0.63m$, $1.07m$, $1.53m$ and $1.93m$. The experiments took place in a laboratory room where desks, chairs, shelves and many other scatterers are located. 100 passive UHF RFID tags were attached on two 2-meters-tall banners and used to evaluate the performance of our method, "Phase ReLock 3D". Additionally, various banner set-ups and robot trajectories were employed in order




Fig. 10: The estimated robot trajectories inside the 2D map created for the experiment with the side-by-side banners.

to test the method under different conditions. The employed experimental set-ups were i) a single banner (50 tags), ii) two banners placed back-to-back (100 tags) and iii) two banners placed side-by-side (100 tags), see Fig. 9, while the robot was forced to move along straight paths, "V"-type paths (i.e. making one turn) and random trajectories such as slaloms ("S"-type). The employed robot's speed was 1.5cm/s .

The poses of the moving robot during the experiments were not known but estimated by a SLAM algorithm [31] - [33]. In addition to RFID equipment, the robot carries a laser and 3D depth cameras that enable it to create a map of the environment and continuously localize itself inside the map, by fusing odometry, lidar and depth-camera data. Fig. 10 depicts the estimated trajectories of the robot (straight, "V"-type" and "S-type") inside the 2D map that was created for one of our experimental set-ups. However, a robot trajectory estimated by any SLAM algorithm is inaccurate and suffers from various errors, caused most importantly by sensor noise and deformations of the created map. Such errors are expected to propagate to the RFID localization error.

Both the robot's (estimated) locations and the measured phase samples are accompanied by their timestamp. The operating system of the robot, which estimates the locations of the latter, acquires the time through a Network Time Protocol (NTP) server that gives a time resolution of μs . Similarly, the RFID reader's clock is synchronized with the robot's to obtain a time resolution of μs as well. Since the speed of the robot is a few cm/s , the ambiguity of the robot's pose that corresponds to each tag measurement, is in practice neglected. Hence, the time resolution of the system is not considered as a source of error that deteriorates the tag's localization accuracy.

A. Phase ReLock 3D against prior art

"Phase ReLock 3D" is then applied to the collected measurements and compared against other SAR based methods that use the Maximum Likelihood approach; i.e. the location

of the tag corresponds to the global solution of a matching function. In particular, the convex optimization applied by Phase ReLock is compared against two different approaches, 1) exhaustive search on an assigned grid of locations and 2) particle swarm optimization along the search space.

1) *Grid-based methods* [10], [14], [15]: According to grid-based methods, the tag's location is sought along a grid of possible locations such that a non convex cost-function is maximized; the cost function essentially represents the probability of a grid-point to be the actual tag position. The presence of multiple local minima and maxima though, necessitates an exhaustive search on the assigned grid, in order to identify the global maximum over the local ones.

A main property of grid-based methods is that both accuracy and execution-time strictly depend on the size of the grid; i.e. the number of possible tag locations explored. For dense and large grids the localization error is expected low; however, the exhaustive search over the grid requires a high computational cost and increased execution-time. On the contrary, a limited and sparse grid can allow for high execution-speeds and real-time applications, but the accuracy can be proven poor.

Among several grid-based methods, we have implemented the holographic method presented in [10] and compared it against "Phase ReLock 3D". Different variations of [10] are expected to have same performance regarding the delivered accuracy and the required estimation-time, since they optimize the same [15], or an equivalent [14], matching function. In the experiments conducted, the space of interest is defined by the length of the robot's trajectory (4m), the reading range (depending on the transmit power) of the reader (6m) and the banner's height (2m). As a result, the assigned three-dimensional grid extends to an area of $4\text{m} \times 6\text{m} \times 2\text{m}$, while different grid-densities were tested; the distance between adjacent grid-points takes the values of 2.5cm , 5cm , 10cm , 20cm and 30cm .

Fig. 11 shows the results of the two compared methods, Phase ReLock 3D and [10]. In particular, it presents the cumulative distribution function (CDF) of the achieved localization error for each experiment. The accuracy of holographic is inversely related to the grid-spacing. As the grid becomes sparser, it is less likely that a grid point is close to the true tag's location and as a consequence, the achieved localization error increases. Worst accuracy is delivered for the sparsest grid (i.e. 30cm), while quite dense grids in the order of 2.5cm and 5cm improved the performance. Concurrently, Phase ReLock 3D being grid-free, outperforms the holographic method for all of the deployed steps.

The results of all experiments are summarized in Fig. 12 (a), where the cumulative distribution function of error for all 550 estimated tag locations is presented. The performance of Phase ReLock can only be compared against [10] when the latter exploits 2.5cm and 5cm grid steps, but still is quite superior, especially for CDF values above 0.7. Fig. 12 (b) corresponds to the execution-time of each of the 550 estimations. Phase ReLock, being independent of a calculation's grid, is executed rapidly and any tag was estimated in ms order. On the contrary, [10]'s speed depends on the density of the grid; the denser the grid, the higher the computational cost of the

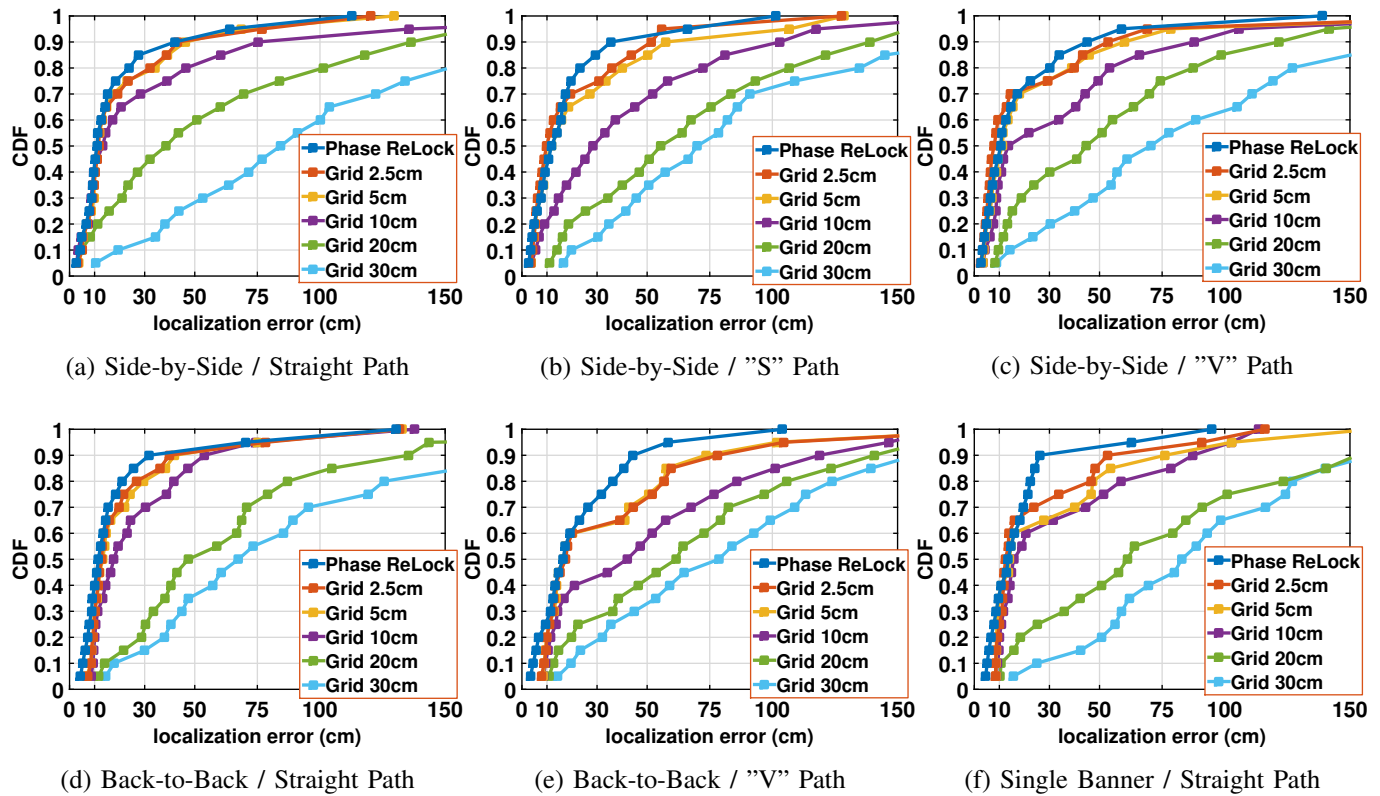


Fig. 11: Cumulative distribution function (CDF) of the achieved localization error by Phase ReLock and grid-based [10] for each of the experiment.

method, making essentially the method extremely slow for large grid densities.

Table I summarizes the performance of each method in terms of accuracy and computation-time per tag. The total time for all 550 tags is also given. Phase ReLock delivers the best accuracy with the lowest computational cost. It accomplished to locate 550 tags in about 1.5min, requiring less than 200ms per tag, with a mean 3D error of slightly less than 20cm. On the contrary, the most accurate grid-based localization (grid of 2.5cm spacing), located all tags with a mean error of 24cm but executed in more than 40 hours; i.e. 1600 times slower than Phase ReLock. The sparse grid of 30cm step managed to locate the tags in comparable time with Phase ReLock (1.8min), but with much worse accuracy (around 90cm).

TABLE I: Phase ReLock 3D vs Grid-Search

method	localization error		execution time	
	mean	std	mean	sum
Phase ReLock 3D	19.5cm	20.1cm	0.17s	1.5min
[10] (step 2.5cm)	24.5cm	27.5cm	4.2min	41.1h
[10] (step 5cm)	26.2cm	29.1cm	36.8s	5.5h
[10] (step 10cm)	38cm	37.8cm	4.7s	43.2min
[10] (step 20cm)	63.7cm	50.2cm	0.6s	5.7min
[10] (step 30cm)	87cm	59.9cm	0.2s	1.8min

2) *Particle swarm optimization* [17]: In order to speed up the identification of the global maximum, [17] proposes the deployment of particle swarm optimization (pso). Each particle essentially represents a candidate tag location. Initially, the algorithm randomly generates the population of particles

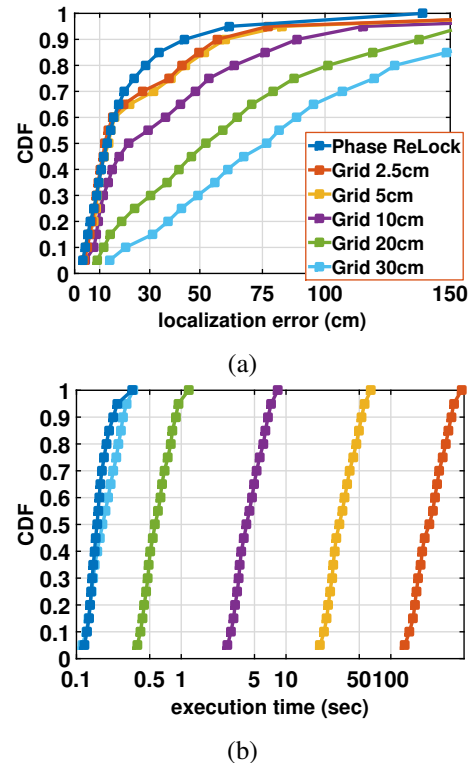


Fig. 12: Cumulative distribution function (CDF) of (a) the achieved localization error (b) and the required estimation-time by Phase ReLock and grid-based [10] for all experiments.

(swarm) within the search-space bounds and it iteratively updates each particle's location according to its individual best location and the best location of its neighbors. This procedure is expected to move the particles toward the global maximum.

Whilst particle swarm optimization reduces the computational cost of localization, compared to grid-based methods, it strictly depends on the initial (randomly generated) locations of the particles and the values of the algorithm's design parameters, e.g. the swarm size, etc. Large search spaces, such as three-dimensional, require a large particle population, while improper selection of the algorithm's parameters makes the swarm trap to a local maximum and the global solution is not identified.

We have implemented the pso algorithm by applying the parameter values as proposed in [17], while the dimensions of the search-space are defined same as before, $4m \times 6m \times 2m$. Fig. 13 investigates the method's performance under the effect of different swarm sizes; 50, 100, 500, 1000 and 2000 particles. The cumulative distribution function of error for all 550 estimated tag locations during all experiments is presented in Fig. 13 (a). As expected, the delivered localization error depends on the population of employed particles. Large populations offer a more thorough search and increase the chance of converging to the actual global maximum of the cost function, whilst deploying few particles fails to identify the global solution and method is trapped to a local maximum. The performance of Phase ReLock can only be compared against [17] when the latter employs a large swarm of particles, but still is quite more accurate. Fig. 13 (b) corresponds to the execution-time of each of the 550 estimations. Thanks to the exploitation of an unwrapped phase model, Phase ReLock preserves convexity of the cost function, which allows the global minimum of the problem to be rapidly found by iterative non-linear optimization. Any tag was estimated in *ms* order. On the contrary, [17]'s speed mainly depends on the swarm size; the more the employed particles, the more the calculations required and the computational cost of the method.

Table II summarizes the performance of each method in terms of accuracy and computation-time. Phase ReLock again delivers the best accuracy with the lowest computational cost. On the contrary, [17] required a swarm of 2000 particles to locate all tags with a mean error of 24cm but processing such a large particle population increased the execution time to 5.5 hours; i.e. more than 200 times slower than Phase ReLock.

TABLE II: Phase ReLock 3D vs Particle Swarm Optimization

method	localization error		execution time	
	mean	std	mean	sum
Phase ReLock 3D	19.5cm	20.1cm	0.17s	1.5min
[17] (2000 particles)	24.6cm	30.5cm	36.9s	5.6h
[17] (1000 particles)	25.2cm	29.8cm	17.2s	2.6h
[17] (500 particles)	33.1cm	46.3cm	7s	1h
[17] (100 particles)	59.2cm	74.8cm	1.7s	15.7min
[17] (50 particles)	83.5cm	89.5cm	0.77s	7min

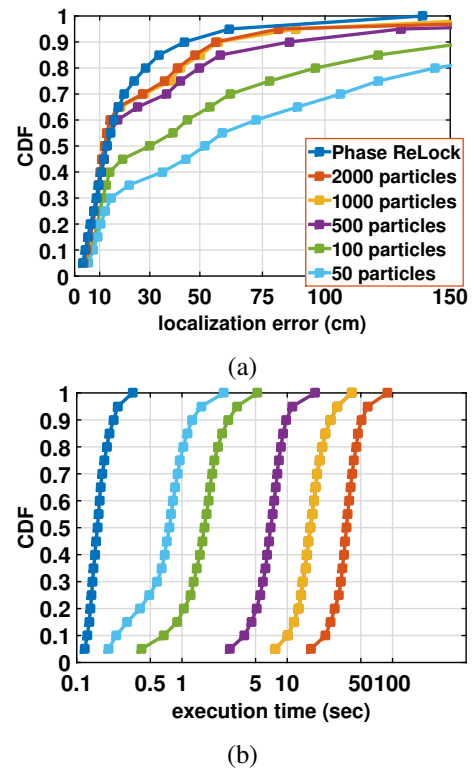


Fig. 13: Cumulative distribution function (CDF) of (a) the achieved localization error (b) and the required estimation-time by Phase ReLock and pso-based [17] for all experiments.

B. Effect of the robot's speed

Next, we evaluate the performance of Phase ReLock 3D for different robot speeds. The speed of the robot directly affects the reader's sampling rate and the number of collected measurements per tag. Hence, it is expected to influence the localization accuracy. When the antenna is moving in relation to the tag, the latter is successfully inquired for as long it is within the antenna's reading range. Thus, the number of the collected data is associated with the length of time the antenna's position allows for tag illumination. Small speeds offer a large collection of data, whilst a poor set of measurements is obtained for quite high speeds. The number of obtained data per tag also depends on the tag population within range; the more the existing tags that compete for an interrogation time-slot, the less the time-slots that each tag will eventually occupy and hence, the less the available measurements of it. As a result, the performance of Phase Relock as the speed increases, is expected to vary depending on the tag population/density of each experimental set-up.

The speed of the robot that was experimentally deployed was $1.5cm/s$. By downsampling the collection of measurements, higher speeds of robot can be assumed; e.g. downsampling the collected set by a factor of two is considered equivalent of doubling the robot's speed, resulting in a speed of $3cm/s$. We have tested values that are increments of the actual speed; i.e. ($\times 2$), ($\times 3$), etc. Fig. 14 demonstrates the effect of increasing the speed on the localization error for the three banner arrangements: (a) a single banner with 50 tags




Fig. 14: Cumulative distribution function (CDF) of the achieved localization error for different robot speeds and experimental tag set-ups.

attached on it, (b) two banners placed side by side (100 tags in total), (c) two banners placed back to back (100 tags in total). In the case of the single banner (lowest tag density), speeds up to $(\times 6)$ the original, seem not to deteriorate the performance of the method, implying that a sufficient set of measurements was available for processing. As the speed further increases, the method finds difficulties in locating the tags accurately, while localization fails for speeds higher than $(\times 10)$. When two banners are placed one next to the other, there is an adequate collection of measurements only when speeds up to $(\times 3)$ the actual one are considered, whilst the method completely fails when the robot moves faster than $(\times 8)$. The worst case scenario where the highest tag density is recorded, is the back-to-back banners. In this case, the collection of data per tag was quite poor and Phase ReLock performed well only for $(\times 2)$ the actual speed, while for speeds higher than $(\times 4)$ the original, the performance abruptly worsens. The above experimental analysis verified the assumption that the method's performance changes depending on the tag density in range, since the increase of speed affected the method's accuracy for each banner set-up differently.

C. Varying the number of employed antennas

The results of sections IV-A and IV-B correspond to the usage of all 4 available antennas carried by the robot. Since the 3D problem requires a minimum of 2 antennas, we investigate the effect of employing 2, 3 and 4 antennas on the performance of Phase ReLock 3D. Fig. 15 (a) represents the achieved localization error for all conducted experiments. In general, the overall accuracy improves with the increasing number of antennas. Specifically, best accuracy is delivered when all 4 antennas are employed, while the performance of 2 antennas is quite inferior. Meanwhile, the number of antennas also influences the computational cost of Phase ReLock, see Fig. 15 (b). The execution time for locating a tag (i.e. the time spent on unwrapping the measured phase curves and solving the optimization problem (4)) increases proportionally to the number of involved antennas. Each antenna introduces an additional curve for phase unwrapping and an additional parameter in the optimization problem; i.e. the constant offset applied in all its collected measurements. As a result, the

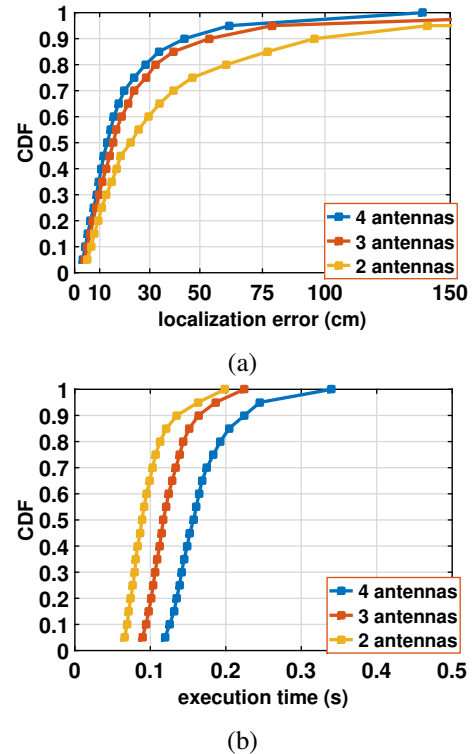


Fig. 15: Cumulative distribution function (CDF) of (a) the achieved localization error and (b) the required estimation-time for an increasing number of deployed antennas

estimation-time per tag increases with the increasing number of antennas linearly.

TABLE III: Effect of different number of employed antennas

Phase ReLock 3D	localization error		execution time	
	mean	std	mean	sum
2 antennas	39.3cm	46.7cm	0.08s	0.8min
3 antennas	24.6cm	29.2cm	0.11s	1min
4 antennas	19.5cm	20.1cm	0.17s	1.5min

Last but not least, Table III summarizes the results of our investigation by presenting the mean and standard deviation of the accomplished localization error, in combination with




Fig. 16: Localization error with respect to the confidence interval of parameters (a) x_{tag} , (b) y_{tag} and (c) z_{tag} for all 550 estimated tags.




Fig. 17: (a) Localization error with respect to the sum of confidence interval of parameters x_{tag} and y_{tag} . The vertical dashed lines represent the thresholds used to identify uncertain tags. (b)-(c) Identification and correction of poor estimations when a threshold of $72.95cm$ (99th percentile) and $14.04cm$ (95th percentile), respectively, is applied.

the mean execution time per tag and the total required time for all 550 tags. It is worth noting that by doubling the number of antennas, the achieved error is two times smaller (4 antennas deliver a mean error of $20cm$, while 2 antennas deliver $40cm$) and the speed of the estimation is two times slower (4 antennas required $1.5min$ for 550 tags while 2 antennas required $0.8min$).

D. Treating poor estimations

An estimation derived by Phase ReLock 3D is accompanied by the 95% confidence interval of each estimated parameter, calculated by the variance in (12). We exploit this reliability metric in order to identify and treat poor estimations. Recalling from section III, the confidence interval is a range of values around the estimated one, which contains the true value with a certain degree of confidence; the narrower that range is, the more reliable the estimation can be considered, and hence, the more accurate. Fig. 16 plots the localization errors for all 550 estimations vs the confidence intervals of their parameters x_{tag} , y_{tag} and z_{tag} ; large values of CI indicate poor confidence and vice versa.

Figs 16 (a) and (b) show that $CI_{x_{tag}}$ and $CI_{y_{tag}}$ demonstrate the expected behavior; increased values of $CI_{x_{tag}}$ or $CI_{y_{tag}}$ indicate poor estimations. On the contrary, the corre-

sponding confidence of z_{tag} parameter, $CI_{z_{tag}}$, is not a reliable indicator of poorly estimated positions (see Fig. 16 (c)), since there is no variation of the antenna's trajectory along the z-axis; i.e. the antenna's height on top of the robot is fixed. For this purpose, we will exploit the sum $CI_{x_{tag}} + CI_{y_{tag}}$ as shown in Fig. 17 (a), such that estimations whose $CI_{x_{tag}} + CI_{y_{tag}}$ is higher than a threshold, will be considered unreliable and treated accordingly. The 99th and 95th percentiles of variable $CI_{x_{tag}} + CI_{y_{tag}}$ were chosen and tested as values for the threshold; the latter are represented in Fig. 17 (a) by the red dashed vertical lines.

Those tags with confidence interval greater than the threshold are then treated as explained next. Phase ReLock 3D is re-applied for all possible antenna-combinations. Since there are 4 available antennas the remaining possible combinations are 10; i.e. antenna-pairs: 1-2-3, 1-2-4, 1-3-4, 2-3-4, 1-2, 1-3, etc. Having now 11 different estimations (including the initial estimation where all 4 antennas were accounted), we select the one with the smallest $CI_{x_{tag}} + CI_{y_{tag}}$. Figs 17 (b) and (c) depict the correction of the estimations for the two applied thresholds, respectively. Notice that each corrected estimation lay on the southwest of the respective initial one, indicating that both localization error and confidence were improved.

Last, Table IV concludes the results of Fig. 17. When the

threshold of 73.95cm (99^{th} percentile) was utilized, 6 tags were identified as poorly located and after the repetition of the optimization for all antenna combinations, the achieved mean error was improved by 60cm . Similarly, for a threshold-value of 14.04cm (95^{th} percentile), 32 unreliable estimations were identified and an improvement of 40cm was recorded.

TABLE IV: Correction of poor estimations

$CI_{x_{tag}} + CI_{y_{tag}}$ threshold			mean error (cm)	
percentile	value	tags	initial	final
99^{th}	72.95	6	96.87	36.88
95^{th}	14.04	32	74.59	37.6

Besides the poor estimations that were successfully identified and corrected, there are others which were not recognized by the variance-threshold, as shown in the northwest corner of Fig. 17 (a). These samples demonstrate estimations with large localization error (about 1m) accompanied by high confidence (i.e. small variance). This effect is mainly resulted by the quality of the collected measurements from the specific tags. In particular, the presence of a reflector in the propagation environment (e.g. a wall behind the tags), could contribute a strong field's component throughout the whole set of measurements (i.e. in all positions of the robot). The strong propagation ray that corresponds to such reflector will be summed to the direct (LOS) ray, affecting all phase measurements in the same manner. As a result, the method would estimate an erroneous location with high confidence. Despite the fact that the experiments were conducted inside a multipath-rich indoor environment, the number of such bad estimations was quite small. It is worth noting that even if Phase ReLock was re-applied for all possible antenna-pairs for those tags, the estimation of their locations would not have been corrected.

V. CONCLUSION

In this paper we have proposed a novel method for 3D localization of UHF RFID tags by a SLAM-enabled robot. Four reader's antennas collect phase measurements. Initially, the phase-measurements are unwrapped. Localization is accomplished by solving the proposed optimization function, through a rapid iterative convergence to its global minimum. The proposed method exploits the ensemble of the collected measurements from all antennas and derives the 3D coordinates of the tag that best match the measurements. As a result, the proposed method is accurate and fast. Furthermore, a new metric for the identification of measurements that reduce the estimations' accuracy is proposed. The metric is quantified by the size of the 95% confidence intervals around the estimated values of the problem's parameters.

Experimental comparison with prior art confirms the superiority of the proposed method with respect to accuracy and computation time. Aggregate results from the experiments indicate a mean expected 3D error less than 20cm , while estimations are derived in 170ms per tag; i.e. a huge improvement compared to prior art.

The high localization accuracy and computational speed of Phase ReLock, in combination with the advantages of a self-navigating ground robot, make the proposed method

beneficial for many practical applications. Among others, it can be applied in logistics management, 24/7 inventorying and localization of goods in large warehouses, instant tracking of misplaced products in retail stores, book management in smart libraries, location-aware games between young visitors and social robot inside a museum, identification of targets for interaction, etc.

However, since in real scenarios the number of tagged objects and hence, the tag population within the reader's reading range, can be uncontrollably huge (e.g. tiny products on shelves of retail store), the read-rate of the employed reader(s) needs to be increased, in order to obtain a fair collection of measurements from all tags. To address that, multiple readers need to be installed, operating at distant central frequencies, in order to ensure a large frequency separation in the front-end electronics of the readers. The read-rate per target volume, would be multiplied by the number of readers. As RFID technology continues to penetrate the market replacing traditional barcode-tagging, additional bandwidth is expected to be given to the technology and possible increase of the reader's read-rate.

Adaptation of the robot's moving strategy can also improve the performance. By adjusting the speed of the robot, such that it travels fast in regions of the environment with few tags and quite slow in tag-crowded regions, a rich set of measurements can be preserved for all tags, while the time required for the robot to traverse the spaces of interest and inquire all tags is optimized.

REFERENCES


- [1] P. V. Nikitin, R. Martinez, S. Ramamurthy, H. Leland, G. Spiess, and K. V. S. Rao, "Phase Based Spatial Identification of UHF RFID Tags," *2010 IEEE International Conference on RFID*, Orlando, Florida, 2010.
- [2] J. Zhou, H. Zhang, and L. Mo, "Two-dimension Localization of Passive RFID Tags Using AOA Estimation," *2011 IEEE Instrumentation and Measurement Technology Conference (I2MTC)*, Binjiang, China, 2011.
- [3] S. Azzouzi, M. Cremer, U. Dettmar, R. Kronberger, and T. Knie, "New Measurement Results for the Localization of UHF RFID Transponders Using an Angle of Arrival (AoA) Approach," *2011 IEEE International Conference on RFID*, Orlando (FL), 2011.
- [4] S. Subedi, E. Pauls, and Y. D. Zhang, "Accurate Localization and Tracking of a Passive RFID Reader based on RSSI Measurements," *IEEE Journal of Radio Frequency Identification*, vol. 1, no. 2, pp. 144-54, 2017.
- [5] J. Zhang, Y. Lyu, J. Patton, S.C. G. Periaswamy, and T. Roppel, "BFVP: A Probabilistic UHF RFID Tag Localization Algorithm Using Bayesian Filter and a Variable Power RFID Model," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 10, pp. 8250-8259, 2018.
- [6] P. Yang, and W. Wu, "Efficient Particle Filter Localization Algorithm in Dense Passive RFID Tag Environment," *IEEE Transactions on Industrial Electronics*, vol. 61, no. 10, pp. 5641-5651, 2014.
- [7] A. Almaaitah, K. Ali, H. S. Hassanein, and M. Ibnkahla, "3D passive tag localization schemes for indoor RFID applications," *2010 IEEE International Conference on Communications*, pp. 1-5, 2010.
- [8] F. Tili, N. Hamdi, and A. Belghith, "Accurate 3D localization scheme based on active RFID tags for indoor environment," *IEEE RFID-TA*, 2012, pp. 378-382.
- [9] T. Liu, L. Yang, Q. Lin, Y. Guo, and Y. Liu, "Anchor-free backscatter positioning for RFID tags with high accuracy," *IEEE INFOCOM*, 2014, pp. 379-387.
- [10] R. Miesen, F. Kirsch, and M. Vossiek, "UHF RFID localization based on synthetic apertures," *IEEE Transactions on Automation Science and Engineering*, vol. 10, no. 3, pp. 807-815, 2013.

- [11] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu, "Tagoram: real-time tracking of mobile rfid tags to high precision using cots devices," In Proceedings of the 20th annual international conference on Mobile computing and networking, pp. 237-248, 2014.
- [12] Y. Zhang, L. Xie, Y. Bu, Y. Wang, J. Wu, and S. Lu, "3-Dimensional Localization via RFID Tag Array," 2017 IEEE 14th International Conference on Mobile Ad Hoc and Sensor Systems, 2017
- [13] L. Qiu, Z. Huang, N. Wirstrom, and T. Voigt, "3DinSAR: Object 3D Localization for Indoor RFID Applications," IEEE International Conference on RFID (RFID), 2016.
- [14] A. Motroni, P. Nepa, P. Tripicchio, M. Unetti, "A Multi-Antenna SAR-based method for UHF RFID Tag Localization via UGV", 2018 IEEE International Conference on RFID Technology & Application (RFID-TA), Macau, China, 2018.
- [15] L. Shangguan and K. Jamieson, "The design and implementation of a mobile rfid tag sorting robot," In Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services, pp. 31-42, 2016.
- [16] E. DiGiampaolo, F. Martinelli, "A Robotic System for Localization of Passive UHF-RFID Tagged Objects on Shelves," IEEE Sensors Journal, vol. 18, no. 20, pp. 8558-8568, 2018.
- [17] F. Bernardini et al., "Particle Swarm Optimization in SAR-Based Method Enabling Real-Time 3D Positioning of UHF-RFID Tags," in IEEE Journal of Radio Frequency Identification, vol. 4, no. 4, pp. 300-313, Dec. 2020, doi: 10.1109/JRFID.2020.3005351.
- [18] E. DiGiampaolo and F. Martinelli, "Range and Bearing Estimation of an UHF-RFID Tag Using the Phase of the Backscattered Signal," in IEEE Journal of Radio Frequency Identification, vol. 4, no. 4, pp. 332-342, Dec. 2020, doi: 10.1109/JRFID.2020.3016168.
- [19] L. M. Ni and Y. Liu, "LANDMARC: indoor location sensing using active RFID," Wireless Networks, 10(6), pp. 701-10, 2004.
- [20] Y. Zhao, Y. Liu, and L. M. Ni, "Vire: Active rfid-based localization using virtual reference elimination," Proc. of IEEE ICCP, 2007.
- [21] S. Megalou, A. Tzitzis, S. Siachalou, T. Yioultsis, J. Sahalos, E. Tsardoulas, A. Filotheou, A. Symeonidis, L. Petrou, A. Bletsas, A. G. Dimitriou, "Fingerprinting Localization of RFID tags with Real-Time Performance-Assessment, using a Moving Robot," 13th European Conference on Antennas and Propagation, Krakow, Poland, March 2019.
- [22] S. Siachalou, S. Megalou, A. Tzitzis, E. Tsardoulas, A. Bletsas, J. Sahalos, T. Yioultsis, A. G. Dimitriou, "Robotic Inventorying and Localization of RFID Tags, Exploiting Phase-Fingerprinting," 2019 IEEE International Conference on RFID-Technology and Applications, RFID-TA 2019, Pisa, Italy, September 2019.
- [23] J. Wang and D. Katabi, "Dude, where's my card?: RFID positioning that works with multipath and non-line of sight," Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM, Hong kong, China, pp. 51-62, 2013.
- [24] A. Tzitzis et al., "Localization of RFID Tags by a Moving Robot, via Phase Unwrapping and Non-Linear Optimization," in IEEE Journal of Radio Frequency Identification, vol. 3, no. 4, pp. 216-226, Dec. 2019, doi: 10.1109/JRFID.2019.2936969.
- [25] A. Tzitzis et al., "Trajectory Planning of a Moving Robot Empowers 3D Localization of RFID Tags with a Single Antenna," in IEEE Journal of Radio Frequency Identification, doi: 10.1109/JRFID.2020.3000332.
- [26] A. Tzitzis, S. Megalou, S. Siachalou, E. Tsardoulas, T. Yioultsis and A. G. Dimitriou, "3D Localization of RFID Tags with a Single Antenna by a Moving Robot and "Phase ReLock"," 2019 IEEE International Conference on RFID Technology and Applications (RFID-TA), Pisa, Italy, 2019, pp. 273-278, doi: 10.1109/RFID-TA.2019.8892256.
- [27] A. Tzitzis et al., "Real-time 3D localization of RFID-tagged products by ground robots and drones with commercial off-the-shelf RFID equipment: Challenges and Solutions," 2020 IEEE International Conference on RFID (RFID), Orlando, FL, USA, 2020, pp. 1-8, doi: 10.1109/RFID49298.2020.9244904.
- [28] M. A. Branch, T. F. Coleman, and Y. Li, "A Subspace, Interior, and Conjugate Gradient Method for Large-Scale Bound-Constrained Minimization Problems," SIAM Journal on Scientific Computing, vol. 21, no. 1, pp. 1-23, 1999.
- [29] R. H. Byrd, R. B. Schnabel, and G. A. Shultz, "Approximate Solution of the Trust Region Problem by Minimization over Two-Dimensional Subspaces," Mathematical Programming, vol. 40, pp. 247-263, 1988.
- [30] J. J. More and D. C. Sorensen, "Computing a Trust Region Step," SIAM Journal on Scientific and Statistical Computing, vol. 3, pp. 553-572, 1983.
- [31] Sebastian Thrun, Wolfram Burgard, and Dieter Fox, "Probabilistic Robotics" (Intelligent Robotics and Autonomous Agents), The MIT Press, 2005

- [32] F. Dellaert, D. Fox, W. Burgard and S. Thrun, "Monte Carlo localization for mobile robots," Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C), Detroit, MI, USA, 1999, pp. 1322-1328 Volume 2, doi: 10.1109/ROBOT.1999.772544
- [33] P. Maybeck, "Stochastic Models, Estimation and Control", Volume 1, Academic Press, New York, 1979



Anastasios Tzitzis was born in Thessaloniki, Greece, in 1994. In 2018 he received the Diploma in Electrical and Computer Engineering from Aristotle University of Thessaloniki, where he is currently working toward the Ph.D. degree. At the same time, he is working as a Research and Teaching Assistant at the Aristotle University. His current research interests include analysis and design of antennas, RFID technology and wave propagation.




Aristidis Raptopoulos Chatzistefanou was born in Florina, Greece, in 1996. He received the Diploma degree in electrical and computer engineering from the Aristotle University of Thessaloniki in 2019, where he is currently pursuing the Ph.D. degree, and working as a Research Assistant with the Aristotle University of Thessaloniki. His main research interests include RFID technology, localization and tracking techniques.



Traianos V. Yioultsis (M'09) received the Diploma and the Ph.D. degrees in electrical and computer engineering from the Aristotle University of Thessaloniki, Greece, in 1992 and 1998, respectively. From 2001 to 2002, he was a Post-Doctoral Research Associate with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign. Since 2002, he has been with the Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, where he is currently a Professor. His current interests

include the analysis and design of antennas and microwave circuits with fast computational techniques, and the modeling of complex wave propagation problems. He has also served as a member of the Editorial Board for IEEE Communications Letters and several international conferences.



Antonis G. Dimitriou (S'01-M06-SM'14) received the diploma and the Ph.D. degree in Electrical and Computer Engineering from the Aristotle University of Thessaloniki, Greece, in 2001, and 2006 respectively. Since 2007, he is with the School of Electrical and Computer Engineering of AUTH, where he currently serves as a teaching and research faculty member.

He has participated in more than 20 research projects, 8 of which since 2007 as a principal investigator in the fields of Robotics, RFIDs, and Wireless Sensor Networks. He is currently the coordinator of project "RELIEF", involving continuous RFID inventorying through robotics and project "CULTUREID", where RFID equipment will be installed inside the Archaeological Museum of Thessaloniki to monitor RFID tagged exhibits and track visitors. He was a Management Committee Member in the ICT COST Action IC301 "Wireless Power Transmission for Sustainable Electronics (WiPE)". He is the author or co-author of approximately 60 journal and conference papers.

Dr. Dimitriou was the recipient of the Ericsson Award of Excellence in Telecommunications in 2001 and co-recipient of the student-paper award in the 2011 IEEE RFID-TA conference. He received the "IEEE Wireless Communications Letters Exemplary Reviewer" award in 2012 and 2014. He is a Senior IEEE Member since February 2014. He also serves as a reviewer for major journals and as a TPC member for international conferences.