

LocAP: Autonomous Millimeter Accurate Mapping of WiFi Infrastructure

Roshan Ayyalasomayajula, Aditya Arun, Chenfeng Wu, Shrivatsan Rajagopalan, Shreya Ganesaraman, Aravind Seetharaman, Ish Kumar Jain, and Dinesh Bharadia (roshana, aaron, chw357, s1rajago, sganesar, arseetha, ikjain, dineshb)@ucsd.edu
University of California, San Diego

Abstract

Indoor localization has been studied for nearly two decades fueled by wide interest in indoor navigation, achieving the necessary decimeter-level accuracy. However, there are no real-world deployments of WiFi-based user localization algorithms, primarily because these algorithms are infrastructure dependent and therefore assume the location of the access points, their antenna geometries, and deployment orientations in the physical map. In the real world, such detailed knowledge of the location attributes of the access Point is seldom available, thereby making WiFi localization hard to deploy. In this paper, for the first time, we establish the accuracy requirements for the location attributes of access points to achieve decimeter level user localization accuracy. Surprisingly, these requirements for antenna geometries and deployment orientation are very stringent, requiring millimeter level and sub- 10° of accuracy respectively, which is hard to achieve with manual effort. To ease the deployment of real-world WiFi localization, we present LocAP, which is an autonomous system to physically map the environment and accurately locate the attributes of existing wireless infrastructure in the physical space down to the required stringent accuracy of 3 mm antenna separation and 3° deployment orientation median errors, whereas state-of-the-art algorithm reports 150 mm and 25° respectively.

1 Introduction

Indoor navigation requires precise indoor maps and accurate user location in these maps. Google, Bing, Apple or Open Street Maps have made considerable progress towards providing precise indoor maps for notable locations like airports and shopping malls [1–4]. On the other hand, there are two decades of research on indoor localization using WiFi infrastructure that achieve decimeter accurate user locations [22, 30, 37, 39, 50, 51, 53, 58, 59, 65–69]. Despite these innovations, we still cannot use our smartphones to navigate in these indoor environments.

The key reason for this inability is the absence of the bridge

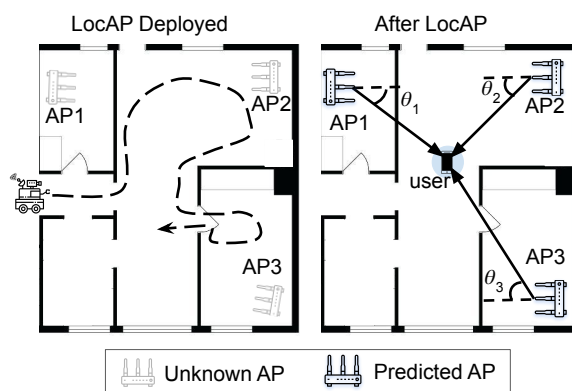


Figure 1: **Implementation of LocAP:** (Left) An unknown environment with unknown AP attributes where LocAP is deployed. (Right) LocAP once deployed determines the AP attributes in the physical map enabling triangulation based user localization.

providing the context of the physical map to the user locations. While there is recent work [8] that bridges this gap, it, like other state-of-the-art localization algorithms [37, 59, 66], is dependant on the accurate location attributes of the WiFi access points (APs) in the physical maps of these airports and malls. To understand what we mean by location attributes, consider the setup shown in Figure 1(right). The smartphone user is triangulated in an indoor environment by estimating the angle subtended by the user at each of the access points. This approach inherently assumes to have accurate knowledge of each access point’s location and its deployment orientation (the angle at which the access point is placed in the given physical map). Further, to estimate the angle made by the user with respect to an access point, the channel state information (CSI) based WiFi localization algorithms need to know the exact antenna placements on these access points.

One can endeavor to manually locate each of these access points in the environment, but that would be labor-intensive, time-consuming and even impossible sometimes because of the following reasons. First, these access points (AP) are

usually not easily visible; they may be located behind a wall or pillar. Second, even if the AP is visible, most of the access points are encased by the manufacturer, making it difficult to know the exact information of the antenna placements on the access point. Third and finally, even if we can estimate the antenna placements on the access point from the datasheet provided by the manufacturer¹, the AP's *deployment orientation* has to be carefully calibrated to the indoor maps within an error of a few degrees. Thus, we need a system that can help in accurate mapping of the existing WiFi infrastructure, which does not involve any manual labor or time.

In this paper, we present LocAP, an autonomous and accurate system to estimate access point location attributes – access point location, antenna placements, and deployment orientation. We call this process of predicting accurate access point attributes as *reverse localization*. LocAP is the first work to establish the requirements for reverse localization as follows:

Accurate Access Point Locations: As shown in Figure 2a, any error in AP location is translated to an error in the location of the user. So, any error exceeding a few tens of centimeters in access points' location is going to adversely affect the decimeter-level user localization. Thus, LocAP needs to locate the access point accurate to within tens of centimeters.

Accurate Antenna Separation: Different APs have different antenna placement configurations and the angle made by the user is measured at the access point using the spacing between antennas. So, any error in measuring antenna placements is going to cause a rotation error at the user. For example, error in antenna separation by 4 mm causes 12° of error in the angle of user measured at the access point, which translates to up to 1 m of error for a user 5 m away from the access point. Thus, LocAP needs to predict the antenna separation accurately to within a few millimeters.

Accurate Deployment Orientation: Finally, the access points can be placed in any orientation in the environment. Any error in measurement of orientation directly translates to the predicted angle subtended by the user at the access point. Hence even 10° of error in deployment orientation causes up to 90 cm of user location error for a user located just 5 m away from the access point. Thus, LocAP should resolve the deployment orientation of the access point accurate to less than 10° of error.

Automation: LocAP's goal is to require no manual effort for the reverse localization, and achieve the stringent requirements discussed earlier. Furthermore, there should be zero effort to associate these positions with the existing indoor maps, ideally in an autonomous way.

LocAP achieves the aforementioned requirements and enables automated and accurate *reverse localization* of the access points. We achieve autonomy by deploying LocAP on a bot retrofitted with a multi-antenna WiFi device used in

¹Datasheets, though publicly available do not talk about antenna placements or dimensions [6, 7, 17, 28, 47, 48, 57].




Figure 2: **LocAP's Motivation:** The user location is predicted wrong due to different errors in access point's estimated details. (a) **Translation:** Predicting the wrong location of the AP. (b) **Dilation:** Predicting wrong antenna separation on the access point results in an error in angle estimated, ($\theta_{obs} \neq \theta_{exp}$) of the user. (c) **Rotation:** Predicting the wrong orientation of the AP.

[8]. When deployed in a new environment, the bot first maps the physical environment. Next, it associates with existing AP infrastructure by collecting multi-antenna channel state information, and pairing it with its predicted location in the physical map as shown in Figure 1(left). LocAP uses this information to build a database of the deployed WiFi infrastructure consisting of all the access point attributes meeting our stringent accuracy requirements. This database of accurate AP location attributes can then be used for decimeter level user localization as depicted in Fig. 1(right)

The main technical contributions of LocAP to achieve the above requirements can be summarized as follows:

cm-accurate Access Point Localization: We make an important observation that accuracy of triangulation based WiFi-localization methods improves with an increasing number of anchor points with known locations. In essence, creating an array of 100's of antennas measuring CSI at known locations achieves cm-level localization, which is not feasible in practice². To overcome this, LocAP leverages the CSI data collected by the bot at 100's of predicted locations, mimicking 100's of virtual antennas with known locations. However, these *predicted* locations suffer from a varying amount of inaccuracy. Hence, LocAP designs a weighted localization algorithm, which weights each location-CSI data-point with a uniquely defined confidence metric capturing the accuracy of the predicted location.

mm-accurate Antenna Geometry Localization: We have seen earlier that both mm-error in antenna separation and error in deployment orientation lead to in-accurate Angle of Arrival (AoA) measurement at the access point, which impedes user-triangulation. Thus, LocAP tackles antenna separation and deployment orientation together by achieving millimeter-

²typical indoor settings are 1000-2000 sq. ft., which would imply deploying an antenna every 100 sq. ft.

level accuracy in predicting the antenna geometry. The first thought would be to use 1000's of virtual antennas to achieve cm-accurate localization [39] by locating individual antenna geometry on the AP. But, this idea can only achieve accuracy at the cm-level and will not suffice to achieve mm-level details of the antenna geometry. Our key observation is to localize the relative antenna geometry between two antennas, primarily because the relative wireless channel between the two antennas can be measured very accurately by measuring their phase information. The phase information is measured at the carrier frequency level ($\lambda=60$ mm equivalent to 360°), hence even phase measurement accurate to 10's of degrees achieves 1-2 mm accuracy. However, this works for only relative antenna separation $d < \frac{\lambda}{2}$. LocAP designs a novel algorithm that uses relative channel information across multiple bot locations to solve for any antenna geometry, unrestricted by antenna separation, to mm-level accuracy.

Automation – Augmenting the SLAM algorithms: To avoid any manual labor and errors, LocAP is deployed on a SLAM (Simultaneous Localization and Mapping) based autonomous bot developed by us [8]. This bot provides us with a physical map and the location and heading of the bot in this physical map at all times. We pair these location-heading measurements with the CSI collected by the mounted WiFi device. However, even the best of SLAM algorithms report the location to be in-accurate up to 10-20 cm, which can have a detrimental effect on the AP location attributes. Therefore, LocAP develops a confidence metric whose core idea to look at the covariance of measurements across consecutive frames.

Further, the implementation of LocAP does not need any modification at the existing access points, as it is deployed on a custom made bot [8] that is mounted with a Quantenna client. The Quantenna client readily reports the channel-state-information (CSI) of the associated access point. We evaluate LocAP in an indoor environment of 1000 sq ft area with multiple off-the-shelf access points and 2 different antenna configurations – rectangular and linear.³ We achieved the following results satisfying the aforementioned accuracy requirements:

Relative Antenna Geometry Prediction: LocAP's relative reverse localization for the antenna separation has a median error of 3 mm ($50\times$ improvement), and a median error of 3° ($8\times$ improvement) for deployment orientation, while state-of-the-art achieves a median error of 150 mm and 25° respectively.

Access Point Localization: LocAP's reverse localization of the access points achieves a median localization error of 13.5 cm improving by 35% over the state-of-the-art WiFi localization algorithms [37].

Case Study-User Localization: State-of-Art user localization is deployed using the access point attributes measured manually and with LocAP. We observe user localization errors of 78 cm and 50 cm respectively, a decrease in the error of about 36%.

³these configurations generalize the more generic antenna deployments found on the commercial off the shelf WiFi access points.

2 Requirement and Motivation

It may seem natural that user localization algorithms [37, 53, 59, 67] could be sufficient for *reverse localizing* the access point's location attributes – location, antenna geometry and deployment orientation. Surprisingly, it turns out that requirements for reverse localization of the access are stringent. To define these requirements, we conduct empirical evaluations from the standpoint on how various errors in AP attributes adversely affect the state-of-the-art decimeter level localization algorithms.

Our empirical setup contains four access points, each with 4 antennas, setup in a 25ft \times 30ft space. The user device is placed at 100 different locations while the access points locate the user using an algorithm similar to [37]. Specifically, we aim to achieve decimeter-level localization accuracies for user WiFi localization algorithms and thus set a hardbound that no more than 50 cm median error for user localization can be tolerated.

Error in the AP's location Firstly, in the above-described setup, we incrementally increase the error in all the access points' locations. Next, we estimate the user location for each of these erroneous access point locations and calculate the user localization error. In Figure 3a, we plot the median user localization error across the access point errors reported. We can see that if the access point locations have an error of more than a few centimeters, the median localization error starts to increase. From this, we can infer that the required level of accuracy for the reverse localization of APs should be in the order of centimeters.

Error in the antenna separation Second, AoA based localization algorithms make use of the relative phase information between two antennas. Earlier, we have seen that the relative antenna position has to be estimated accurately to have exact measurements of angles. Even when the access point positions are reported correctly, we can observe that the localization error increases with just a few millimeters of errors in the reported relative antenna positions as shown in Figure 3b. This observation is intuitive because the relative antenna distances are usually of the order of a wavelength of the transmitted signal, which in the case of WiFi is 6cm. So, any error which is greater than a few millimeters is going to make a huge difference in the relative phase measured at the access point.

Error in the Deployment Orientation Finally, the antenna array can be oriented in any direction. It is also important to know the exact deployment orientation of the antenna array. Errors in this orientation will proportionately affect the angle of arrival measurements made at the access points. We observe that the greater the error in deployment orientation prediction, the higher the median localization error becomes as shown in Figure 3c. From this plot, we can see that even 7° will degrade the median user localization accuracy to more than 50 cm.




Figure 3: **Robustness of localization accuracy to Access Point (AP) location errors:**(a) Shows that median localization error increases with increase in error of estimated AP location.(b) Shows how median localization error increases with increase in error of estimated antenna locations. (c) Shows how median localization error increases with increase in error of estimated antenna deployment orientation.

In summary, we should locate the access point’s location with less than 30 cm of error, the antenna separation within 5 mm of error and the deployment orientation to less than 7° of error. While these locations are typically mapped manually by humans using specialized equipment like VICON [55] or laser-based range finders [11], this process is time-consuming, labor-intensive and error-prone. So, we need a system that can accurately localize access points attributes satisfying these stringent requirements. Note that the most stringent requirements are the mm-accurate antenna separation and sub-7 degree deployment orientation. The state-of-the-art [37, 39, 53, 67] localization algorithms can locate the individual antennas to within a few 10 centimeters even by deploying hundreds of AP’s in a given environment, which is insufficient to determine the antenna geometry as per required specifications established earlier in this section. Further, there are relative localization algorithms [38, 61, 64] which track the user’s location across contiguous observations few millimeters apart. These ideas could potentially be used to find the relative antenna geometry. But, these tracking algorithms assume that the two relative locations are less than $\lambda/2$ apart [6, 7, 17, 28, 47, 48, 57] but the antenna separations on most access points are more than $\lambda/2$ apart, where there is an ambiguity that cannot be resolved. So, we design a system, LocAP, which fulfills these requirements and locates the access points and their antennas with the desired level of accuracy

3 Design

In this section, we present the design of LocAP. Recall that our main goal is to autonomously determine access points’ location attributes within the reference coordinates of the physical map to enable easily deploy-able WiFi-localization. LocAP deploys a SLAM based autonomous bot developed in [8] to map the environment. The autonomous bot provides it’s location and heading with respect to the environment’s map. Simultaneously, a four antenna WiFi device retrofitted

on the bot, connects with the existing WiFi infrastructure, all the while reporting the CSI information at each instance. Furthermore, to avoid changes to deployed AP infrastructure, we perform all the processing on the bot. LocAP, therefore, is provided with the location and orientation of the bot with respect to the physical map and the CSI data from the WiFi device on the bot, which connects with the existing WiFi infrastructure. We design LocAP to use these inputs to provide accurate access point attributes—location, antenna separation, and deployment orientation with respect to the physical map.

First, we discuss how to achieve the cm-level accurate location of the AP that also accounts for inaccuracies in reported bot poses. Second, we present LocAP’s algorithm to estimate the antenna separation and deployment orientation of all the APs that needs to achieve the stringent requirement of mm-level accuracy. In both of these scenarios, we assume we have the CSI corresponding to the direct path and later in Section 3.3 we discuss how we tackle the presence of multipath in the environment and recover the direct path’s CSI. Finally, we present the SLAM-based bot design, which does the best effort to provide the necessary measurements mentioned above. But often, these measured poses are not accurate. So, LocAP builds an algorithm which reports a confidence metric for each measured pose. This confidence metric helps us surmount the errors in the bot locations to calculate AP location attributes.

3.1 Locating the Access Point

In this subsection, we focus on identifying the position of one of the access point’s antenna. This position of the antenna would then be representative of the access point’s location and we refer to this as the first antenna in the subsequent text. Recall that the access point’s location has to be estimated accurately to cm-level. A simple solution can be to utilize the existing WiFi localization approaches to locate one of the antennas on the access point, which would then become the access point’s location. Unfortunately, state-of-




Figure 4: **First Antenna Localization:** Gives an overview of how triangulation from 10s of bot locations locates the access point accurately to within few centimeters.

the-art localization algorithms only report decimeter level location estimates. However, we make an interesting observation: these algorithms show increasing location accuracies with an increase in the number of access points deployed in an environment. In our scenario, we have a mobile bot which collects CSI data from the deployed access point at multiple anchors. This bot covers a large area setting up 100's of anchors which aids in cm-accurate first antenna localization.

Owing to this setup of LocAP, we can employ an angle of arrival estimation algorithm similar to [37] and estimate the direct path's AoA, α_p^{bot} for p^{th} bot location ($p = 1, 2, \dots, P$). We measure these AoA's with respect to the bot's local axis ($X'-Y'$) corresponding to the first antenna's transmission for each bot location $\mathbf{u}_p = [u_p, v_p]$. To enable this AoA based first antenna triangulation, we should also know the direction of the bot's heading (θ_p) with respect to the global axis (denoted by $X-Y$ in Figure 4), which is reported by the bot as mentioned earlier. With $(\mathbf{u}_p, \theta_p$ and $\alpha_p^{bot})$ we can find the first antenna location as an intersection of P lines:

$$\text{Line}_p \equiv (y_1 - v_p) = \tan(90^\circ - (\alpha_p^{bot} + \theta_p))(x_1 - u_p) \quad (1)$$

Ironically, the AoA based triangulation accuracy is bounded by the errors in the bot's reports of its location, (u_p, v_p) and heading, θ_p^{bot} . Clearly, this creates a vicious unending loop – to predict the antenna locations we need accurate bot measurements and vice-versa to predict the bot's locations. To overcome this problem, we take advantage of SLAM algorithms [21] to get accurate ground truth estimates of the bot location and heading. Unfortunately, SLAM-based bots do not have 100% confidence in all the location estimates they report, forcing us to only cherry-pick the measurements which we believe are accurate. Based on this intuition, we design a confidence metric, $w_p \in [0, 1]$ for each bot location \mathbf{u}_p . Further details on the design of the confidence metric are discussed in Section 3.4. This confidence metric implies that the bot is more confident with the reported pose the closer it is to one. We thus implement a low-confidence rejection algorithm, which rejects the measurements with confidences, w_p , in the lowest 20% (Using only $[0.8 \times P]$ lines).

We use these confidences in combination with the rest of our $[0.8 \times P]$ line equations to define a weighted least squares problem to optimally solve for the first antenna location as

follows:

$$\min_{\mathbf{x}_1} \|W(S\mathbf{x}_1 - \mathbf{t})\|^2 \quad (2)$$

where $\mathbf{x}_1 = [x \ y]^T$ is the first antenna location, $W = \text{diag}(w_1, w_2, \dots, w_{0.8P})$ is the weight matrix, $S(p, :) = [\cos(\alpha_p^{bot} + \theta_p) \ -\sin(\alpha_p^{bot} + \theta_p)]^T$ and $\mathbf{t}(p) = [u_p \cos(\alpha_p^{bot} + \theta_p) - v_p \sin(\alpha_p^{bot} + \theta_p)]$. Thus, we estimate of the first antenna's location \mathbf{x}_1 which corresponds to the access points location.

3.2 Determining Antenna Separation and Deployment orientation

As described above, we can leverage the motion of the bot to identify the accurate location of one antenna on the access point. One might wonder if it is possible to apply this algorithm iteratively to identify the location of each antenna on the access point and hence recover the relative placement of antennas. However, it is not so straightforward. In particular, the geometry prediction needs to be an order of magnitude more accurate than the location prediction. While it suffices to measure the location of the access point to cm-level, the geometry, i.e. the relative position of antennas, needs to be mm-accurate. While combining across 10s of bot locations provides antenna location accurate to cm-level, it does not extend to mm-accurate antenna geometry by combining across 100s or even 1000s of bot locations as shown in the prior art [39]. This problem occurs owing to the asynchronous clocks between the access point and the bot's WiFi device when measured at a single antenna at the access point.

To overcome this problem we make a key observation - in contrast to the phase measured at one antenna on the access point, the relative phase across two antennas is rid of synchronization errors as they share the same clock. Further, at WiFi 11ac's 5GHz carrier frequency, a wavelength of 6 cm corresponds to a phase difference of 2π radians. Empirically, we have observed that we can easily resolve phase differences up to $\pi/18$ radians (10°), which facilitates measurement of the distance between two antennas with a resolution of 2 mm, thus enabling us to locate the antenna geometry accurately to within few millimeters. Hence, our first key insight is to measure the relative antenna separations, d_i , and deployment orientations, ψ_i , for all the N_{AP} antennas on the access point with respect to the first antenna ($i = 2, 3, \dots, N_{AP}$).

Unfortunately, although the relative phase information can resolve relative antenna separation to within 2mm, it cannot resolve for antenna separations greater than $\lambda/2$. To further understand this, consider an example scenario where the bot is moving in a circular arc about the two-antenna access point in steps of small angles as shown in Figure 5a. To avoid overcrowding of subscripts, we consider a two antenna access point and drop the access point's antenna indexing, i . Similar analysis can be performed pairwise on all the antennas on the access point with respect to the first antenna. Now, to




Figure 5: **Estimating AoD from phase difference:** (a) A sample case where the bot is in circular arc around the AP (b) Phase difference $\Delta\phi$ vs the orientation of the bot β assuming the deployment point of the AP, $\psi = 0$ (c) Slope $\frac{d\Delta\phi}{d\beta}$ vs the orientation of the bot β when compensated for the orientation of the access point.

locate the second antenna with respect to the first antenna, we analyze relative phase across these two antennas. We know that for the bot's location, \mathbf{u}_p , the phase difference between these two antennas corresponding to the direct path, $\Delta\phi_p$, can be estimated as:

$$\Delta\phi_p = \text{mod} \left(\frac{2\pi d}{\lambda} \sin(90^\circ - (\beta_p - \psi)), 2\pi \right) \quad (3)$$

where, the parameters of interest ψ and d are antenna deployment orientation (with respect to the X-axis) and antenna separation respectively. β_p is the angle subtended by the bot's location at the access point with respect to the global X-axis. From the inset in Figure 5a, we can see that the angle of departure from the AP is given by

$$\alpha_p^{\text{AP}} = 90^\circ - (\beta_p - \psi) \quad (4)$$

and the extra distance travelled (represented by the red-dashed segment) is given by $d \sin(\alpha_p^{\text{AP}})$. This extra distance travelled induces the phase difference given in Equation 3. Thus, the phase difference across two antennas can help us estimate the antenna separation, d and deployment orientation ψ . To better understand this relation, we plot $\Delta\phi_p$ for all the bot locations along the circular arc against the angle subtended by the bot, β_p , for various antenna separations d in Figure 5b. From this plot we can see that for $d \leq \lambda/2$, we have a unique mapping between the phase difference, $\Delta\phi_p$, and the bot's location, but for $d > \lambda/2$ we have ambiguous solutions that prevents us from estimating d and ψ . The ambiguity occurs because the phase difference we measured is a modulus of 2π , which means for a given $\Delta\phi_p$, the actual phase difference can be $2n_p\pi + \Delta\phi_p$, where n_p is any positive integer. This means we have three unknowns, (d, ψ, n_p) to solve for, given a single phase difference value, $\Delta\phi_p$. Furthermore, even for each additional bot location we have a new $\Delta\phi_{p+1}$ estimate, we also add an extra unknown n_{p+1} making it impossible to uniquely solve for d and ψ . LocAP's key insight is that, in contrast to the phase difference $\Delta\phi_p$, the differential phase difference

with respect to the bot's angle at the AP (β_p) for optimally small increments of β_p , has a unique one-to-one mapping as shown in Figure 5c. So, the second key observation we make is that while the phase difference is not uniquely solvable for $d > \lambda/2$, the differential phase difference is uniquely solvable. Intuitively, two close bot positions will have the similar phase wrap-around's, and hence, taking the difference of the phase differences, $\Delta\phi_{p2} - \Delta\phi_{p1}$, can eliminate the ambiguity.

So far we have considered that the bot is moving along a circular trajectory. In fact, LocAP does not restrict the bot's motion to a circular arc and can work with arbitrary motion, as long as the CSI is measured regularly. To understand the exact implementation of LocAP's relative antenna geometry prediction, we consider a more free-flow path as shown in Figure 6. Concretely, determining the relative antenna geometry requires two parameters – the distance between antennas, d , and the deployment orientation of the antenna array, ψ , as can be seen from Figure 6. The bot moves to P distinct locations along a pre-determined trajectory about the AP and collects a series of P CSI measurements, H_p ($p = 1, 2, \dots, P$), while simultaneously reporting the bot's locations, \mathbf{u}_p . The bot makes an angle β_p with respect to the global X-axis. Next, for each position of the bot, \mathbf{u}_p , we evaluate the differential phase difference $\frac{d\Delta\phi_p}{d\beta_p}$ between the two antennas on the access point. Differentiating Equation 3, we get

$$\frac{d\Delta\phi}{d\beta} = -\frac{2\pi d}{\lambda} \cos(90^\circ - (\beta - \psi)) = -\frac{2\pi d}{\lambda} \sin(\beta - \psi) \quad (5)$$

But, for incremental movements of the bot, the differential phase difference in Equation 5 can be approximated as

$$\frac{d\Delta\phi_p}{d\beta_p} \approx \frac{\Delta\phi_{p+1} - \Delta\phi_p}{\beta_{p+1} - \beta_p} \quad (6)$$

The bot traces $P (> 3)$ positions as it moves, which enables us to obtain the solution from an over-determined system of equations, consequently reducing the noise level. Thus achieving highly accurate relative antenna position and orientation,




Figure 6: **Relative Geometry Prediction:** Shows the same setup as in Figure 5a with a two antenna AP making angle ψ with the positive x-axis and the bot moving about the located first antenna of the AP in an arbitrary path.

and thereby achieving millimeter-level accuracy for relative antenna localization. Now to solve for (d, ψ) uniquely as an over-determined system, it is easier to work with Cartesian co-ordinates than polar coordinates. So, we fix the location of the first antenna of the AP, the antenna on the left in Figure 6, as (x_1, y_1) and represent the second antenna (x, y) defined in the global coordinate system as:

$$(x, y) = (x_1 + d \cos(\psi), y_1 + d \sin(\psi))$$

We rewrite Equation (5) in terms of (x, y) as follows:

$$\frac{d\Delta\phi}{d\beta} = \frac{2\pi}{\lambda} [-(x - x_1) \sin(\beta_p) + (y - y_1) \cos(\beta_p)] \quad (7)$$

for $p = 1, 2, \dots, P-1$

Next, we represent these P set of linear equations in matrix-vector form as follows,

$$A \begin{bmatrix} x - x_1 \\ y - y_1 \end{bmatrix} = \mathbf{b} \quad (8)$$

where A is a $(P-1) \times 2$ matrix and \mathbf{b} is a $(P-1)$ sized column vector defined as

$$A(p, :) = [-\sin(\beta_p) \quad \cos(\beta_p)] \quad (9)$$

$$\mathbf{b}(p) = \frac{\lambda}{2\pi} \frac{\Delta\phi_{p+1} - \Delta\phi_p}{\beta_{p+1} - \beta_p}, \quad p = 1, 2, \dots, P-1 \quad (10)$$

We further denote $\mathbf{x} = [x \quad y]^T$ and $\mathbf{x}_1 = [x_1 \quad y_1]^T$. We estimate \mathbf{x} to the following least squares problem:

$$\min_{\mathbf{x}} \|A(\mathbf{x} - \mathbf{x}_1) - \mathbf{b}\|^2 \quad (11)$$

In this way we can uniquely solve for the cartesian coordinates of the second antenna with respect to the first antenna.

Note that the two measurements $\{\beta_p, \Delta\phi_p\}$ and $\{\beta_{p+1}, \Delta\phi_{p+1}\}$ should not be very close to avoid noise amplification. On the other hand, the measurements should not be very

far apart to cause an error in the estimation of the derivative. A large separation between consecutive measurements can increase the phase difference to more than 2π , thus creating discontinuities across the series of P measurements. Our experiments suggest that around 5° of angular separation $(\beta_{p+1} - \beta_p)$ provides the best results for an antenna separation in $d = [0, 4\lambda]$, where $\lambda = 6\text{cm}$ is the minimum wavelength in the 5GHz frequency band. We emphasize the estimated value of ψ will be in the range of $0 \leq \psi \leq \pi$ because the orientation of the antenna array can be defined uniquely in $0 \leq \psi \leq \pi$.

Generalizing Equation 11, we locate the relative location of each antenna on the access point as $\mathbf{x}_i = [x_i \quad y_i]^T$, where, $i = 2, 3, \dots, N_{AP}$, where N_{AP} is the number of antennas on the AP. We finally find the antenna separations as $d_i = \sqrt{(x_i - x_1)^2 + (y_i - y_1)^2}$, and the deployment orientation as $\psi_i = \tan^{-1} \frac{y_i - y_1}{x_i - x_1}$, for all the antennas with respect to the first antenna, \mathbf{x}_1 . Thus, we accurately predict the location, antenna separation and deployment orientation of the access point.

3.3 Multipath

So far, in both Section 3.1 and Section 3.2, we have assumed only one single path from the AP to the bot to solve for the access point attributes. However, the environment creates multipath which would cause the previous algorithms to fail by distorting the phase measurements. We leverage multipath rejection algorithm from [37] to estimate the direction of direct path for AP localization (Section 3.1) and build a novel algorithm to recover direct path phases as required in Section 3.2.

Recall from Section 3.1 that locating the first-antenna on the AP requires direct path AoA information at the bot. However, the received signal at the bot is usually a mix of signals arriving from different directions. We leverage multiple antennas on the bot along with the channel information across multiple subcarriers of the WiFi signal to identify the direct path and isolate it from other paths similar to prior art [37]. As first step, we collect $N_{bot} \times N_{sub}$ CSI-matrix (across N_{bot} bot client's antennas and N_{sub} subcarriers) as shown in Figure 7(a). We then apply 2D-FFT transform to estimate the AoA and Time-of-Flight (ToF) for each arriving path to the bot (Figure 7(b)). Finally, we estimate the direct path AoA by observing the signal, which has the least ToF. Intuitively, the direct path signal travels the shortest distance and thus has the lowest ToF. Thus, we can use these direct path AoA estimates to run our AP localization algorithm, as discussed in Section 3.1.

Note, however, that the direct path AoA information is not enough for estimating AP's antenna geometry (Section 3.2). In this case, our algorithm requires relative phase information across multiple AP antennas corresponding to the direct path signal. Our first insight is to estimate the direct path channel individually for each AP antenna and use them to re-




Figure 7: **Multipath rejection:** (a) Shows the measured $N_{bot} \times N_{sub}$ complex channel matrix. (b) We perform 2D FFT based transform [9] to estimate the 2D AoA-ToF profile within which we identify the direct path as the least ToF path. (c) We then perform a windowing around this peak to obtain direct path filtered AoA-ToF profile. (d) Finally, direct path’s $N_{bot} \times N_{sub}$ complex channel is estimated by performing a 2D-IFFT on the windowed AoA-ToF profile.

cover the relative phase information. We take the $N_{bot} \times N_{sub}$ CSI-matrix for a fixed AP antenna and estimate the AoA-ToF profile using the same procedure as described in the previous paragraph and Figure 7(a),(b). From [37], we know that the direct path signal is concentrated around the first ToF peak (in the AoA-ToF domain). So, our insight is to apply appropriate window function in the AoA-ToF domain to remove the adulteration due to multipath (Figure 7(c)) and use this information to extract the channel corresponding to direct path. Finally, to extract the direct path signal from this windowed AoA-ToF profile, we perform 2D-IFFT on this windowed signal, as shown in Figure 7(d). As we established before, the same process can be repeated for each AP antenna to finally obtain accurate AP antenna geometries, as discussed in Section 3.2.

3.4 Autonomous Bot and Confidence Metrics

In the following section, let us look more closely at the confidence metric we mentioned in Section 3.1. We deploy RevBot largely to automate our data collection pipeline and further implementation details can be found in [8]. The key pieces of data we need to collect are the bot’s pose information (provided by SLAM algorithms), and time-synchronized CSI estimates for each AP in the environment (provided by an onboard access point). Unfortunately, the position and heading reported by SLAM algorithms are not completely error-free, and the measurements can be adversely affected by the movement of the bot and the surroundings resulting in errors from 20-25 cm. These particularly worse, low-confidence measurements, need to be discarded to obtain accurate AP

geometry predictions. But, most SLAM algorithms do not expose the accurate confidences of a particular reported pose. Fortunately, we can manufacture a pseudo-confidence metric by comparing the match of a current measurement with its surroundings. We make these comparisons using 3D pointclouds generated using an RGB-D camera. Pointclouds are to a 3D space what pixels are to a 2D image – each point carries an (x, y, z) coordinate and color information. We make the following observation - by looking at the registration accuracy of the point-clouds generated by consecutive pose measurements, we can estimate the quality of the relative transformation in question.

More concretely, let us consider two consecutive measurement frames \mathcal{F}_i and \mathcal{F}_{i+1} . We determine the relative transformation T_i between the two frames by looking at their pose estimates. Hence, T_i takes us from \mathcal{F}_i to \mathcal{F}_{i+1} . Furthermore, from the RGB-D images captured at these frames, we can generate point-clouds. By applying T_i to the point-cloud from \mathcal{F}_i , we get an estimate of \mathcal{F}_{i+1} and we can stitch these two point-clouds together. If T_i is accurate, then we will get a perfect overlap of these pointclouds over all the points visible in both the frames. Based on this intuition, we use the covariance matrix \mathcal{V}_i as implemented by [16]. Now, this covariance matrix accommodates all six degrees of freedom as found in a 3D environment, three belonging to each direction of translation and three for each axis of rotation, hence $\mathcal{V}_i \in \mathbb{R}^{6 \times 6}$. The first two diagonal elements give us the variance in the x and y position and $\mathcal{V}_i[1, 2]$ gives us the co-variance between x and y . The variance in $(x+y)$ tells us how much wiggle room there is for the pose in question. Hence, the larger the wiggle room, the less confident we are in our poses. Furthermore, we observe that these variances vary in orders of magnitude, and to linearize our confidence metric, we take the log of the variance. We calculate the pseudo-confidence metric for \mathcal{F}_i as

$$C_i = \log(\text{var}(x+y)) \quad (12)$$

$$= \log(\text{var}(x) + \text{var}(y) - 2\text{cov}(x,y)) \quad (13)$$

Finally, we normalize $C_i, \forall i = 1, 2, \dots, P$, between 0 and 1 to determine w_i , which are confidences we use in Equation 2 used to filter out the low confidence bot locations.

4 Micro-benchmarks

Before evaluating LocAP’s performance, we must understand how the error in the ground truth locations reported by the autonomous bot is affecting the algorithm. We have utilized the robot implementation described in [8], while replacing the single antenna client Quantenna platform with a 4 antenna linear array Quantenna station as shown in Figure 8a. For that, we first estimate the bot’s location error and analyze its effects on the accurate prediction of the location of the access point and the relative antenna geometry on the access point.




Figure 8: **Accuracy of the bot's ground truth movement:** (a) The bot we used for our experiments, a Turtlebot-2 equipped with a 4 antenna Quantenna board, LIDAR, RGB-D camera. (b) Depiction on how bot's error can effect the relative antenna localization algorithm.

4.1 Error in Bot's ground truth Location

Since, we are using the same bot setup described in [8] we use the median localization error reported for the bot in their experiments. We can observe that the median error Δr is around 6cm in this case. Further, we study the orientation errors within the same setup. We find that the median error $\Delta\beta$ in orientation is 3° .

Next, we quantify the effect of this error on the accuracy of locating the access point and determining the relative antenna geometry.

4.2 Effects of Bot's Error

First, we estimate the location of the access point. For this step, we use both the bot's location and orientation. Hence, we must look at the errors in both these measurements. We observe that an error of Δr in bot's location error directly corresponds to an error of Δr in the access point's location prediction, which is 6cm in our scenario. Next, assuming an orientation error of $\Delta\beta$, we observe that the error will be $R\Delta\beta$ in the access point's location, where R is the estimate of the distance to the access point. Hence, the upper-bound on the total error propagated will be $\Delta r + R\Delta\beta$, which for an average indoor distance of $R = 5\text{m}$ would be 32cm.

Second, for the relative antenna location estimation, from Figure 8b we can see that the error in bot's location, Δr , translates to error in the angle estimated at the access point, $\beta_i + \Delta\beta_i$, where approximately $\Delta\beta_i = \frac{\Delta r}{R}$. Hence, we redefine A from Equation 9 as $A' = A \begin{bmatrix} 1 & \frac{\Delta r}{R} \\ -\frac{\Delta r}{R} & 1 \end{bmatrix}$, while b remains unchanged. Thus we can re-write Equation 11, assuming $\mathbf{x}_1 = 0$, as

$$\min_{\mathbf{x}'} \|\mathbf{A}'\mathbf{x}' - \mathbf{b}\|^2 \quad (14)$$

where $\mathbf{x}' = \mathbf{x} + \Delta\mathbf{x}$, and $\Delta\mathbf{x} = [\Delta x \ \Delta y]^T$. Solving for $\Delta\mathbf{x}$ from the Equations 11 and 14, and simplifying by neglecting higher order error polynomial terms we can see that $\Delta x = \frac{\Delta r}{R}y$, $\Delta y = \frac{\Delta r}{R}x$. We know that $\mathbf{x} = [x \ y]^T$ is of the order of few centimeters, while Δr is of the order of few centimeters and R of the order of few meters, which reduces the whole expression for Δx and Δy to be of order of $\frac{1}{10}$ millimeter, which is well within limits of the tolerance for relative antenna localization. Thus we observe that the relative antenna geometry on the access points can be estimated accurately to within few millimeters using LocAP and its implementation on our autonomous system.

5 Evaluation

Now that we have seen all the components of LocAP, we evaluate LocAP's performance in a real world deployment to see if it has conformed to the stringent requirements we established in Section 2. For this we have deployed our autonomous bot in two different indoor environments, as shown in [8], that span 1000 sq. ft. in area, and have 8 different access points deployed at different locations, heights and orientation. Across these 8 different access points, we have covered two standard antenna geometries, linear and square antenna arrays, and covered 5 different antenna separations, $\{\lambda/2, \lambda, 3\lambda/2, 2\lambda, 5\lambda/2\}$, where $\lambda = 6\text{ cm}$ is the minimum wavelength in the 155 channel of the 5GHz frequency band. Throughout this experiment, we collect CSI from multiple access points across space and time which is used to implement LocAP. The ground truth for all the evaluations are measured accurately with a commodity laser range finder [11], that is accurate up to 1mm, after carefully marking the axes on the ground and labeling the 1000 sq ft space of experimentation. This entire process of labeling the experimental space of 1000 sq ft takes a minimum of one hour spent by a group of at least three people. While there is two decades of CSI based WiFi localization, LocAP is the first work to tackle the problem of *reverse localization* of the WiFi access points and thus is compared with a state-of-the-art AoA based user localization algorithm [37], SpotFi, which combines data across multiple anchor locations.

With the given setup the overview of LocAP's results are as follows: LocAP achieves 5 cm of median localization error for the first antenna localization utilizing the weighted least squares formulation while a simple least-squares problem achieves just 8 cm of median localization error. Further, the relative geometry prediction algorithm of LocAP locates the access points in this setup accurately with a median antenna separation error of 3 mm and a median orientation error of 3° , whereas the state-of-the-art localization algorithms achieve a 150 mm median error for antenna separation and 25° median deployment orientation error as shown in Figure 9.

A final case study of user localization with the updated LocAP's AP attributes showed a reduction of 28 cm in me-




Figure 9: **Single Antenna Localization accuracy:** Shows the localization error of locating a single antenna on each AP (a) for various bandwidths and (b) for various number of antennas on the client on the autonomous bot. (c) **Antenna Separation:** CDF plot of error in measuring antenna separation across 8 different access point deployments. (d) **Deployment Orientation:** CDF plot of error in measuring deployment orientation across 8 different Access point deployments. The black vertical lines in the plots represent the requirements established in Section 2

dian user localization compared to the manual AP attribute mapping.

5.1 AP Location accuracy

To evaluate the access point localization accuracy, we deploy it in 8 different test scenarios across various heights of access points, different locations, environments and distances from the bot. To get a statistically accurate estimate of these locations, we have collected the CSI corresponding to each of these manually determined locations at 20 different time instants. With this data, we have estimated the location of each individual antenna on these access points using a least-squares triangulation algorithms employing [37]. As shown in Figure 9a, we find that the median error is 5 cm, well below the established threshold. Unfortunately, manually measuring locations takes hours of manual time and thus defeats the purpose of LocAP.

Hence, we deploy LocAP on our autonomous platform [8] that collects the same amount of data within 5 minutes. We use the SpotFi algorithm [37] as a comparative baseline model for the bot data. SpotFi assumes accurate ground truth locations of the anchors unlike LocAP’s implementation that smartly rejects anchor locations that are unreliable. We observed that while the baseline model provides a median AP localization error of 20.5 cm, our weighted least squares with smart-rejection achieves 13.5 cm showing an improvement of 36% in AP localization.

Further, the bandwidth assumed for these initial results is 80MHz, while the commodity WiFi access points hardly operate at these bandwidths. These WiFi access points usually use either 20MHz or 40MHz bandwidths. To mimic this, we also collect CSI data with the same setup for both 40MHz and 20MHz bandwidths. These CSI estimates have then been utilized to test our algorithm at different WiFi bandwidths. The CDF plot for variation of localization accuracy across different bandwidths can be seen in Figure 9a. It is seen that at higher bandwidths, the localization accuracy is marginally better, while LocAP still attains centimeter-level accuracy for

localizing the access point.

The design of LocAP relies on the angles estimated from the CSI data received. While the above-reported results are for a 4-antenna station, a commodity off-the-shelf WiFi device does not always have 4 antennas. Hence, we performed another experiment to observe the effect of change in the number of antennas on LocAP. This was done by changing the number of antennas present on the station mounted on the mobile robot. The CDF plot for the localization error with the increasing number of antennas can be seen in Figure 9a. The localization accuracy increases with the increasing number of antennas on the client mounted on the mobile robot. This is evidenced by the lower median error observed with 3 antennas present on the mobile robot as seen in Figure 9b. We further observe that a 2 antenna WiFi device significantly hurts the performance of LocAP. This performance degradation is because for a 2 antenna system, the multipath need to be at least 90° apart for the two different paths to be resolved.

5.2 Relative Antenna Geometry Accuracy

After the location of the first antenna of the AP is obtained, LocAP finds the positions of the other antennas of the AP relative to the first antenna. This is achieved by traversing around the reverse localized antenna of the AP, as described in Section 3.2. To test this algorithm, we deploy APs with a linear antenna array and a square antenna array AP in the two aforementioned environments. Similar to AP location estimation, we have collected data for each antenna setup at 40 different time instances to obtain statistically accurate results. The relative antenna locations on these APs were measured using LocAP and then compared with the ground truth to get the relative antenna localization errors and the deployment orientations. We further compare these results with that derived by state-of-the-art localization algorithm, SpotFi [37].

Relative Antenna Separation: We first measure the relative antenna separation of all the antennas on the access point with respect to the first antenna and the CDF plot for the




Figure 10: **User Localization accuracy:** Shows the CDF of localization accuracy after localizing the access points with LocAP and compared with those results of the manually labeled APs.

errors in relative antenna localization is shown in Figure 9c. We can see that the median error is about 3 mm for the relative antenna localization of LocAP while the state-of-the-art WiFi localization algorithm combined over multiple bot locations and time instances achieves 20 cm of median antenna separation error. Thus we show that LocAP achieves millimeter-level accuracy and meets the 5 mm error threshold set in Section 2 for predicting the antenna separation of the access point.

Deployment Orientation: We also measure the deployment orientation of all the antennas on the access point with respect to the first antenna and the CDF plot for the errors in the deployment orientation is shown in Figure 9d. We can see that while the state-of-the-art localization algorithm has a median error of 25° , LocAP’s deployment orientation prediction algorithm achieves a median orientation error of just 3° , meeting the 7° limit set in Section 2.

5.3 Case Study: User Localization

So far we have seen the performance of LocAP in accurately predicting the access point attributes. We implement LocAP, to enable CSI based indoor user localization. Further LocAP is automated by deploying on a bot to remove any manual labor and time or human errors. As discussed in Section 1, human based measurements lead to high degree of errors, especially in the antenna separation measurements that are needed to be accurate to less than 5mm of errors, especially when the antennas are housed in a casing whose datasheets provided by the chip designers do not contain information regarding the antenna placements on board [28, 47, 48]. Further, the antenna placement is determined mostly by the manufacturer, and the vast cardinality of the available vendors and their models make it impossible to estimate the antenna geometry from their datasheets, which also mostly do not discuss about the antenna placements on board [6, 7, 17, 57]. Additionally, deployment orientation has to be measured accurate to less

than 7° of error, which becomes extremely impossible for manual measurements. While we have shown 3mm ($<5\text{mm}$) error in predicting antenna separation and 3° ($<7^\circ$) error in orientation deployment predictions for LocAP. To verify the effect of both LocAP mapped AP attributes and manually mapped AP attributes on the state-of-the-art indoor WiFi localization algorithms [37, 53, 67], we have asked a group of 25 people to measure the first antenna locations, relative antenna separations and the deployment orientation of the access points deployed in a realistic scenario. Users have been provided with a laser range finder [11] and compass based apps used in smartphones.

From these user measurements, we have observed that manual mapping can make their best efforts to map the AP locations accurate with 21 cm median error, the antenna separation accurate to within 4 mm of median error, and a median absolute error of about 13° in measuring the deployment orientation of the access point.

We then deploy LocAP in a 1000 sq ft environment and locate the 4 access points’ attributes. A moving user that covers 300 different marked locations in this environment is then localized using both the manually mapped and LocAP’s mapped AP attributes and the corresponding CDF is shown in Figure 10c. From this plot, we can see that while human mapped AP attributes have a median localization error of 78 cm, LocAP’s AP locations achieve 50 cm median error. Thus we can see that LocAP solves for the fundamental dependency of CSI based user localization algorithms by accurately predicting the AP attributes within the physical map.

6 Related Work

There has been significant work in the field of localization and LocAP’s implementation work on reverse localizing the access points is closely related to the work in the following three fields:

Indoor Localization: Wide-scale deployment of WiFi based infrastructure and WiFi chips on hand-held devices makes indoor localization promising for various indoor navigation applications. There has been extensive research in WiFi based indoor localization algorithms over the past two decades [10, 15, 22, 23, 30, 37, 39, 41, 43, 45, 51, 53, 58–60, 65–69, 72, 74]. While most of the initial work was based on the Received Signal Strength Information [10, 15, 45, 60, 74] these algorithms do not achieve meter-level localization, or require extensive fingerprinting to achieve desired decimeter-level localization. Thus, most of the later work has been focused on CSI based localization algorithms [22, 30, 37, 39, 51, 53, 58, 59, 65–69]. LocAP which leverages the idea of Angle of Arrival based localization. Some such algorithms which have been developed in the past few years [37, 67] achieve decimeter-level localization and extend it to achieve centimeter-level localization accuracy. However, these WiFi based localization algorithms assume the knowledge of the location of the AP to measure

the user's location with respect to the AP location. In contrast to the above work, LocAP builds a relative localization technique which provides millimeter-level accuracy for the antenna geometry on the AP. Furthermore, we also demonstrate that LocAP can solve for the antenna separation values larger than a single wavelength (λ).

Source Localization: Solving the problem of accurate knowledge of the WiFi AP locations have been attempted for RSSI based [26] and CSI based [54] systems. But these algorithms do not achieve centimeter-level localization for APs, but solve for the general regional mapping of these access points. These works are limited by the available bandwidth and thus there has also been significant work on ultra-wideband (UWB) based localization [5, 13, 14, 18, 34, 46, 49] and anchor localization algorithms [12, 19, 20, 34–36]. But these UWB systems require new infrastructure deployment. Similarly, there has been significant work towards a beacon based localization system [9, 29, 32, 33, 42, 52, 62, 63, 70, 71, 73] which have been shown to achieve decimeter-level localization but also need additional deployment of infrastructure. LocAP solves the problems of exact WiFi access point localization and exact antenna placements on WiFi Access Points.

Relative Localization: LocAP solves for millimeter-level accurate antenna placements on any given WiFi access point by borrowing and extending the principles from wireless tracking. Wireless tracking or relative localization is a well-solved problem unlike localization, with reported accuracies up to few centimeters and few millimeters [38, 61, 64]. Though all of these algorithms would need the separation between two consecutive locations to be tracked to be less than $\lambda/2$ distance apart, LocAP solves for relative localization of two antennas that are at any arbitrary distance from each other, including for distances greater than $\lambda/2$ apart. Thus LocAP can enable high mobility tracking for indoor WiFi devices.

SLAM Automation: There has been exhaustive research conducted in graph based SLAM algorithms [24]. In LocAP we employ a SLAM based autonomous bot to report ground truth and also design a metric to understand the confidence of the bot for a given ground truth. Confidences for reported measurements can be extracted from the marginal co-variances of the nodes used to describe these variables and are used to perform data association [27, 31, 44, 56]. Though these numerical methods are valid, most of them are not implemented on standard SLAM platforms, to the best of our knowledge. Furthermore, commonly used frameworks [25, 40] do not readily expose these marginal co-variances. We extend the methods described in [16] as a proxy for these internal co-variance metrics.

7 Conclusion and Future Work

We presented, LocAP, an automated reverse localization system of the existing WiFi APs that was successful in achieving the requirements for accurate localization of AP position,

antenna separation and deployment orientation. After the mobile robot is allowed to traverse the unknown environment, we have a map of the indoor environment and the reverse localized positions of all the APs in this environment. If we consider the map to be part of a coordinate system, we can provide each access point with its coordinate in the environment, such that the AP becomes self-aware about its location. When a new user enters this environment, and associates with one of these APs, they can locate the user in turn almost instantaneously relative to their position.

Using the mapping and reverse localization information, we can provide accurate indoor localization and navigation for large indoor environments. These accurate AP location attributes aids many of the networking issues like user location based smart hand-off, network load balancing utilizing both AP locations and client locations and other networking services based on AP and client locations. Further, with the emergence of 5G and 11ad/ax wireless protocols, where directional beams become more and more important, these angle of arrival estimates that are provided by LocAP, can be further used to perform smart-beamforming at both the client and the AP side.

In LocAP we have analyzed the 2D scenario when the access point is in the same plane as the user to be located. In a real world deployment the access point is placed at least a meter above the user height thus subtending a non-zero polar angle at the access point. This does not affect LocAP's algorithm on relative geometry prediction as the cartesian co-ordinates defined absorb the polar angular term. Thus unchanging the formulation of the relative antenna geometry prediction algorithm enabling LocAP to perform accurately under 3D deployments.

Acknowledgements– We thank anonymous reviewers and our shepherd, Ellen Zegura, for their insightful comments and feedback. We thank Deepak Vasishth and the members of WCSNG for their comments and feedback throughout the process of the paper.

References

- [1] Apple Maps. <https://www.apple.com/ios/maps/>.
- [2] Bing Maps. www.bing.com/maps.
- [3] Google Maps. www.maps.google.com.
- [4] Open Street Map. www.openstreetmap.org.
- [5] N. A. Alsindi, B. Alavi, and K. Pahlavan. Measurement and modeling of ultrawideband toa-based ranging in indoor multipath environments. *IEEE Transactions on Vehicular Technology*, 58(3):1046–1058, 2009.
- [6] Apple. Airport Support. <https://support.apple.com/airport>.

- [7] Aruba Networks. DS-AP303Series. https://www.arubanetworks.com/assets/ds/DS_AP303Series.pdf.
- [8] R. Ayyalasomayajula, A. Arun, C. Wu, S. Sanatan, S. Abhishek, D. Vasisht, and D. Bharadia. Deep learning based wireless localization for indoor navigation. In *The 26th Annual International Conference on Mobile Computing and Networking (MobiCom '20)*. ACM, 2019.
- [9] R. Ayyalasomayajula, D. Vasisht, and D. Bharadia. Bloc: Csi-based accurate localization for ble tags. In *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*, pages 126–138. ACM, 2018.
- [10] V. Bahl and V. Padmanabhan. RADAR: An In-Building RF-based User Location and Tracking System. INFOCOM, 2000.
- [11] BOSCH. Laser Measure DLE40 professional. <http://www.bosch-professional.com/ma/en/laser-measure-dle-40-131500-0601016300.html>.
- [12] M. Cao, B. D. Anderson, and A. S. Morse. Sensor network localization with imprecise distances. *Systems & control letters*, 55(11):887–893, 2006.
- [13] Y.-T. Chan, W.-Y. Tsui, H.-C. So, and P.-c. Ching. Time-of-arrival based localization under nlos conditions. *IEEE Transactions on Vehicular Technology*, 55(1):17–24, 2006.
- [14] H. Chen, G. Wang, Z. Wang, H.-C. So, and H. V. Poor. Non-line-of-sight node localization based on semi-definite programming in wireless sensor networks. *IEEE Transactions on Wireless Communications*, 11(1):108–116, 2012.
- [15] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan. Indoor Localization Without the Pain. *MobiCom*, 2010.
- [16] S. Choi, Q.-Y. Zhou, and V. Koltun. Robust reconstruction of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5556–5565, 2015.
- [17] Cisco. catalyst-9120ax. <https://www.cisco.com/c/en/us/products/collateral/wireless/catalyst-9120ax-series-access-points/datasheet-c78-742115.html#Aestheticallyredesignedfornextgenerationenterprise>.
- [18] L. Cong and W. Zhuang. Nonline-of-sight error mitigation in mobile location. *IEEE Transactions on Wireless Communications*, 4(2):560–573, 2005.
- [19] C. Di Franco, A. Prorok, N. Atanasov, B. Kempke, P. Dutta, V. Kumar, and G. J. Pappas. Calibration-free network localization using non-line-of-sight ultra-wideband measurements. In *Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks*, pages 235–246. ACM, 2017.
- [20] Y. Diao, Z. Lin, and M. Fu. A barycentric coordinate based distributed localization algorithm for sensor networks. *IEEE Transactions on Signal Processing*, 62(18):4760–4771, 2014.
- [21] H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.
- [22] J. Gjengset, J. Xiong, G. McPhillips, and K. Jamieson. Phaser: Enabling Phased Array Signal Processing on Commodity Wi-Fi Access Points. *MobiCom*, 2014.
- [23] A. Goswami, L. E. Ortiz, and S. R. Das. Wigem: A learning-based approach for indoor localization. In *Proceedings of the Seventh Conference on emerging Networking EXperiments and Technologies*, page 3. ACM, 2011.
- [24] G. Grisetti, R. Kummerle, C. Stachniss, and W. Burgard. A tutorial on graph-based slam. *IEEE Intelligent Transportation Systems Magazine*, 2(4):31–43, 2010.
- [25] G. Grisetti, C. Stachniss, W. Burgard, et al. Improved techniques for grid mapping with rao-blackwellized particle filters. *IEEE transactions on Robotics*, 23(1):34, 2007.
- [26] D. Han, D. G. Andersen, M. Kaminsky, K. Papagiannaki, and S. Seshan. Access point localization using local signal strength gradient. In *International Conference on Passive and active network measurement*, pages 99–108. Springer, 2009.
- [27] V. Ila, L. Polok, M. Solony, and P. Svoboda. Highly efficient compact pose slam with slam++. *arXiv preprint arXiv:1608.03037*, 2016.
- [28] Intel. dual-band-wireless-ac-9260.
- [29] V. Iyer, V. Talla, B. Kellogg, S. Gollakota, and J. Smith. Inter-technology backscatter: Towards internet connectivity for implanted devices. In *SIGCOMM*, 2016.
- [30] K. Joshi, S. Hong, and S. Katti. PinPoint: Localizing Interfering Radios. NSDI, 2013.
- [31] M. Kaess and F. Dellaert. Covariance recovery from a square root information matrix for data association. *Robotics and autonomous systems*, 57(12):1198–1210, 2009.

- [32] B. Kellogg, A. Parks, S. Gollakota, J. R. Smith, and D. Wetherall. Wi-fi backscatter: Internet connectivity for rf-powered devices. In *ACM SIGCOMM Computer Communication Review*, 2014.
- [33] B. Kellogg, V. Talla, S. Gollakota, and J. R. Smith. Passive wi-fi: Bringing low power to wi-fi transmissions. In *NSDI*, 2016.
- [34] B. Kempke, P. Pannuto, and P. Dutta. Harmonium: Asymmetric, bandstitched uwb for fast, accurate, and robust indoor localization. In *2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 1–12. IEEE, 2016.
- [35] U. A. Khan, S. Kar, and J. M. Moura. Distributed sensor localization in random environments using minimal number of anchor nodes. *IEEE Transactions on Signal Processing*, 57(5):2000–2016, 2009.
- [36] U. A. Khan, S. Kar, and J. M. Moura. Diland: An algorithm for distributed sensor localization with noisy distance measurements. *IEEE Transactions on Signal Processing*, 58(3):1940–1947, 2010.
- [37] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti. SpotFi: Decimeter Level Localization Using Wi-Fi. SIGCOMM, 2015.
- [38] M. Kotaru and S. Katti. Position tracking for virtual reality using commodity wifi. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 68–78, 2017.
- [39] S. Kumar, S. Gil, D. Katabi, and D. Rus. Accurate Indoor Localization with Zero Start-up Cost. *MobiCom*, 2014.
- [40] M. Labbe and F. Michaud. Rtab-map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation. *Journal of Field Robotics*, 2019.
- [41] A. M. Ladd, K. E. Bekris, A. Rudys, L. E. Kavraki, and D. S. Wallach. Robotics-based location sensing using wireless ethernet. *Wireless Networks*, 11(1-2):189–204, 2005.
- [42] Y. Ma, N. Selby, and F. Adib. Drone relays for battery-free networks. In *SIGCOMM*.
- [43] A. T. Mariakakis, S. Sen, J. Lee, and K.-H. Kim. Sail: Single access point-based indoor localization. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*, pages 315–328. ACM, 2014.
- [44] J. Neira and J. D. Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on robotics and automation*, 17(6):890–897, 2001.
- [45] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 32–43. ACM, 2000.
- [46] A. Prorok and A. Martinoli. Accurate indoor localization with ultra-wideband using spatial models and collaboration. *The International Journal of Robotics Research*, 33(4):547–568, 2014.
- [47] Qualcomm. [ar6004. https://www.qualcomm.com/media/documents/files/ar6004-datasheet.pdf](https://www.qualcomm.com/media/documents/files/ar6004-datasheet.pdf).
- [48] Quantenna. [QSR10GU-AX. http://www.quantenna.com/wp-content/uploads/2018/04/QSR10GU-AX-V1.1.pdf](http://www.quantenna.com/wp-content/uploads/2018/04/QSR10GU-AX-V1.1.pdf).
- [49] Z. Sahinoglu. *Ultra-wideband positioning systems*. Cambridge university press, 2008.
- [50] S. Sen, J. Lee, K.-H. Kim, and P. Congdon. Avoiding multipath to revive inbuilding wifi localization. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*, pages 249–262. ACM, 2013.
- [51] S. Sen, B. Radunovic, R. R. Choudhury, and T. Minka. You are facing the mona lisa: Spot localization using phy layer information. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 183–196. ACM, 2012.
- [52] 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*, pages 31–42, 2016.
- [53] E. Soltanaghaei, A. Kalyanaraman, and K. Whitehouse. Multipath triangulation: Decimeter-level wifi localization and orientation with a single unaided receiver. In *MobiSys*, 2018.
- [54] A. P. Subramanian, P. Deshpande, J. Gao, and S. R. Das. Drive-by localization of roadside wifi networks. In *IEEE INFOCOM 2008-The 27th Conference on Computer Communications*, pages 718–725. IEEE, 2008.
- [55] T-Series. [VICON. www.vicon.com/products/documents/Tseries.pdf](http://www.vicon.com/products/documents/Tseries.pdf).

- [56] G. D. Tipaldi, G. Grisetti, and W. Burgard. Approximate covariance estimation in graphical approaches to slam. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3460–3465. IEEE, 2007.
- [57] UniFi. UAP-AC-HD-DS. https://dl.ubnt.com/datasheets/unifi/UniFi_UAP-AC-HD_DS.pdf.
- [58] M. C. Vanderveen, C. B. Papadias, and A. Paulraj. Joint angle and delay estimation (jade) for multipath signals arriving at an antenna array. *IEEE Communications Letters*, 1(1):12–14, 1997.
- [59] D. Vasisht, S. Kumar, and D. Katabi. Decimeter-Level Localization with a Single Wi-Fi Access Point. NSDI, 2016.
- [60] H. Wang, S. Sen, A. Elgohary, M. Farid, M. Youssef, and R. R. Choudhury. No need to war-drive: Unsupervised indoor localization. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 197–210, 2012.
- [61] J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus. RF-compass: Robot object manipulation using rfids. In *Proceedings of the 19th annual international conference on Mobile computing & networking*, pages 3–14. ACM, 2013.
- [62] J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus. RF-compass: Robot Object Manipulation Using RFIDs. MobiCom, 2013.
- [63] J. Wang, H. Jiang, J. Xiong, K. Jamieson, X. Chen, D. Fang, and B. Xie. LiFS: Low Human-effort, Device-free Localization with Fine-grained Subcarrier Information. MobiCom, 2016.
- [64] J. Wang, D. Vasisht, and D. Katabi. Rf-idraw: Virtual touch screen in the air using rf signals. *ACM SIGCOMM*, 2014.
- [65] Y. Xie, J. Xiong, M. Li, and K. Jamieson. xd-track: leveraging multi-dimensional information for passive wi-fi tracking. In *HotWireless*, pages 39–43. ACM, 2016.
- [66] Y. Xie, J. Xiong, M. Li, and K. Jamieson. md-track: Leveraging multi-dimensionality in passive indoor wi-fi tracking. *arXiv preprint arXiv:1812.03103*, 2018.
- [67] J. Xiong and K. Jamieson. ArrayTrack: A Fine-grained Indoor Location System. NSDI, 2013.
- [68] J. Xiong, K. Jamieson, and K. Sundaresan. Synchronicity: Pushing the envelope of fine-grained localization with distributed mimo. In *HotWireless*, 2014.
- [69] J. Xiong, K. Sundaresan, and K. Jamieson. ToneTrack: Leveraging Frequency-Agile Radios for Time-Based Indoor Wireless Localization. MobiCom, 2015.
- [70] C. Xu, B. Firner, Y. Zhang, R. Howard, J. Li, and X. Lin. Improving RF-based Device-free Passive Localization in Cluttered Indoor Environments Through Probabilistic Classification Methods. IPSN, 2012.
- [71] 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. MobiCom, 2014.
- [72] M. Youssef and A. Agrawala. The Horus WLAN Location Determination System. MobiSys, 2005.
- [73] P. Zhang, D. Bharadia, K. Joshi, and S. Katti. Hitchhike: Practical backscatter using commodity wifi. In *SenSys*, 2016.
- [74] X. Zhu and Y. Feng. Rssi-based algorithm for indoor localization. *Communications and Network*, 5(02):37, 2013.