Privacy Preserving WiFi Sensing Using Differential Privacy

Holden Wozniak
University at Buffalo
Buffalo, NY, USA
holdenwo@buffalo.edu

Roshan Ayyalasomayajula
University at Buffalo
Buffalo, NY, USA
roshana@buffalo.edu

Abstract

Channel state information (CSI)-based WiFi sensing enables fine-grained indoor localization by exploiting phase-coherent structure across antennas and subcarriers. However, this same structure makes CSI inherently vulnerable to passive inference attacks capable of recovering precise user location from routine wireless measurements.

This paper presents a differentially private WiFi sensing framework that perturbs CSI directly in the physical phase domain prior to beamforming rather than applying post-processing or output-level noise. The mechanism injects calibrated Gaussian perturbations into frequency-dependent (distance) and antenna-dependent (angle) phase components, inducing controlled distortion in the resulting distance–angle multipath representation while preserving a physically consistent channel structure.

We derive a closed-form bound linking differential privacy parameters to expected Cartesian localization error, showing that localization uncertainty scales as $O(\epsilon^{-2})$ and increases with both angular and distance sensitivity. Across real-world CSI measurements from the WILD dataset, empirical results closely match the theoretical prediction and validate the proposed error model over multiple privacy budgets.

Experiments further show that the induced angular perturbations survive common CSI post-processing and refinement techniques, preventing reliable recovery of precise angle-of-arrival estimates even after adversarial filtering. Finally, we introduce a reverse reconstruction procedure that preserves forward consistency of perturbed multipath profiles while maintaining plausible CSI structure.

Overall, this work demonstrates that differential privacy can be embedded directly into the wireless sensing pipeline through phase-level channel perturbations, enabling quantifiable and controllable tradeoffs between localization accuracy and privacy leakage.

1 Introduction

Indoor localization and wireless sensing technologies have become increasingly common as these technologies provide location awareness utility to systems without having to come up with entirely separate infrastructure [14, 19, 34]. While GPS technologies perform well outdoors, signal attenuation and multipath propagation significantly limit their effectiveness in indoor environments [31]. Because of this, other methods are used for indoor localization. In the case of WiFi sensing, channel state information (CSI) that already exists on routers can be leveraged to compute user locations relative to the routers position which provides better performance and scalability [16, 20].

Bluetooth beaconing is a widely used indoor localization approach that relies on low-energy Bluetooth (BLE) transmitters placed in known fixed positions to periodically broadcast identifiers. A receiver estimates its location by measuring the received signal strength (RSSI) from multiple beacons and applying either fingerprinting or trilateration-based techniques [13]. While Bluetooth beacons have the benefits of low power consumption, low cost, and ease of deployment, they suffer from several limitations. RSSI measurements are highly sensitive to multipath fading, body blockage, device orientation, and environmental dynamics which leads to significant instability and reduced localization accuracy[13]. Additionally, beacon-based systems require dense infrastructure deployment to achieve good accuracy which increases maintenance overhead and limits scalability in large or complex environments [22, 34].

Camera-based tracking is another method that can be used for localization that also comes with its own concerns. Camera-based tracking often requires line-of-sight and often has to fall back to secondary sensing and predictive algorithms to maintain tracking. Camera-based tracking is also sensitive to lighting conditions, as the same scene under different lighting conditions can appear different to the camera and break tracking[35]. However, the biggest concern with camera-based tracking is privacy as it requires constant scanning of environments. Coupled with modern advancements in machine learning, profiles on non-consenting users become easier to create and put those users at risk. While effective in security applications, camera-based tracking for localization comes with surveillance concerns when applied to public spaces [35].

In WiFi sensing, CSI-based localization has emerged as a powerful technique due to its ability to capture fine-grained wireless channel characteristics [20, 26]. It provides detailed

information about the amplitude and phase response of subcarriers across multiple antennas; this allows systems to exploit multipath propagation patterns and derive highly accurate spatial estimation [15, 16]. Through sampling CSI, angle vs. distance multipath profiles can be constructed which ultimately map into cartesian coordinate estimates[17, 24, 30]. Compared to RSSI-based approaches, CSI-based methods offer substantially higher space resolution and robustness in densely populated environments [17, 26].

Despite the advantages, the strengths of CSI-based methods enabling stronger localization accuracy also come with significant privacy concerns. CSI measurements contain distinctive channel characteristics that can be used to track specific user locations over time [28, 36]. Furthermore, because wireless signals are capable of propagating through walls in indoor spaces, users may unknowingly expose sensitive location information to nearby access points without their consent [2, 3]. As WiFi sensing systems become more widely adopted, the issue of privacy has become a crucial challenge that needs addressing [36].

Traditional privacy-preserving approaches often rely on access control, encryption mechanisms, and anonymization. However, these methods do not directly address the leakage of information that is contained in the channel measurements themselves. Even with labels removed tying users to specific channel information, CSI still enables behavior inference as more packets are sampled over time [4, 25, 27]. Therefore, the focus of research must shift to creating a privacy preserving mechanism that operates directly on sensing data while still maintaining sufficient utility for applications.

Differential privacy is a widely studied framework for providing mathematically quantifiable privacy guarantees, especially in statistical and machine learning systems [10, 11]. Rather than trying to completely hide data (thus resulting in destroyed sensing utility), differential privacy aims to limit the influence of any single user's data on the output of a system by introducing carefully calibrated noise. This is achieved typically though the addition of Gaussian noise scaled according to a tunable privacy budget and sensitivity parameters[1, 9, 21]. While differential privacy has been studied extensively in database and machine learning systems, applying it to WiFi sensing has remained an open problem due to the complex relationship between perturbations in physical data and downstream estimation accuracy [12].

As such, this work presents a WiFi sensing framework that applies differentially private perturbations directly to phase information through the manipulation of steering vectors. Gaussian noise calibrated according to privacy parameters is injected into the phase components of both the angle and distance steering vectors. In doing so, localization peaks in the distance-angle multipath profile are broadened and shifted; this introduces uncertainty into the recovered estimates while still preserving the observed structure of the wireless channel. Furthermore, this work also provides a reverse reconstruction mechanism capable of generating plausible CSI measurements from the spoofed multipath profiles to further hide the perturbation process.

Through the combination of differential privacy theory, wireless channel modeling, and experimental validation, this work provides a mathematically based framework for privacy-preserving WiFi localization systems that enable control for trading off sensing utility and user privacy protection. This work proves the error introduced to the sensing pipeline is bounded to $O(\epsilon^{-2})$ where ϵ is a tunable privacy budget; it verifies it empirically through repeated averaging of samples with noise applied to them and comparing with known true locations. Additionally, this work shows the angle-of-arrival error introduced survives common refinement techniques as well as verifies a spoofing mechanism to further hide the privacy providing mechanism at work.

2 Background

2.1 WiFi Localization

WiFi localization systems commonly utilize CSI which is fundamental to the beamforming process of sending signals in the direction of connected users [15, 16]. CSI provides detailed measurements of the wireless channel response from multiple subcarriers across antennas. Unlike RSSI methods that compress channel behavior into a single power measurement, CSI preserves both the amplitude and phase information for each subcarrier which enables significantly higher spatial resolution as well as deployment in more variable environments [20, 26]. CSI for a given subcarrier k can be represented as

$$H_k = A_k e^{j\phi_k} \tag{1}$$

where A_k represents the amplitude response and ϕ_k represents the phase response of the wireless channel at subcarrier k. Because routers divide wireless transmissions across many narrowband subcarriers, CSI samples the channel frequency response at multiple frequencies simultaneously and provides more complete information about the characteristics of signal propagation [15].

The phase component of CSI is influenced by both the propagation distance traveled by a signal as well as the angle at which the signal reaches the antenna array [17, 30]. As a wireless signal propagates through a space, its phase is proportional to the total path length traveled; additional phase offsets are introduced across antennas due to additional travel time from the first antenna to other antennas due to the arrival angle. Therefore, the measured phase is the sum of these two components

$$\phi_k = \phi_d + \phi_\theta \tag{2}$$

CSI for a given subcarrier can therefore be rewritten as the product of two exponential terms to enable independent exploitation of range and direction information when constructing angle-distance multipath profiles:

$$H_k = A_k e^{j\phi_d} e^{j\phi_\theta} \tag{3}$$

In CSI-based localization, steering vectors are used to model how a signal arriving to an antenna array at different times can be translated into the angle it arrives at and is represented by [17, 24]

$$a_\theta(\theta) = \begin{bmatrix} 1 \\ e^{j\frac{2\pi}{c} f_c d_{ant} \cdot 1 \cdot \sin(\theta)} \\ \vdots \\ e^{j\frac{2\pi}{c} f_c d_{ant} \cdot (N-1) \cdot \sin(\theta)} \end{bmatrix}. \tag{4}$$

where d_{ant} is the antenna spacing, N is the number of antennas, f_c is the carrier frequency of the signal, and c is the speed of light. This vector encodes how a signal from direction θ would appear across the antenna array. The system also involves a distance steering vector that captures phase due to propagation distance:

$$a_d(d) = \begin{bmatrix} e^{j\frac{2\pi}{c} f_1 d} \\ e^{j\frac{2\pi}{c} f_2 d} \\ \vdots \\ e^{j\frac{2\pi}{c} f_N d} \end{bmatrix}. \tag{5}$$

where f_N is the subcarrier frequency of subcarrier N and d is the total distance traveled in meters. Because each subcarrier experiences a slightly different wavelength, the received phase rotates at a frequency-dependent rate [17].

When applied to raw CSI, these steering vectors map the CSI to distance-angle pairs

$$MP(\theta, d) = a_\theta^H(\theta)\, H\, a_d(d), \tag{6}$$

and large values of $MP(\theta, d)$ indicate the strongest path from a user to a router [17, 30].

2.2 Threat Model

This work considers a passive adversarial sensing model in which an attacker attempts to infer sensitive user location information using CSI-based WiFi localization techniques [28, 36]. The attacker is assumed to possess access to CSI measurements collected at a multi-antenna WiFi access point and is capable of constructing distance-angle multipath profiles using standard beamforming and steering vector based localization methods. Using these measurements, the attacker aims to recover precise user position and track movement over time to create tracking profiles for specific users.

The attacker is assumed to have knowledge of the wireless sensing pipeline including antenna geometry, steering-vector construction, and localization algorithms used to generate multipath profiles. Furthermore, the attacker may collect CSI measurements continuously over time in order to average noise, refine localization estimates, or apply post-processing techniques intended to improve sensing accuracy. The attacker, however, does not possess direct knowledge of the privacy perturbations applied within the proposed system; they are generally aware that there is a privacy protecting mechanism in place but have no knowledge of how noise injection occurs.

The primary privacy objective of this work is to reduce the attacker's confidence in recovering precise spatial information from CSI measurements while preserving sufficient wireless sensing utility for legitimate applications. Rather than completely destroying localization capability, the goal is to introduce controlled uncertainty into the localization pipeline such that inferred user positions become probabilistically ambiguous within bounded spatial error regions.

Unlike client-side privacy mechanisms that require end-user devices to actively manipulate transmitted signals or beamforming behavior [6], the approach considered in this work operates at the router level. Therefore, users are protected without requiring modifications to client hardware, firmware, or communication protocols.

Under this threat model, the effectiveness of the proposed mechanism is evaluated according to the following criteria: (1) successful reduction in localization accuracy achieved by the attacker and bounding the error in terms of the privacy budget, (2) persistence of localization uncertainty after post-processing or refinement operations, and (3) preservation of physically plausible CSI structure and wireless communication utility despite the introduction of privacy-preserving perturbations.

2.3 Phase Perturbation

Recent work enabling privacy in WiFi sensing have shown how perturbations applied to physical-layer channel measurements can allow for users to opt-out of sensing systems while still preserving typical WiFi functionality for data transfer. Specifically, MIRAGE[6] is a system that introduces obfuscating user location by adding small amounts of delay in the direct path channel such that the length of the direct path appears longer than the multipath due to the phase rotation the signal arrives at. In doing so, localization algorithms take the multipath (now appearing as the shortest path) as the true user location and creates a false location. Because the direct path is just delayed but still intact, communication throughput is still mostly sustained.

The underlying rationale behind phase perturbation is that CSI-based localization is heavily dependent on coherent spatial phase alignment to construct accurate multipath profiles [17, 24]. As such, small perturbations in CSI phase propagate through beamforming and angle estimation pipelines which causes heatmap peaks to shift, widen, or even disappear completely. Localization algorithms estimate user position by identifying coherent maxima in the multipath profile; disrupting phase consistency directly reduces the certainty sensing systems can recover location information.

It is important to note, however, that protections such as MIRAGE are client-side [6]; as such, it puts the responsibility of protecting privacy onto users. While effective at protecting users through physical modification of the signal, it is not reasonable to assume users wish to automatically opt-in to sensing and then give them the option to opt-out. Therefore, users should not be expected to take additional measures to protect themselves from publicly deployed systems, thus making the burden of providing privacy router-side.

2.4 Differential Privacy

The differential privacy (DP) framework is designed to provide quantifiable privacy guarantees while preserving the overall utility of a system [10, 11]. Formally, a randomized mechanism M satisfies (ϵ, δ)-DP if, for any pair of neighboring datasets D and D' differing by at most one observations, and for any possible output set S:

$$\Pr[M(D) \in S] \leq e^{\varepsilon} \Pr[M(D') \in S] + \delta. \tag{7}$$

where ϵ represents the privacy budget and controls the strength of the privacy guarantee and δ represents a tiny probability that the privacy fails. Smaller values of ϵ correspond to stronger privacy guarantees as mechanism outputs become increasingly indistinguishable between neighboring datasets.

DP enabled systems are able to achieve this by introducing randomized noise carefully calibrated according to environmental sensitivities [10]. This sensitivity measures the maximum change in the output of a function caused by omitting any single observation in the input dataset. For a function $f(D)$, the L_2-sensitivity is defined as

$$\Delta_2 f = \max_{D \sim D'} \|f(D) - f(D')\|_2. \tag{8}$$

The Gaussian mechanism has nice properties such as appearing as natural noise and easier translation of DP guarantees and is adopted by many implementations of DP-enabled systems [1, 9, 21]. Noise samples are drawn from a zero-mean Gaussian distribution

$$\eta \sim \mathcal{N}(0, \sigma^2), \tag{9}$$

where the standard deviation σ is selected by

$$\sigma \geq \frac{\sqrt{2 \ln(1.25/\delta)}\, \Delta_2 f}{\varepsilon}. \tag{10}$$

The relationship of σ and ϵ establishes the privacy-utility tradeoff inherent to DP.

3 Localization Pipeline

The localization pipeline considered in this work makes location estimates by transforming raw CSI measurements into angle-distance multipath profiles [17, 30]. The system assumes that multi-antenna access points capable of measuring CSI across multiple subcarriers are used for sensing [15]. For each received packet, the wireless channel response is sampled across antennas and subcarriers, producing a complex-valued CSI measurement of the form

$$H_k = A_k e^{j\phi_d} e^{j\phi_\theta} \tag{11}$$

where A_k represents the amplitude response and ϕ_k represents the phase response associated with subcarrier k. To estimate spatial location, the measured CSI is correlated against steering vectors corresponding to hypothesized propagation angles and distances [17, 24]. Previously, using steering vectors we show that raw CSI can be translated to a distance-angle heatmap:

$$MP(\theta, d) = a_\theta^H(\theta)\, H\, a_d(d), \tag{12}$$

where H represents the measured CSI vector and $(\cdot)^H$ denotes the Hermitian transpose. Peaks in the multipath profile correspond to the most likely propagation paths of the wireless channel [17, 30]. By identifying dominant peaks within the distance-angle heatmap, the system estimates the relative spatial location of targets within the environment.

This translation is closely related to foundational beamforming and subspace-based localization methods commonly used in wireless sensing systems [18, 23]. By correlating measured CSI against hypothesized steering vectors, the localization pipeline effectively performs filtered matching across candidate distance-angle pairs. Peaks in the resulting multipath profile correspond to constructive phase alignment between the measured channel response and the hypothesized propagation geometry, thus enabling estimation of dominant propagation paths and user locations.

The final stage of the localization pipeline converts the detected distance-angle coordinates into Cartesian coordinate estimates. For a detected multipath peak (d, θ), the estimated location is computed as

$$x = d \cos\theta, \qquad y = d \sin\theta \tag{13}$$

allowing the system to map wireless propagation measurements into two-dimensional Cartesian coordinates. Because the localization process depends heavily on coherent phase

relationships across antennas and subcarriers, perturbations introduced into CSI phase measurements propagate directly into multipath peak estimation and downstream spatial localization accuracy [17, 24].

4 Proposed Privacy Mechanism

The proposed privacy-preserving mechanism operates by introducing differentially private perturbations directly into the phase structure of CSI measurements prior to multipath profile generation. Rather than perturbing the final localization estimate itself, the mechanism modifies the intermediate steering-vector phase relationships used for angle and distance estimation [17, 24].

The proposed mechanism perturbs phase relationships by introducing randomized phase slopes across both subcarriers and antennas. The perturbation process is calibrated using the Gaussian differential privacy mechanism such that the injected noise magnitude is controlled by a privacy budget (ϵ, δ) [1, 10].

The Gaussian noise scale is determined according to

$$\sigma = \frac{\sqrt{2\ln(1.25/\delta)}\Delta}{\epsilon} \tag{14}$$

where Δ represents the sensitivity of the localization system. Because the localization process depends jointly on propagation distance and angle estimation, separate sensitivity parameters are defined for distance and angle perturbation:

$$\sigma_d = \Delta_d \frac{\sqrt{2\ln(1.25/\delta)}}{\epsilon} \tag{15}$$

$$\sigma_\theta = \Delta_\theta \frac{\sqrt{2\ln(1.25/\delta)}}{\epsilon} \tag{16}$$

where Δ_d represents the distance sensitivity and Δ_θ represents the angular sensitivity of the system. Gaussian perturbations are then sampled independently as

$$\delta_\tau \sim \mathcal{N}(0, \sigma_\tau^2) \tag{17}$$

$$\delta_\theta \sim \mathcal{N}(0, \sigma_\theta^2) \tag{18}$$

where $\sigma_\tau = \sigma_d / c$ converts the distance perturbation into an equivalent time-of-flight phase slope.

These perturbations are applied directly to the CSI phase structure through two complementary mechanisms. A frequency-dependent phase slope is introduced across subcarriers to perturb the estimated propagation distance:

$$\phi_{\text{freq}}(f) = 2\pi f \delta_\tau \tag{19}$$

An antenna-dependent phase slope is also introduced across antenna elements to perturb the angle-of-arrival (AoA) estimation:

$$\phi_{\text{ant}}(m) = 2\pi \frac{d_{ant}}{\lambda} m \delta_\theta \tag{20}$$

where m denotes the antenna index, d_{ant} is antenna spacing, and λ is the carrier wavelength. The combined perturbation is then applied multiplicatively to the CSI matrix:

$$\tilde{H} = H \odot e^{j(\phi_{\text{freq}} + \phi_{\text{ant}})} \tag{21}$$

where $\odot$ denotes element-wise multiplication. This operation preserves the overall complex structure of the wireless channel while introducing controlled uncertainty to the phase information of CSI.

The perturbed CSI measurements are then processed using the same steering-vector beamforming pipeline described previously [17, 23]. Because the perturbation alters phase coherence across both subcarriers and antennas, dominant localization peaks within the multipath heatmap become wider and shifted thus reducing localization certainty while preserving the overall physical structure of the CSI translation.

To further obscure the perturbation process and maintain physically plausible channel structure, this work additionally introduces an approximate reverse reconstruction mechanism capable of generating CSI-like measurements from a perturbed multipath profile. Let $\overline{MP}(d, \theta)$ denote a spoofed angle–distance heatmap. Distance and angle steering matrices are first constructed as

$$A_d(d) = e^{j\frac{2\pi f}{c}d} \tag{22}$$

$$A_\theta(\theta) = e^{j\frac{2\pi d}{\lambda}\sin\theta} \tag{23}$$

The CSI estimate is then reconstructed through pseudoinverse projection:

$$\hat{H} = A_\theta^\dagger \overline{MP} A_d^\dagger \tag{24}$$

where $(\cdot)^\dagger$ denotes the Moore–Penrose pseudoinverse. This reconstruction produces a CSI estimate whose forward transformation through a non-noisy mechanism approximately reproduces the perturbed multipath profile while maintaining realistic channel characteristics. Therefore, the resulting spoofed CSI appears physically consistent with the modified localization output, making direct detection of the perturbation process significantly more difficult.

It is important to note that the proposed mechanism does not physically modify the underlying wireless channel or transmitted communication signal itself. Instead, the perturbation is applied only to the CSI measurements reported to

the sensing and localization pipeline after channel estimation has already occurred. This differs fundamentally from systems such as MIRAGE [6], which introduce physical delay perturbations directly into the wireless propagation process to manipulate localization outcomes. In the proposed framework, the perturbed CSI exists purely as a sensing-layer representation used for localization processing and is not retransmitted over the air. Directly transmitting heavily perturbed CSI-equivalent phase distortions through the communication channel would significantly disrupt coherent demodulation, beamforming, channel equalization, and spatial stream recovery processes required for reliable wireless communication [15, 33].

5 Theoretical Analysis

To characterize the impact of differentially private phase perturbations on localization accuracy, this section derives an explicit bound relating Gaussian perturbation parameters to Cartesian localization error. The objective is to quantify how uncertainty introduced into angle and distance estimation propagates into final spatial coordinate estimates.

Consider a user located relative to the access point using polar coordinates (d, θ), where $d > 0$ denotes propagation distance and θ denotes AoA. The corresponding Cartesian coordinates are given by

$$x = d\cos\theta, \qquad y = d\sin\theta \tag{25}$$

Under the proposed privacy mechanism, Gaussian perturbations are independently introduced into both the estimated distance and angle measurements following the Gaussian differential privacy mechanism defined in Equation 10. Let

$$N_d \sim \mathcal{N}(0, \sigma_d^2) \tag{26}$$

$$N_\theta \sim \mathcal{N}(0, \sigma_\theta^2) \tag{27}$$

represent the induced perturbations, where σ_d and σ_θ are defined by the sensitivity-scaled Gaussian mechanism in Equations 15–16. The perturbed estimates become

$$\tilde{d} = d + N_d \tag{28}$$

$$\tilde{\theta} = \theta + N_\theta \tag{29}$$

and the resulting noisy Cartesian coordinates are

$$\tilde{x} = \tilde{d}\cos\tilde{\theta} \tag{30}$$

$$\tilde{y} = \tilde{d}\sin\tilde{\theta} \tag{31}$$

The localization error and its expectation are defined as

$$\|\mathbf{e}\|^2 = (\tilde{x} - x)^2 + (\tilde{y} - y)^2 \tag{32}$$

$$\mathbb{E}[\|\mathbf{e}\|^2] = \mathbb{E}\left[(\tilde{x} - x)^2 + (\tilde{y} - y)^2\right] \tag{33}$$

Without loss of generality, set $\theta = 0$, yielding

$$x = d, \qquad y = 0 \tag{34}$$

so that

$$\tilde{x} = (d + N_d)\cos N_\theta, \qquad \tilde{y} = (d + N_d)\sin N_\theta \tag{35}$$

Substituting into the error expression gives

$$\|\mathbf{e}\|^2 = (d - (d + N_d)\cos N_\theta)^2 + ((d + N_d)\sin N_\theta)^2 \tag{36}$$

$$\|\mathbf{e}\|^2 = d^2 - 2d(d + N_d)\cos N_\theta + (d + N_d)^2 \tag{37}$$

Taking expectation over independent perturbations yields

$$\mathbb{E}[\|\mathbf{e}\|^2] = d^2 - 2d\,\mathbb{E}[(d + N_d)\cos N_\theta] + \mathbb{E}[(d + N_d)^2] \tag{38}$$

Using independence and zero-mean properties,

$$\mathbb{E}[(d + N_d)^2] = d^2 + \sigma_d^2 \tag{39}$$

$$\mathbb{E}[(d + N_d)\cos N_\theta] = d\,\mathbb{E}[\cos N_\theta] \tag{40}$$

For Gaussian $N_\theta \sim \mathcal{N}(0, \sigma_\theta^2)$,

$$\mathbb{E}[\cos N_\theta] = e^{-\sigma_\theta^2/2} \tag{41}$$

Thus,

$$\mathbb{E}[\|\mathbf{e}\|^2] = \sigma_d^2 + 2d^2\left(1 - e^{-\sigma_\theta^2/2}\right) \tag{42}$$

For small perturbations,

$$1 - e^{-\sigma_\theta^2/2} \approx \frac{\sigma_\theta^2}{2} \tag{43}$$

$$\therefore\ \mathbb{E}[\|\mathbf{e}\|^2] \approx \sigma_d^2 + d^2\sigma_\theta^2 \tag{44}$$

Substituting the Gaussian mechanism noise scales defined in Equation 15–16 yields

$$\mathbb{E}[\|\mathbf{e}\|^2] \approx \frac{\left(\Delta_d^2 + d^2\Delta_\theta^2\right) 2\ln(1.25/\delta)}{\epsilon^2} \tag{45}$$

This establishes a direct relationship between differential privacy parameters and expected localization error, showing that uncertainty scales inversely with ϵ^2 while increasing with both distance and angular sensitivities.

6 Experimental Results

This section evaluates the effectiveness of the proposed differentially private WiFi sensing mechanism through experimental analysis of localization error, persistence of angular perturbation after refinement, and consistency of CSI reconstruction behavior. Experiments were performed using real-world WiFi channel measurements obtained from the Wireless Indoor Localization Dataset (WILD) dataset[5], and CSI-based localization traces following standard WiFi sensing methodologies[15, 17]. The results demonstrate that the proposed mechanism produces localization uncertainty consistent with the derived theoretical bounds while maintaining physically plausible channel structure after perturbation and reconstruction.

Unlike traditional localization systems that focus on improving absolute positioning accuracy against prior systems, the goal of this work is to validate the derived theoretical relationship between differentially private perturbations and localization error. Therefore, the evaluation is designed as a self-consistency study where empirical measurements are directly compared against the derived error bound rather than against external localization baselines.

6.1 Experimental Noise Bounds

To experimentally validate the derived theoretical localization error bound, evaluations were conducted using the `channels_jacobs_July28.mat` dataset from the WILD dataset. Starting from the first packet in the dataset, every 100th packet was sampled in order to reduce temporal correlation between consecutive CSI measurements while capturing environmental diversity across the dataset[15].

For each sampled packet, the proposed differentially private phase perturbation mechanism was applied using varying privacy budgets $\epsilon \in \{0.8, 1, 1.5, 2.2, 3, 4\}$, following the Gaussian differential privacy mechanism[1, 10]. For each privacy budget, Gaussian perturbations were independently sampled according to the mechanism described previously and injected into both the distance and angular phase steering components of the CSI measurements. The resulting noisy multipath profiles were then converted into Cartesian coordinate estimates, and localization error was measured using mean squared error (MSE) in the XY-coordinate space. [1]

To reduce stochastic variance caused by random Gaussian sampling, the experiment was repeated 50 times for each privacy budget for each sampled packet and the resulting localization errors were averaged. Furthermore, injected noise variance was clipped to three standard deviations to avoid outlier influence on the results. Figure 1 shows the distribution of observed MSE values across the tested privacy budgets.

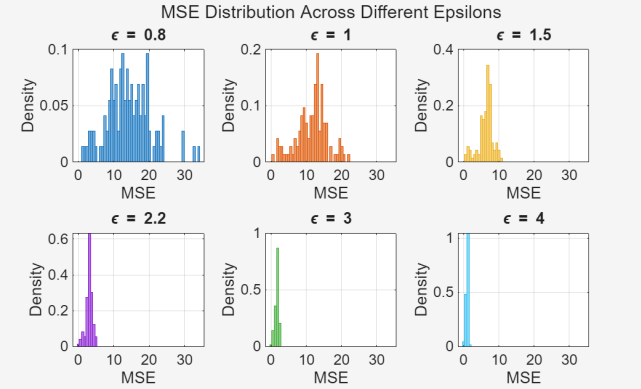

Figure 1: Distribution of localization mean squared error under varying differential privacy budgets ϵ. Lower values of ϵ introduce stronger perturbations and correspondingly larger localization uncertainty.

The results closely follow the theoretical relationship derived in the analytical bound:

$$\mathbb{E}[\|\mathbf{e}\|^2] \approx \frac{(\Delta_d^2 + d^2\Delta_\theta^2)2\ln(1.25/\delta)}{\epsilon^2} \tag{46}$$

which predicts inverse quadratic scaling between localization error and the privacy budget ϵ[9, 10, 21]. As expected, smaller privacy budgets produce significantly larger localization uncertainty due to increased Gaussian perturbation magnitude. Additionally, the variance of the localization error distributions decreases as ϵ increases, reflecting reduced stochastic influence under weaker privacy constraints.

Although individual trials occasionally exceed the predicted expected bound due to the probabilistic nature of Gaussian sampling, the aggregate experimental behavior remains strongly consistent with the derived theoretical approximation. These results demonstrate that the proposed perturbation mechanism provides controllable and mathematically predictable spatial uncertainty through direct manipulation of CSI phase structure.

6.2 Persistence of AoA Error

An important consideration for privacy-preserving WiFi sensing systems is whether an attacker can remove or refine perturbations through post-processing techniques. Many WiFi sensing systems utilize multiple access points and triangulate user positions through angle-of-arrival (AoA) estimation. To evaluate robustness, we test whether injected perturbations persist after refinement.

[1] All implementation code and experimental evaluation scripts used in this work were written in MATLAB and are publicly available for reproducibility at: https://github.com/WIRES-UB/CSI_to_DLocFeatures.git.

To evaluate perturbation persistence, the proposed mechanism was applied to CSI measurements prior to AoA estimation and multipath profile generation. The resulting noisy profiles were subjected to several refinement operations intended to simulate attacker recovery attempts, including common-phase removal and subspace-based denoising techniques such as singular value decomposition (SVD) filtering and phase slope regression[15, 30].

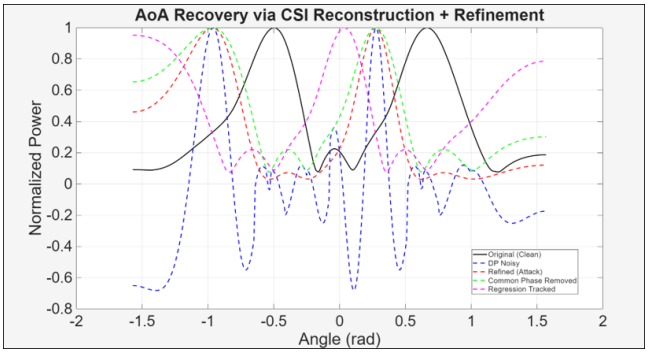

Figure 2: AoA recovery experiment comparing original CSI localization peaks against differentially private perturbations and post-processing refinement attempts.

The clean localization profile exhibits sharp and well-defined peaks corresponding to dominant propagation paths. After applying the proposed perturbation mechanism, the AoA response becomes broadened and shifted, introducing ambiguity into direction estimation. While refinement operations partially recover coarse structure, the original localization peaks are not fully restored, and residual phase inconsistencies continue to degrade precise angle estimation.

6.3 Spoofing Raw CSI

In addition to perturbing localization outputs, the proposed framework introduces a reverse reconstruction mechanism capable of generating plausible CSI measurements from perturbed multipath profiles. This reconstruction is intended to obscure the existence of the perturbation mechanism by producing channel measurements that remain physically consistent with the modified localization output.

To evaluate reconstruction quality, a noisy angle-distance multipath profile generated using the proposed mechanism was reconstructed into CSI using the steering-vector formulation described in prior work on CSI-based localization[17]. The reconstructed CSI was then passed back through the original localization pipeline to produce a second multipath profile. Figure 3 compares the original noisy multipath profile against the reconstructed CSI projection.

The reconstructed multipath profile preserves the dominant structure of the original perturbed heatmap, including the locations and relative intensities of major peaks. This demonstrates that the reverse reconstruction mechanism produces CSI measurements whose forward projection closely matches the intended perturbed sensing output.

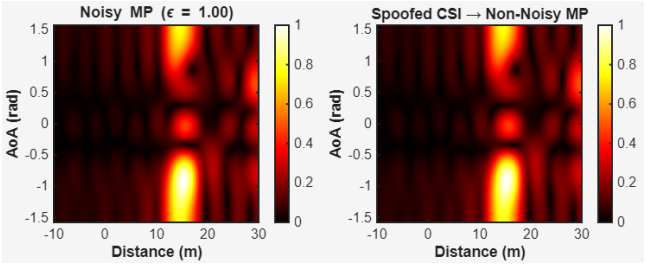

Figure 3: Comparison between a noisy multipath profile and the resulting multipath reconstruction obtained after CSI spoofing and re-projection through the localization pipeline.

From a privacy perspective, this is significant because it enables generation of physically plausible CSI that is consistent with observed localization behavior. Prior work on WiFi sensing and CSI-based inference highlights that such channel statistics remain sufficient for accurate localization in the absence of perturbations[16, 20].

7 Related Works

7.1 Wireless Localization

Indoor localization and wireless sensing technologies have become increasingly common as they provide location awareness without requiring dedicated sensing infrastructure [22, 34]. While GPS performs well outdoors, its performance degrades significantly indoors due to signal attenuation, multipath propagation, and frequent non-line-of-sight conditions [14, 19]. As a result, a wide range of alternative localization approaches have been proposed, including Bluetooth beacons, RFID, ultra-wideband (UWB), inertial sensing, vision-based tracking, and WiFi-based systems [16, 31]. Among these, WiFi-based localization is particularly attractive because it leverages already-deployed infrastructure and avoids the need for additional hardware deployment [7, 26].

Bluetooth beacon-based localization is a widely studied approach that relies on low-energy Bluetooth transmitters placed at known locations, where receivers estimate position using RSSI measurements [13]. Although BLE systems are low-cost and easy to deploy, RSSI-based measurements are highly sensitive to multipath fading, device orientation, and environmental dynamics, which leads to significant instability in localization performance [13, 16]. Additionally,

achieving high accuracy typically requires dense beacon deployments, increasing infrastructure and maintenance overhead in large-scale environments [31].

Camera-based localization systems provide an alternative sensing modality that can achieve high spatial accuracy under controlled conditions, but they require line-of-sight and are sensitive to illumination changes and occlusions [35]. More importantly, vision-based systems introduce significant privacy concerns because they continuously capture visual data of the environment, enabling potential identification and tracking of non-consenting users [35]. These limitations restrict their applicability in public or privacy-sensitive environments despite their strong performance in controlled settings.

WiFi-based localization systems originally relied on RSSI fingerprinting, where signal strength measurements are collected during an offline calibration phase and later matched to estimate user location [8, 32]. While these systems are inexpensive and simple to deploy, they suffer from poor robustness in dynamic environments due to multipath effects and temporal variations in signal strength [13, 19]. Furthermore, fingerprinting approaches require extensive site surveys and periodic recalibration, which limits scalability in large or frequently changing environments [16]. These challenges motivated the development of more physically expressive wireless representations.

CSI-based methods have emerged as stronger alternatives to RSSI by exposing fine-grained amplitude and phase information across subcarriers and multiple antennas [15]. Unlike RSSI, CSI captures detailed multipath structure, enabling significantly improved spatial resolution and robustness in complex indoor environments [26, 29]. CSI-based methods have enabled a wide range of applications including localization, motion tracking, and device-free sensing, extending beyond traditional positioning tasks [20, 28]. By exploiting phase coherence and spatial diversity, CSI-based systems can infer fine-grained spatial information that is not observable in RSSI-based systems [7].

Several landmark systems have demonstrated the effectiveness of CSI for high-accuracy localization. SpotFi introduced joint angle-of-arrival and time-of-flight estimation using commodity WiFi hardware, achieving decimeter-level localization accuracy by exploiting multipath structure [17]. Chronos further demonstrated sub-nanosecond ToF estimation using frequency hopping across channels, enabling precise distance estimation without specialized hardware [24]. ArrayTrack leveraged antenna arrays and spatial signal processing techniques to achieve fine-grained indoor tracking accuracy [30]. Together, these systems demonstrate that multipath propagation, traditionally treated as interference, can instead be exploited as a rich source of spatial information for localization [17, 24].

Beyond device-based localization, device-free WiFi sensing systems have been developed to infer human presence, motion, and activity directly from wireless channel variations [3, 28]. These systems eliminate the need for users to carry devices by exploiting changes in signal propagation caused by human movement and environmental dynamics [20]. Prior work has demonstrated applications including gait recognition, gesture recognition, occupancy estimation, and human identification using CSI measurements [4, 27]. More recent surveys further categorize these systems into passive sensing, activity recognition, and behavioral inference frameworks [20, 28]. However, as sensing accuracy improves, the same CSI measurements increasingly expose sensitive user behavior and spatial information [36].

The fine-grained nature of CSI makes it particularly powerful for inference tasks, but also introduces significant privacy risks. Wireless signals propagate through walls and obstacles, enabling sensing systems to observe users without direct interaction or consent [2]. Prior work has shown that WiFi-based sensing systems can reconstruct human motion, infer identity from gait patterns, and detect fine-grained activity signatures from CSI measurements [4, 27]. As a result, WiFi sensing systems inherently introduce a tension between sensing utility and user privacy, motivating the need for privacy-preserving approaches that operate directly on physical-layer measurements rather than application-level data [36].

7.2 Privacy Implementations

Privacy risks in WiFi-based localization systems have been shown in a variety of implemented CSI-based sensing systems, where CSI is directly used to infer user position through AoA, ToF, or fingerprinting methods [17, 24, 30]. Systems such as SpotFi explicitly recover AoA and ToF information from CSI measurements to estimate user location at decimeter-level accuracy using commodity WiFi hardware [17]. Similarly, Chronos implements fine-grained ToF estimation by sweeping across multiple frequency channels to improve temporal resolution, enabling accurate distance estimation from standard WiFi cards [24]. These implementations demonstrate that CSI inherently encodes precise spatial information, which can be exploited for localization without requiring user participation.

One of the earliest practical privacy protections in WiFi sensing systems is access restriction at the driver or firmware level, where CSI extraction is only exposed through modified network interfaces such as the Intel 802.11n CSI Tool [15]. In these implementations, raw CSI samples are only accessible to privileged applications, effectively limiting who can perform localization. However, systems built on top of this

infrastructure still directly use unmodified CSI for localization tasks, meaning that privacy is enforced only through access control rather than data protection. Once CSI is exposed, systems such as ArrayTrack and SpotFi demonstrate that it can be directly used for accurate spatial inference without further safeguards [17, 30].

CSI fingerprinting systems such as FIFS and deep learning-based localization frameworks explicitly store or learn mappings from CSI measurements to physical locations [26, 29]. In these implementations, localization is performed by training classifiers or regression models on CSI feature vectors collected from known reference points. While these systems achieve high localization accuracy, they also implicitly encode location information into model parameters or fingerprint databases, making them vulnerable to inversion or inference attacks if the trained models or datasets are exposed. This highlights that privacy risks arise not only from raw CSI but also from learned representations derived from it.

Device-free WiFi sensing systems further demonstrate privacy risks in implemented localization pipelines by enabling position inference without requiring users to carry any transmitting device. Systems such as WiWho and other CSI-based passive localization frameworks infer user position or identity based on perturbations in wireless channels caused by human presence and movement [27, 28]. These implementations show that even when no explicit device is associated with a user, CSI measurements can still be used to reconstruct spatial trajectories; while this paper focuses solely on systems with device involvement, this point reinforces the inherent privacy leakage in wireless localization systems.

7.3 Differential Privacy

A major advantage of differential privacy over heuristic privacy-preserving approaches is the ability to provide mathematically quantifiable guarantees on information leakage through probabilistic noise mechanisms [10, 11]. In particular, the Gaussian mechanism has become one of the most widely studied and implemented approaches for achieving (ϵ, δ)-differential privacy due to its strong composition properties and compatibility with high-dimensional optimization problems [1, 11]. Prior work has shown that Gaussian noise calibrated according to the sensitivity of a function can bound the distinguishability between neighboring datasets while still preserving sufficient utility for downstream inference tasks [9, 21]. Because of these properties, Gaussian-based DP mechanisms have become foundational in modern privacy-preserving machine learning systems, specifically in iterative optimization frameworks such as DP-SGD where repeated queries to sensitive data require careful accounting of cumulative privacy loss [1, 12]. These theoretical guarantees make Gaussian differential privacy particularly attractive for wireless localization systems, where privacy mechanisms must balance formal privacy protections against degraded utility in localization.

Differential privacy has been most commonly implemented in machine learning systems through differentially private stochastic gradient descent (DP-SGD), which introduces gradient clipping and Gaussian noise during model training [1]. In practical implementations, DP-SGD is used to train neural networks while limiting the influence of any individual training sample on the final model parameters. This mechanism has been deployed in large-scale systems for privacy-preserving analytics and learning tasks, demonstrating that DP can be effectively integrated into iterative optimization pipelines.

To improve practical utility in deployed systems, Rényi Differential Privacy (RDP) has been introduced as a tighter privacy accounting mechanism that tracks cumulative privacy loss more efficiently than classical (ϵ, δ)-DP [21]. In implemented machine learning systems, RDP is commonly used to monitor privacy budgets during long training procedures involving repeated gradient updates, enabling more accurate control over privacy-utility tradeoffs compared to standard composition theorems.

In applied machine learning systems, differential privacy has been integrated into end-to-end training pipelines where noise injection is performed either at the input level, gradient level, or output level depending on system design constraints [1, 11]. These implementations demonstrate that DP is not limited to theoretical formulations but can be embedded directly into practical learning systems, particularly in settings where data sensitivity is high and repeated queries to private data occur.

Despite these advances, existing DP implementations are primarily designed for structured data domains such as tabular datasets, images, or gradients in deep learning models [1]. On the other hand, CSI-based localization systems operate on highly structured physical-layer signals where small perturbations in phase or amplitude can lead to disproportionate effects on AoA and ToF estimation outcomes [17, 24]. As a result, directly applying standard DP implementations to wireless localization pipelines is not straightforward and requires mechanisms that account for the structure of wireless propagation itself.

8 Conclusion

This work presented a privacy-preserving framework for CSI-based WiFi localization through the application of differentially private phase perturbations directly within the sensing pipeline. By introducing calibrated Gaussian perturbations

into the steering-vector phase structure used for angle and distance estimation, the proposed mechanism produces controlled uncertainty within distance-angle multipath profiles while preserving the overall physical structure of the wireless environment. Unlike traditional privacy mechanisms that focus primarily on encryption, access control, or anonymization in the communication or data layer[11, 12, 36], the proposed approach directly targets the sensing information encoded within wireless channel measurements themselves.

This work derives an explicit analytical relationship between differential privacy parameters and downstream Cartesian localization error. Theoretical analysis demonstrated that the expected localization error scales proportionally with both angular and distance sensitivity while increasing inversely with the square of the privacy budget ϵ[9, 10, 21]. Experimental evaluation using real-world CSI measurements further validates the derived bounds and shows strong agreement between theoretical predictions and observed localization behavior across multiple privacy budget configurations.

The experimental results additionally demonstrate that phase perturbations remain effective even after refinement and post-processing operations commonly used in CSI-based localization systems[15, 30]. Although adversarial recovery methods are capable of partially restoring coarse spatial structure, the perturbations continue to introduce persistent localization uncertainty due to the sensitivity of steering-vector beamforming to coherent phase alignment.

Furthermore, the proposed reverse reconstruction mechanism successfully generates physically plausible CSI measurements whose forward transformation reproduces the intended perturbed multipath profiles, leveraging standard CSI modeling and beamforming formulations from prior WiFi sensing systems[17, 20]. This helps further obscure the existence of the underlying privacy mechanism while maintaining consistency with observed channel statistics.

While the results presented in this work establish a strong foundation for privacy-preserving WiFi sensing, several important directions remain for future research. One promising direction involves implementing the proposed mechanism directly within commodity router firmware or lightweight access-point hardware in order to evaluate real-time deployment feasibility and computational overhead in practical wireless environments[15, 24]. Prior work has demonstrated that practical WiFi sensing systems often require careful optimization at the hardware and firmware level to achieve real-time performance under commodity constraints[24].

Additionally, future work may extend this framework to multi-router or distributed localization systems where multiple synchronized access points jointly estimate user position[17, 30]. Such distributed sensing architectures are widely used in high-accuracy indoor localization systems to improve spatial resolution and robustness through spatial diversity.

In such systems, deriving formal bounds on angle-of-arrival recovery and localization uncertainty under coordinated sensing architectures remains an open problem, particularly when multiple separated observations are combined under noise or perturbation models[31, 34].

Overall, this work demonstrates that differential privacy can be integrated into CSI-based WiFi sensing systems through physically meaningful perturbations applied at the channel level. By establishing both theoretical error guarantees and experimentally validated privacy behavior, the proposed framework provides a foundation for future wireless sensing systems capable of controlling the utility–privacy tradeoff of sensing systems.

9 Acknowledgement

In the writing of this paper, ChatGPT was used translate handwritten mathematical notes and MATLAB scripts into the syntax used for formatting equations and expressions throughout.

References

[1] M. Abadi et al. 2016. Deep Learning with Differential Privacy. In *Proceedings of ACM CCS*. 308–318.

[2] Fadel Adib, Zachary Kabelac, Dina Katabi, and Robert Miller. 2015. Capturing the Human Figure Through a Wall. In *Proceedings of NSDI*. 1–13.

[3] Fadel Adib and Dina Katabi. 2013. See Through Walls with WiFi. In *Proceedings of ACM SIGCOMM*. 75–86.

[4] Khalid Ali et al. 2015. Keystroke Recognition Using WiFi Signals. In *Proceedings of ACM MobiCom*. 90–102.

[5] Aditya Arun, Akshaj Bharadwaj, and Roshan Ayyalasomayajula. 2022. Wi-Fi Indoor Localization Dataset (WILD-v2). https://kaggle.com/competitions/wild-v2. Kaggle.

[6] Roshan Ayyalasomayajula, Aditya Arun, Wei Sun, and Dinesh Bharadia. 2022. Users are Closer than they Appear: Protecting User Location from WiFi APs. *arXiv preprint arXiv:2211.10014* (2022). doi:10.48550/arXiv.2211.10014

[7] Roshan Ayyalasomayajula, Aditya Arun, Chenfeng Wu, Sanatan Sharma, Abhishek Rajkumar Sethi, Deepak Vasisht, and Dinesh Bharadia. 2020. Deep learning based wireless localization for indoor navigation. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–14.

[8] Paramvir Bahl and Venkata Padmanabhan. 2000. RADAR: An In-Building RF-Based User Location and Tracking System. In *Proceedings of IEEE INFOCOM*. 775–784.

[9] Mark Bun and Thomas Steinke. 2016. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Theory of Cryptography Conference (TCC)*. Springer, 635–658. doi:10.1007/978-3-662-53641-4_24

[10] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography Conference (TCC)*. Springer, 265–284. doi:10.1007/11681878_14

[11] Cynthia Dwork and Aaron Roth. 2014. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer

Science, Vol. 9. Now Publishers Inc. 211–407 pages. doi:10.1561/0400000042

[12] Ahmed El Ouadrhiri and Ahmed Abdelhadi. 2022. Differential Privacy for Deep and Federated Learning: A Survey. *IEEE Access* 10 (2022), 22359–22380. doi:10.1109/ACCESS.2022.3151671

[13] R. Faragher and R. Harle. 2015. Location Fingerprinting With Bluetooth Low Energy Beacons. *IEEE Journal on Selected Areas in Communications* 33, 11 (2015), 2418–2428. doi:10.1109/JSAC.2015.2430281

[14] Yanying Gu, Anthony Lo, and Ignas Niemegeers. 2009. A Survey of Indoor Positioning Systems for Wireless Personal Networks. *IEEE Communications Surveys & Tutorials* 11, 1 (2009), 13–32.

[15] Daniel Halperin et al. 2011. Tool Release: Gathering 802.11n Traces with Channel State Information. In *ACM SIGCOMM Computer Communication Review*, Vol. 41. 53.

[16] Suining He and S.-H. Gary Chan. 2016. Wi-Fi Fingerprint-Based Indoor Positioning: Recent Advances and Comparisons. *IEEE Communications Surveys & Tutorials* 18, 1 (2016), 466–490.

[17] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. In *Proceedings of ACM SIGCOMM*. 269–282. doi:10.1145/2785956.2787487

[18] Hamid Krim and Mats Viberg. 1996. Two Decades of Array Signal Processing Research: The Parametric Approach. *IEEE Signal Processing Magazine* 13, 4 (1996), 67–94. doi:10.1109/79.526899

[19] Hui Liu, Houshang Darabi, Pat Banerjee, and Jing Liu. 2007. Survey of Wireless Indoor Positioning Techniques and Systems. *IEEE Transactions on Systems, Man, and Cybernetics* 37, 6 (2007), 1067–1080.

[20] Yongsen Ma et al. 2019. WiFi Sensing with Channel State Information: A Survey. *Comput. Surveys* 52, 3 (2019), 1–36.

[21] Ilya Mironov. 2017. Rényi Differential Privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. IEEE, 263–275. doi:10.1109/CSF.2017.11

[22] Huthaifa Obeidat, Waseem Shuaieb, Omar Obeidat, and Raed Abd-Alhameed. 2021. A Review of Indoor Localization Techniques and Wireless Technologies. *Wireless Personal Communications* 119, 1 (2021), 289–327.

[23] Ralph O. Schmidt. 1986. Multiple Emitter Location and Signal Parameter Estimation. *IEEE Transactions on Antennas and Propagation* 34, 3 (1986), 276–280. doi:10.1109/TAP.1986.1143830

[24] Deepak Vasisht, Swarun Kumar, and Dina Katabi. 2016. Chronos: Sub-Nanosecond Time of Flight on Commercial Wi-Fi Cards. In *Proceedings of ACM MobiCom*. 1–14.

[25] Wei Wang et al. 2017. E-Eye: Device-Free Location-Oriented Activity Identification Using Fine-Grained WiFi Signatures. *IEEE Transactions on Mobile Computing* 17, 11, 2568–2582.

[26] Xuyu Wang et al. 2016. CSI-Based Fingerprinting for Indoor Localization: A Deep Learning Approach. *IEEE Transactions on Vehicular Technology* 66, 1 (2016), 763–776.

[27] Chenshu Wu et al. 2016. WiWho: WiFi-Based Person Identification Using Human Gait. In *Proceedings of IEEE INFOCOM*. 1–9.

[28] Chenshu Wu et al. 2017. Device-Free Wireless Sensing and Localization with WiFi Signals: A Survey. *IEEE Transactions on Vehicular Technology* 67, 9 (2017), 7967–7986.

[29] Jiang Xiao et al. 2012. FIFS: Fine-Grained Indoor Fingerprinting System. In *Proceedings of IEEE ICCCN*. 1–7.

[30] Jie Xiong and Kyle Jamieson. 2013. ArrayTrack: A Fine-Grained Indoor Location System. In *Proceedings of USENIX NSDI*. 71–84.

[31] Ahmad Yassin et al. 2017. Recent Advances in Indoor Localization: A Survey on Theoretical Approaches and Applications. *IEEE Communications Surveys & Tutorials* 19, 2 (2017), 1327–1346.

[32] Moustafa Youssef and Ashok Agrawala. 2005. The Horus WLAN Location Determination System. In *Proceedings of ACM MobiSys*. 205–218.

[33] Heejung Yu and Taejoon Kim. 2014. Beamforming Transmission in IEEE 802.11ac under Time-Varying Channels. *The Scientific World Journal* 2014 (2014), 920937. doi:10.1155/2014/920937

[34] Faheem Zafari, Athanasios Gkelias, and Kin K. Leung. 2019. A Survey of Indoor Localization Systems and Technologies. *IEEE Communications Surveys & Tutorials* 21, 3 (2019), 2568–2599.

[35] Y. Zhan et al. 2020. Vision-Based Indoor Localization: A Survey. *IEEE Access* 8 (2020), 194353–194375. doi:10.1109/ACCESS.2020.3032777

[36] Jie Zhang et al. 2018. Privacy Risks in WiFi-Based Sensing Systems. *IEEE Security & Privacy* 16, 5 (2018), 70–77.