

Guaranteeing Spoof-Resilient Multi-Robot Networks

Stephanie Gil^{†1} Swarun Kumar^{†1} Mark Mazumder² Dina Katabi¹ Daniela Rus¹
¹Massachusetts Institute of Technology ²MIT Lincoln Laboratory
 {swarun, sgil, dk, rus}@mit.edu mazumder@ll.mit.edu
[†]Co-primary authors

Abstract—Multi-robot networks use wireless communication to provide wide-ranging services such as aerial surveillance and unmanned delivery. However, effective coordination between multiple robots requires trust, making them particularly vulnerable to cyber-attacks. Specifically, such networks can be gravely disrupted by the Sybil attack, where even a single malicious robot can spoof a large number of fake clients. This paper proposes a new solution to defend against the Sybil attack, without requiring expensive cryptographic key-distribution. Our core contribution is a novel algorithm implemented on commercial Wi-Fi radios that can “sense” spoofers using the physics of wireless signals. We derive theoretical guarantees on how this algorithm bounds the impact of the Sybil Attack on a broad class of robotic coverage problems. We experimentally validate our claims using a team of AscTec quadrotor servers and iRobot Create ground clients, and demonstrate spoofer detection rates over 96%.

I. INTRODUCTION

Multi-robot networks rely on wireless communication to enable a wide range of tasks and applications: coverage [28, 5, 31], disaster management [6], surveillance [3], and consensus [27] to name a few. The future promises an increasing trend in this direction, such as delivery drones which transport goods (e.g., Amazon Prime Air [1]) or traffic rerouting algorithms (e.g., Google Maps Navigation) that rely on broadcasted user locations to achieve their goals. Effective coordination, however, requires trust. In order for these multi-robot systems to perform their tasks optimally, transmitted data is often assumed to be accurate and trustworthy; an assumption that is easy to break. A particularly challenging attack on this assumption is the so-called “Sybil attack.”

In a Sybil attack a malicious agent generates (or spoofs) a large number of false identities to gain a disproportionate influence in the network.¹ These attacks are notoriously easy to implement [33] and can be detrimental to multi-robot networks. An example of this is coverage, where an adversarial client can spoof a cluster of clients in its vicinity in order to create a high local demand, in turn denying service to legitimate clients (Figure 1). Although a vast body of literature is dedicated to cybersecurity in general multi-node networks (e.g., a wired LAN), the same is not true for multi-robot networks [15, 30], leaving them largely vulnerable to attack. This is because many characteristics unique to robotic networks make security more challenging; for example, traditional key passing or cryptographic authentication is difficult to maintain




Fig. 1: **Sybil Attack on Coverage:** A server robot provides locational coverage to legitimate clients when no attack is present. In a Sybil attack, an adversary spoofs many fake clients to draw away coverage from the legitimate clients.

due to the highly dynamic and distributed nature of multi-robot teams where clients often enter and exit the network.

This paper addresses the challenge of guarding against Sybil attacks in multi-robot networks. We focus on the general class of problems where a group of server robots coordinate to provide some service using the broadcasted locations of a group of client robots. Our core contribution is a novel algorithm that analyzes the received wireless signals to detect the presence of spoofed clients spawned by adversaries. We call this a “virtual spoofer sensor” as we do not use specialized hardware nor encrypted key exchange, but rather a commercial Wi-Fi card and software to implement our solution. Our virtual sensor leverages the rich physical information already present in wireless signals. At a high level, as wireless signals propagate, they interact with the environment via scattering and absorption from objects along the traversed paths. Carefully processed, these signals can provide a unique signature or “spatial fingerprint” for each client, measuring the power of the signal received along each spatial direction (Fig. 2). Unlike message contents such as reported IDs or locations which adversaries can manipulate, spatial fingerprints rely on physical signal interactions that cannot be exactly predicted [13, 24].

Using these derived fingerprints, we show that a confidence weight, $\alpha \in (0, 1)$ can be obtained for each client in the network. We prove that these confidence weights have a desirable property where legitimate clients have an expected confidence weight close to one, while spoofed clients will have an expected confidence weight close to zero. A particularly attractive feature of confidence weight α is that it can be readily integrated as a per-client weighting function into a wide variety of multi-robot controllers. More importantly, the analytical bounds on these weights can provably limit the ill-effects of spoofers on the performance of these controllers. This paper demonstrates this capability in the context of the

¹Please refer to [7, 26] for a detailed treatment of this class of cyber attacks.

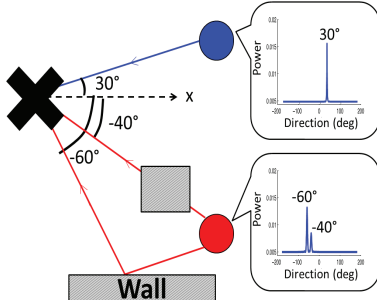


Fig. 2: **Spatial Fingerprints:** A quadrotor server measures the directional signal strength of each client (here, simplified to 2-D). The blue client has one line-of-sight peak; the other, 2 signal paths.

well-known locational coverage algorithm [5, 31].

We provide an extensive experimental evaluation of our theoretical claims using a heterogeneous team of air/ground robots consisting of two AscTec Hummingbird platforms and ten iRobot Create platforms. We conduct our experiments in general indoor settings with randomly placed clients and demonstrate a spoofer detection rate of 96%. For the case of coverage we find that the converged positions of the service robots is on average 3 cm from optimal even when more than 75% of total clients in the network are spoofed.

Contributions of this paper: We develop a virtual sensor for spoofing detection which provides performance guarantees in the presence of Sybil attacks and is applicable to a broad class of problems in distributed robotics. We show that the influence of spoofers is analytically bounded under our system in a coverage context, where each robotic node providing coverage remains within a radius of its position in the absence of an attack. Our theoretical results are validated extensively through experiments in diverse settings.

II. RELATED WORK

The problem of Sybil attacks has been studied in general multi-node, often static, networks, and many tools have been developed for these settings. Past work falls under three categories: (1) Cryptographic authentication schemes can be used to prevent Sybil attacks (Table 7 in [39]). These require trusted central authorities and computationally expensive distributed key management, to account for dynamic clients that enter and leave the network [39]. (2) Non-cryptographic techniques in the wireless networking community leverage wireless physical-layer information to detect spoofed client identities or falsified locations [16, 42, 40, 41]. These rely on bulky and expensive hardware like large multi-antenna arrays, that cannot be mounted on small robotic platforms. (3) Recent techniques have attempted to use wireless signal information like received signal strength (RSSI) [37, 29] and channel state information [22]. Such techniques need clients to remain static, since mobility can cause wireless channels to fluctuate rapidly [2]. In addition, they are susceptible to power-scaling attacks, where clients scale power differently to imitate different users. In sum, the above systems share one or more of the following characteristics making them ill-suited to multi-robot networks: (1) require computationally-intensive

key management; (2) rely on bulky and expensive hardware; (3) assume static networks. Indeed past work has highlighted the gravity and apparent sparsity of solutions to cyber-security threats in multi-robot networks [15, 30, 4].

Unlike past work, our solution has three attributes that particularly suit multi-robot networks. (1) It captures physical properties of wireless signals and therefore does not require distributed key management. (2) It relies on cheap commodity Wi-Fi radios, unlike hardware-based solutions [40, 42]. (3) It is robust to client mobility and power-scaling attacks.

Finally, our system builds on Synthetic Aperture Radar (SAR) to construct signal fingerprints [8]. SAR has been widely used for radar imaging [8, 17] and indoor positioning [19, 18, 36, 12]. In contrast, this paper builds upon SAR to provide cyber-security to multi-robot networks. In doing so, it provides theoretical security guarantees that are validated experimentally. These integrate readily with performance guarantees of existing multi-robot controllers, like the well-known robotic coverage controllers [5, 31] as shown in Sec. §VI.

III. PROBLEM STATEMENT

This paper focuses on problems where the knowledge of agent positions facilitates some collaborative task. Specifically, it assumes two groups of agents, “clients” requiring some type of location-based service such as coverage or goods delivery and “servers” whose positions are optimized in order to provide the service to its clients. Let $P := \{p_1, \dots, p_c\}$ denote the client positions in \mathbb{R}^3 . Let $X := \{x_1, \dots, x_m\}$ be the positions of the servers in \mathbb{R}^3 and the notation $[m] = \{1, \dots, m\}$ denote their indices. We consider the case where a subset of the clients, $S \subset P$ (with $s := |S|$) are “spoofed” clients.

Definition 3.1 (Spoofed Client): A single malicious client may generate multiple unique identities, each with a fabricated position. Each generated, or “spawned” identity is considered a *spoofed client*. By spoofing multiple clients, the malicious client gains a disproportionate influence in the network. All clients which are not spoofed are considered *legitimate clients*.

Threat Model: Our threat model considers one or more adversarial robot clients with one Wi-Fi antenna each. The adversaries can be mobile and scale power on a per-packet basis. We only consider adversarial clients.² Adversarial clients perform the “Sybil Attack” to forge packets emulating s non-existent clients, where s can exceed the number of legitimate clients. More formally:

Definition 3.2 (Sybil Attack): Define a network of client and server positions as $P \cup X$, where a subset S of the clients are spoofed, such that $P = S \cup \tilde{S}$. We assume that set P is known but knowledge of which clients are spoofed (i.e., in S) is unknown. This attack is called a “Sybil Attack.”

To counter the Sybil attack, this paper has two objectives. First, we find a relation capturing directional signal strength between a client i and a server l . We seek a mapping $F_{il} : [0, \frac{\pi}{2}] \times [0, 2\pi] \mapsto \mathbb{R}$ such that for any 3D direction (θ, ϕ) defined in Fig. 4, the value $F_{il}(\theta, \phi)$ is the power of

²The case of adversarial server robots is left for future work although many of the concepts in the current paper are extensible to this case as well.

the received signal from client i along that direction. Using this mapping, or “fingerprint”, our first problem is to derive a *confidence weight* whose expectation is provably bounded near 1 for legitimate clients and near 0 for spoofed clients. Further, we wish to find these bounds analytically from problem parameters like the signal-to-noise ratio of the received wireless signal. We summarize this objective as Problem 1 below:

Problem 1: Spoofers Detection Let \mathcal{F}_i be the set of fingerprints measured from all clients $j \in [c]$ and servers $l \in [m]$ in the neighborhood, \mathcal{N}_i , of client i .³ Here, a neighborhood of client i , \mathcal{N}_i , are all agents that can receive Wi-Fi transmissions sent by client i . Using \mathcal{F}_i , derive a confidence weight $\alpha_i(\mathcal{F}_i) \in (0, 1)$ and a threshold $\omega_i(\sigma_i^2) > 0$ where σ_i^2 represents error variances such as the signal-to-noise ratio that are assumed to be given. Find $\omega_i(\cdot)$ to have the provable property of differentiating spoofed clients whereby spoofed clients are bounded below this threshold, i.e., $E[\alpha_i] \leq \omega$, and legitimate clients are bounded above this threshold $E[\alpha_i] \geq 1 - \omega$.

Our second objective is to apply our spoofer detection method to multi-robot control problems. We consider the well-known coverage problem in [5, 31]. We show that by integrating the confidence weight from Problem 1, we can analytically bound the error in performance caused by spoofed clients in the network. We consider the coverage problem where an importance function is defined over an environment and where the positions of the clients correspond to peaks in the importance function. Here, servers position themselves to maximize their proximity to these peaks, to improve their coverage over client robots. If $C_V = \{x_1^*, \dots, x_m^*\}$ is the set of server positions optimized by the coverage controller with zero spoofers, we wish to guarantee that server positions optimized with spoofers present, C_{V_α} , is “close” to C_V . We state this second objective more specifically as Problem 2 below:

Problem 2: Sybil-resilience in Multi-Robot Coverage

Consider a locational coverage problem where an importance function $\rho(q) > 0$ is defined over an environment $\mathcal{Q} \subset \mathbb{R}^3$ and $q \in \mathcal{Q}$. Specifically, consider an importance function that can be decomposed into terms, $\rho_i(q)$, depending on each client’s position, $i \in [c]$ (for example, each client position corresponds to a peak), i.e., $\rho(q) = \rho_1(q) + \dots + \rho_c(q)$. Let $C_V = \{x_1^*, \dots, x_m^*\}$ be the set of server positions returned by an optimization of $\rho(q)$ over X , where there are zero spoofed clients in the network. Under a Sybil attack, let $C_{V_\alpha} = \{x_1, \dots, x_m\}$ be the set of server positions returned by an optimization of an α -modified importance function $\rho(q) = \alpha_1 \rho_1(q) + \dots + \alpha_c \rho_c(q)$ where the importance weight terms α_i satisfy the bounds stated in Problem 1. We wish to find an $\epsilon(\mathcal{P}) > 0$ such that the set C_{V_α} is within a distance $\epsilon(\mathcal{P})$ to C_V . C_{V_α} is within a distance $\epsilon(\mathcal{P})$ to C_V if $\forall x \in C_{V_\alpha}$ there exists a unique $y \in C_V$ where $\text{dist}(x, y) < \epsilon(\mathcal{P})$. Here, \mathcal{P} is a set of problem parameters that we wish to find.

³Detecting if a client i is spoofed becomes easier given more servers communicating with i (i.e., a larger neighborhood \mathcal{N}_i). But even with a single server, this determination can be made. A theoretical treatment of this point is given in Sec. §V and experimental results (§VII-A) use as little as one server.

Intuitively, solutions to Problem 2 guarantee that under a Sybil attack, all server positions computed using an α -modified coverage controller are within a computable distance $\epsilon(\mathcal{P})$ from their optimal positions (i.e., in the absence of spoofers). Sec. §VI derives a closed-form for $\epsilon(\mathcal{P})$ and shows the set \mathcal{P} of problem parameters to be the number of spoofers, the footprint of the environment covered, and signal noise.

IV. FINGERPRINTS TO DETECT MALICIOUS CLIENTS

Here we construct a *fingerprint*, a directional signal strength profile for a communicating server-client pair. Our choice of signal fingerprints have many desirable properties that enable us to derive a robust spoof-detection metric: they 1) capture directional information of the transmitted signal source and thus are well-suited for flagging falsely reported client positions, 2) can be obtained for a single server-client pair, unlike location estimation techniques such as triangulation which require multiple servers to coordinate, 3) cannot be manipulated by the client, since the occurrence of each signal path is due to environment reflections, 4) are applicable in complex multipath environments where a transmitted signal is scattered off of walls and objects; since these scattered signals manifest themselves as measurable peaks in the fingerprint, complex multipath contributes significantly to fingerprint uniqueness.

We construct fingerprints using wireless channels h , complex numbers measurable on any wireless device characterizing the attenuation in power and the phase rotation that signals experience as they propagate over the air. These channels also capture the fact that wireless signals are scattered by the environment, arriving at the receiver over (potentially) several different paths [35]. Fig. 3 is an example 2D schematic of a wireless signal traversing from a client robot to a server robot arriving along two separate paths: one attenuated direct path at 40° and one reflected at 60° . If the server robot had a directional antenna, it could obtain a full 3D profile of power of the received signal (i.e., $|h|^2$) along *every* spatial direction. We use such a 3-D profile as a “spatial fingerprint” that can help distinguish between different clients.

Unfortunately directional antennas are composed of large arrays of many antennas that are too bulky for small agile robot platforms. Luckily, a well-known technique called Synthetic Aperture Radar [8] (SAR) can be used to emulate such an antenna using a commodity Wi-Fi radio. Its key idea is to use small local robotic motion, such as spinning in-place, to obtain multiple snapshots of the wireless channel that are then processed like a directional array of antennas. SAR can be implemented using a well-studied signal processing algorithm called MUSIC [14] to obtain spatial fingerprints at each server robot.

Mathematically, we obtain a spatial fingerprint for each wireless link between a server l and client i as a matrix $F_{il} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$. For each spatial path represented as (θ, ϕ) (see Fig. 4), F_{ij} maps to a scalar value representing the signal power received along that path. More formally:

$$F_{il}(\phi, \theta) = 1/|Eig_n(\hat{\mathbf{h}}_{il}\hat{\mathbf{h}}_{il}^\dagger)e^{\sqrt{-1}\Psi_{il}(\phi, \theta)}|^2 \quad (1)$$

Where $\hat{\mathbf{h}}_{il}$ is a vector of the ratio of wireless channel snapshots between two antennas mounted on the body of the server l and

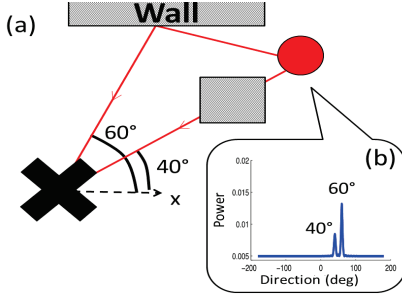


Fig. 3: **Example Signal Fingerprint:** (a) A server (x) receives a client (•) signal on 2 paths: direct along 40° attenuated by an obstacle (shaded) and reflected by a wall along 60°. (b) is a corresponding fingerprint: peak heights at 40° and 60° correspond to their relative attenuations.

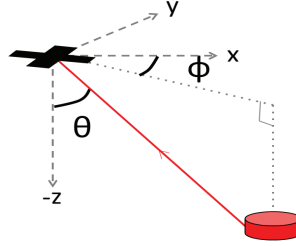


Fig. 4: **3-D Angles:** The figure depicts the notation for the azimuthal angle ϕ and polar angle θ for the direct path from a ground client (•) to aerial server robot (x) in 3 dimensions. More generally, the set of all angles between client i and server l are denoted as Φ_{il}, Θ_{il} respectively.

$\Psi_{il}(\phi, \theta) = \frac{2\pi r}{\lambda} \cos(\phi - \mathbf{B}_1) \sin(\theta - \mathbf{\Gamma}_1)$, λ is the wavelength of the signal and r is the distance between the antennas, $\mathbf{B}_1, \mathbf{\Gamma}_1$ are the server's angular orientation, $Eig_n(\cdot)$ are noise eigenvectors, $(\cdot)^\dagger$ is conjugate transpose, and k is the number of signal eigenvectors, equal to the number of paths.

While our above formulation is derived from MUSIC [14], it varies in one important way: while MUSIC uses a single-antenna channel snapshot h_{il} , we use the channel ratio $\hat{h}_{il} = h_{1il}/h_{2il}$ between two antennas. This modification provides resilience to intentional power scaling by the sender since scaling his transmit power by χ yields a measured ratio $\hat{h}_{il} = \chi h_{1il}/(\chi h_{2il})$; a value unaffected by power scaling.

V. CONSTRUCTING A CLIENT CONFIDENCE WEIGHT

Given a client fingerprint $F_{il}(\phi, \theta)$ for each client i relative to a robotic server l , we wish to generate a confidence weight $\alpha_i \in [0, 1]$ that approaches 1 for legitimate clients, and 0 otherwise. We achieve this by defining α_i as the product of two terms β_i and γ_{ij} that go to 0 if a client reports a falsified location or has the same fingerprint as another client j respectively. In particular, β_i is termed the *honesty* metric and is the likelihood (Eq. (2)) that client i is indeed along its reported direction (ϕ_{il}, θ_{il}) with respect to each server l in its neighborhood. The second term γ_{ij} is the *similarity* metric - the likelihood that client i 's fingerprint as seen by server l is not unique compared to that of a different client j of server l . Finally, α_i is the product of 1) β_i and 2) $(1 - \gamma_{ij})$ over all $j \neq i$, which compares client i 's fingerprint with all other clients in its neighborhood and approaches 0 if client i 's profile is not unique. Therefore if either the honesty term or similarity term goes to 0, the weight α_i for client i also approaches zero.

$$\alpha_i = \beta_i \prod_{j \neq i} (1 - \gamma_{ij}) \quad \text{where,} \quad \beta_i = \prod_{l \in \mathcal{N}_i} \mathcal{L}(i \text{ is at } (\phi_{il}, \theta_{il}) | F_{il})$$

$$\gamma_{ij} = \prod_{l \in \mathcal{N}_i} \mathcal{L}(i \text{ spoofs } j | F_{il}, F_{jl}) \quad (2)$$

Here, $\mathcal{L}(\cdot)$ denotes an event likelihood, (ϕ_{il}, θ_{il}) is the reported direction of client i with respect to server l , and the neighborhood \mathcal{N}_i are servers communicating with client i .

Defining Honesty and Similarity Metrics: The honesty metric β_i and similarity metric γ_{ij} are derived using peak

Symbol	Meaning
m, c, s	No. of servers, clients, spoofer
p_i, x_l	Position of client i / server l
F_{il}, k	Fingerprint of i at l , k peaks
$\hat{\mathbf{h}}_{il}$	$M \times 1$ channel ratios of i to l
$f(\cdot; \mu, \sigma^2)$	PDF of normal distribution
$g(\cdot; \mu, \sigma^2)$	$\min(1, \sqrt{2\pi} f(x; \mu, \sigma^2))$
κ	Constant = $((\sqrt{2} + \sqrt{\pi})/\pi)^2$
α_i, β_i	confidence, honesty metric of i
γ_{ij}	Similarity metric of client i, j
SNR	Signal-to-noise ratio
RSSI	Received Signal Strength
$\sigma_\theta^2, \sigma_\phi^2$	Variance in peak shifts of F_{il}
$\hat{\sigma}_\theta^2, \hat{\sigma}_\phi^2$	$\sigma_\theta^2, \sigma_\phi^2$ plus measurement error
C_{V_L}, C_{V_α}	Coverage centroid of optimal, our system; error \vec{e} within ϵ
$L(Q), \rho(q)$	Footprint, Mass function

Fig. 5: **Table of Most Common Notations**

locations in client fingerprints. In practice however, peaks may have slight shifts owing to noise. Thus, any comparison between peak locations must permit some variance due to these shifts. Fortunately, noise in wireless environments can be modeled closely as additive white-Gaussian [35]. As the following lemma shows, this results in peak shifts that are also Gaussian, meaning that their variance is easy to model and account for. More formally, the lemma states that shifts are normally distributed with zero mean and well-defined variance, based on the wireless medium's signal-to-noise ratio (SNR):

Lemma 5.1: Let $\Delta\theta_i, \Delta\phi_i$ denote the error between the azimuthal and polar angle of the uncorrelated i^{th} path of a (potentially multipath) source and the corresponding angles of the (local) maximum in the fingerprint $F(\phi, \theta)$, over several uniformly gathered packets (i.e., SAR snapshots) for $\theta \in (10^\circ, 80^\circ)$. Then $\Delta\theta_i$ and $\Delta\phi_i$ are normally distributed with a mean 0, and expected variance σ_θ^2 and σ_ϕ^2 :

$$\sigma_\theta^2 = \sigma_\phi^2 = 9\lambda^2 / (8M\pi^2 r^2 SNR)$$

Where, λ is the wavelength of the signal, SNR is the signal-to-noise ratio in the network⁴, M is the number of packets per-rotation, and r is the distance between the antennas. \square

The above lemma follows from well-known Cramer-Rao bounds [25, 10, 9] shown previously for linear antenna movements in SAR [34] but readily extensible to circular rotations (proof in supplementary text [11]). Using this lemma, we can define the honesty metric β_i as the likelihood that the client is at its reported location, subject to this Gaussian error and additional measurement error in reported locations.

Definition 5.2: (β_i) Let $\phi_{F_{il}}$ and $\theta_{F_{il}}$ denote the closest maximum in $F_{il}(\phi, \theta)$ to (ϕ_{il}, θ_{il}) . We denote $\hat{\sigma}_\phi^2$ and $\hat{\sigma}_\theta^2$ as the variances in angles σ_ϕ^2 and σ_θ^2 plus any variance due to measurement error of reported locations that can be calibrated from device hardware. We define β_i for client i as:

$$\beta_i = \prod_l g(\phi_{il} - \phi_{F_{il}}; 0, \hat{\sigma}_\phi^2) \times g(\theta_{il} - \theta_{F_{il}}; 0, \hat{\sigma}_\theta^2) \quad (3)$$

Where $g(x; \mu, \sigma^2) = \min(1, \sqrt{2\pi} f(x; \mu, \sigma^2))$ is a normalized Gaussian PDF $f(x; \mu, \sigma^2)$ with mean μ and variance σ^2 . \square

⁴For clarity, we drop dependence on i, l for SNR, σ_θ and σ_ϕ

In practice, reported client locations are subject to measurement errors due to position sensor inaccuracies. Our definition of β_i above accounts for this by using the effective variances $\hat{\sigma}_\phi^2$ and $\hat{\sigma}_\theta^2$ that are the sum of the variance in angles, σ_ϕ^2 and σ_θ^2 , in addition to the variances due to measurement error.

Using Lemma 5.1 we define the similarity metric γ_{ij} as the likelihood that two client fingerprints share identical peaks:

Definition 5.3: (γ_{ij}) Let (Φ_{il}, Θ_{il}) and (Φ_{jl}, Θ_{jl}) denote the set of local maxima, ordered by non-decreasing angle values, in fingerprints F_{il} and F_{jl} . We define γ_{ij} for client i relative to client j as:

$$\gamma_{ij} = \prod_{\phi_i \in \Phi_{il}, \phi_j \in \Phi_{jl}} g(\phi_i - \phi_j; 0, 2\sigma_\phi^2) \prod_{\theta_i \in \Theta_{il}, \theta_j \in \Theta_{jl}} g(\theta_i - \theta_j; 0, 2\sigma_\theta^2) \quad (4)$$

Where $g(\cdot; \mu, \sigma^2)$ is from Definition. 5.2, and the factor of 2 in the variance accounts for computing the difference of two normally distributed values. \square

Defining the Confidence Weight: We notice that Eqn. 2, 3 and 4 fully define α_i for each client i . In summary, the confidence weight is computed in three steps: (1) Obtain the client fingerprint using SAR on wireless signal snapshots. (2) Measure the variance of peak locations of these client fingerprints using their Signal-to-Noise Ratio. (3) Compute the similarity and honesty metrics using their above definitions to obtain the confidence weight. Algorithm 1 below summarizes the steps to construct α_i for a given client i .

Algorithm 1 Algorithm to Compute Client Confidence Weight

▷ Input: Ratio of Channels \hat{h}_{H1} and SNR
 ▷ Output: Confidence Weight, α_i for client i
 ▷ Step (1): Measure fingerprints for client i
for $l = 1, \dots, m$ **do**
 for $\phi \in \{0^\circ, \dots, 360^\circ\}; \theta \in \{0^\circ, \dots, 360^\circ\}$ **do**
 Find $F_{il}(\phi, \theta)$ using a single spin to get \hat{h}_{H1} (Eqn. 1)
 end for
end for
 ▷ Step (2): Measure variances in peak locations using SNR
 $\sigma_\theta^2 = \sigma_\phi^2 =$ Apply Lemma 5.1 SNR
 ▷ Step (3): Find honesty, similarity and confidence weight
 $\beta_i =$ Apply Defn. 5.2 using $\sigma_\theta^2, \sigma_\phi^2$, peaks of F_{il}
for $j = \{1, \dots, c\} \setminus \{i\}$ **do**
 $\gamma_{ij} =$ Apply Defn. 5.3 using $\sigma_\theta^2, \sigma_\phi^2$, peaks of F_{il}, F_{jl}
end for
 $\alpha_i = \beta_i \prod_{j \neq i} (1 - \gamma_{ij})$

We now present our main result that solves Problem 1 in the problem statement (Sec. §III). The following theorem says the expected α_i 's of legitimate nodes approach 1, while those of spoofers approach 0, allowing us to discern them under well-defined assumptions: (A.1) The signal paths are independent. (A.2) Errors in azimuth and polar angles are independent. (A.3) The clients transmit enough packets to emulate a large antenna array (in practice, 25 – 30 packets per second).⁵

Theorem 5.4: Consider a network with m servers and c clients. A new client i either: 1) spoofs s clients reporting

⁵This is a mild requirement since 25 – 30 packets can be transmitted in tens of milliseconds, even at the lowest data rate of 6Mb/s of 802.11n Wi-Fi.

a random location, potentially scaling power, or; 2) is a uniformly randomly located legitimate client. Let $\alpha_{spoofer}, \alpha_{legit}$ be the confidence weights in either case. Assume that the client obtains its signals from servers along k paths (where the number of paths k is defined by Eqn. §1 in Sec. §IV). Under A.1-A.3, the expected $\alpha_{spoofer}, \alpha_{legit}$ are bounded by:

$$\begin{aligned} E[\alpha_{spoofer}] &\leq \left[\sqrt{\hat{\sigma}_\theta \hat{\sigma}_\phi \kappa} \right]^m [2mk\sigma_\theta\sigma_\phi]^s \\ E[\alpha_{legit}] &\geq 1 - cm\hat{\sigma}_\theta\hat{\sigma}_\phi \left[\sqrt{2\sigma_\theta\sigma_\phi\kappa} \right]^{mk} \end{aligned} \quad (5)$$

Where $\kappa = ((\sqrt{2} + \sqrt{\pi})/\pi)^2$, $\sigma_\theta, \sigma_\phi, \hat{\sigma}_\theta, \hat{\sigma}_\phi$ are the variances defined in Lemma 5.1 that depend on signal-to-noise ratio (the latter include measurement error in reported locations).

Proof Sketch: To give some intuition on why the theorem holds, we provide a brief proof sketch (proof in supplementary text [11]). To begin with, notice from their definitions that both the honesty metric β_i and confidence metric γ_{ij} inspect peaks in fingerprints F_{il} (Lemma 5.1). For the honesty metric β_i of a legitimate node, this peak location should be normally distributed (subject to noise, measurement error) around the reported location. For a spoofer that reports a random location, the peak location is uniformly distributed. A similar (but inverse) argument holds for γ_{ij} . Hence, we simply need to show is that the definitions of β_i and γ_i which are both products of the form $g(X)$ can be bounded in expectation if X is uniform or normally distributed.

To this end, consider two random variables u and ν which are respectively uniform and normally distributed between 0 and 2π with mean 0 and variance σ^2 . Let $S = \sqrt{2}\sigma(\ln \frac{1}{\sigma})^{0.5}$, the value at which the minimization in $g(x)$ is triggered. $E[g(\nu)]$ and $E[g(u)]$ are as follows:

$$\begin{aligned} E[g(\nu)] &= \int_{-S}^S f(x; 0, \sigma^2) dx + \sqrt{8\pi} \int_{-\infty}^{-S} [f(x; 0, \sigma^2)]^2 dx \\ &\geq \int_{-S}^S f(x; 0, \sigma^2) dx = \text{erf}\left(\frac{S}{\sigma\sqrt{2}}\right) \geq 1 - \sigma \end{aligned} \quad (6)$$

Where $\text{erf}(\cdot)$ is the well known Error function and using $1 - \text{erf}(x) < e^{-x^2}$. Similarly, we can evaluate $E[u(n)]$ as:

$$\begin{aligned} E[g(u)] &= \int_{-S}^S \frac{1}{2\pi} dx + 2\sqrt{2\pi} \int_{-2\pi}^{-S} \frac{1}{2\pi} f(x; 0, \sigma^2) dx \\ &\leq \frac{S}{\pi} + \frac{1}{\sqrt{2\pi}} \left(1 - \text{erf}\left(\frac{S}{\sigma\sqrt{2}}\right) \right) \leq \sqrt{\sigma}\kappa \end{aligned} \quad (7)$$

By assumptions A.1-A.3, we can apply these bounds to write the expectation of the honesty metric β_i as a product of those of the independent variables:

$$\begin{aligned} E[\beta_{spoofer}] &= \prod_l E[g(u; 0, \hat{\sigma}_\phi^2)] E[g(u; 0, \hat{\sigma}_\theta^2)] \leq \left[\sqrt{\hat{\sigma}_\theta \hat{\sigma}_\phi \kappa} \right]^m \\ E[\beta_{legit}] &= \prod_l E[g(\nu; 0, \hat{\sigma}_\phi^2)] E[g(\nu; 0, \hat{\sigma}_\theta^2)] \geq 1 - m\hat{\sigma}_\theta\hat{\sigma}_\phi \end{aligned}$$

Applying a similar argument, the similarity metric γ is:

$$E[\gamma_{spoof}] = \prod_{p=1}^k E[f(\nu; 0, 2\sigma_\phi^2)f(\nu; 0, 2\sigma_\theta^2)] \geq 1 - 2mk\sigma_\theta\sigma_\phi$$

$$E[\gamma_{legit}] = \prod_{p=1}^k E[g(u; 0, 2\sigma_\phi^2)g(u; 0, 2\sigma_\theta^2)] \leq [\sqrt{2\sigma_\theta\sigma_\phi\kappa}]^{mk}$$

Combining the above equations, we prove Eqn. 5. \square

A natural question one might ask is if the above lemma holds in general environments, where its assumptions A.1-A.3 may be too stringent. Our extensive experimental results in Sec. VII show that our bounds on α approximately predict performance in general environments. Further, Sec. §VII-A shows that results from an anechoic chamber, which emulate free-space conditions where the lemma's assumptions can be directly enforced, tightly follow the bounds of Lemma 5.1.

In sum, one can adopt the above lemma to distinguish adversarial nodes from legitimate nodes, purely based on α . However, an interesting alternative is to incorporate α directly into multi-robot controllers to give provable service guarantees to legitimate nodes. The next section show how α_i readily integrates with robotic coverage controllers, in particular.

VI. THREAT-RESISTANT DISTRIBUTED CONTROL

This section describes how our spoof detection method from Sec. §V integrates with well-known coverage controllers from [5, 31, 32]. The area coverage problem deals with positioning server robots to minimize their Euclidean distance to certain areas of interest in the environment. These areas are determined by an importance function $\rho(q)$ that is defined over the environment $\mathcal{Q} \subset \mathbb{R}^3$ of size $L(\mathcal{Q})$. For our coverage problem, the peaks of the importance are determined by client positions P , e.g., $\rho(q, P) = \rho_1(q) + \dots + \rho_c(q)$ where $\rho_i(q)$ quantifies the influence of client i 's position on the importance function. Using [5, 31, 32], server robot positions optimizing coverage over $\rho(q, P)$ will minimize their distance to clients.

To account for spoofed clients, we modify the importance function $\rho(q, P)$ using the α_i for each client $i \in [c]$ that is computed by Algorithm 1. E.g., we can multiply each client-term in $\rho(q, P)$ by its corresponding confidence weight: $\rho(q, P)_\alpha = \alpha_1\rho_1(q) + \dots + \alpha_c\rho_c(q)$. Given the properties of these weights derived in Theorem 5.4, i.e., α_i is bounded near zero for a spoofed client and near one for a legitimate client, the effect of multiplication by the α 's is that terms corresponding to spoofed clients will be bounded to a small value (see Fig. 6); providing resilience to the spoofing attack.

For simplicity, we assume the importance function $\rho(q)$ is static (from [5]) and α 's from Algorithm 1 are computed once, at the beginning of the coverage algorithm. We note that our approach readily extends to the adaptive case in [31, 32] when the importance function (and location of clients) change, by having the service robots exchange their learned importance function. This in turn can trigger a re-calculation of α values.

We now show that computed server positions are impacted by spoofers to within a closed-form bound, that depends on

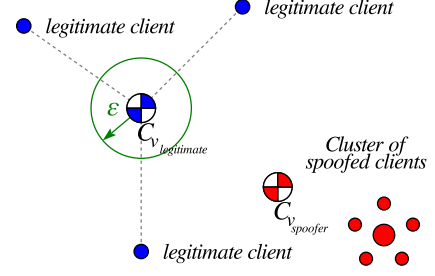


Fig. 6: **Coverage guarantee:** An ϵ ball around the ground-truth centroid, $C_{V_{\text{legitimate}}}$, is shown in green. Theorem 6.1 finds $\epsilon(\mathcal{P})$ so that server positions remain in this ball in the presence of spoofed clients.

problem parameters like signal-to-noise ratio. Theorem 6.1 below solves Problem 2 of our problem statement (Sec. §III).

Theorem 6.1: Let X be a set of server robot positions and $P = S \cup \tilde{S}$ be a set of client positions where S is the set of spoofed client positions, and \tilde{S} is the set of legitimate clients. The identities of the clients being spoofed is assumed unknown. Let $\{\alpha_1, \dots, \alpha_c\}$ be a set of confidence weights satisfying Theorem 5.4 and assume a known importance function $\rho(q, P) = \rho_1(q) + \dots + \rho_c(q)$ that is defined over the environment $\mathcal{Q} \subset \mathbb{R}^3$ of size $L(\mathcal{Q})$. Define $C_V = \{x_1^*, \dots, x_m^*\}$ to be the set of server positions optimized over $\rho(q, \tilde{S})$, i.e., where there are zero spoofed clients and C_{V_α} to be the set of server positions optimized over $\rho(q, P)_\alpha = \alpha_1\rho_1(q) + \dots + \alpha_c\rho_c(q)$ where there is at least one spoofed client, i.e. $|S| \geq 1$. If $\{\alpha_1, \dots, \alpha_c\}$ satisfy Theorem 5.4, we have that $\forall x \in C_{V_\alpha}$ there exists a unique $y \in C_V$, where in the expected case $\text{dist}(x, y) \leq \epsilon(m, s, \sigma_\phi, \sigma_\theta, \kappa)$

$$\epsilon = \max \left\{ [\sqrt{\hat{\sigma}_\theta \hat{\sigma}_\phi \kappa}]^m [2mk\sigma_\theta\sigma_\phi]^s, cm\hat{\sigma}_\theta\hat{\sigma}_\phi [\sqrt{2\sigma_\theta\sigma_\phi\kappa}]^{mk} \right\} L(\mathcal{Q})$$

and $m, s, \sigma_\phi, \sigma_\theta, \kappa$ are problem parameters as in Theorem 5.4.

Proof: We make an important observation that $E[\alpha_i] \leq a$ if client i is a spoofed node, and $E[\alpha_i] \geq b$ otherwise; hence:

$$\rho(q, P)_\alpha = a(\rho_1(q) + \dots + \rho_s(q)) + b(\rho_{s+1}(q) + \dots + \rho_c(q))$$

is the maximal effect that the presence of spoofed clients can have on the importance function. Intuitively, all spoofed clients have a weight of *at maximum* a and all legitimate clients have a reduced weight of *at minimum* b . Using this observation we can bound the influence of the spoofed clients on computed server control inputs (see Fig. 6). Specifically, recall from [5] that the position control for each server is: $u_l = -2M_V(C_V - c_l)$, where $M_V = \int_V \rho(q) dq$, $C_V = \frac{1}{M_V} \int_V q \rho(q) dq$ and V is the voronoi partition for server l defined as all points $q \in \mathcal{Q}$ with $\text{dist}(q, x_l) < \text{dist}(q, x_g)$ where $g \neq l$. Using the importance function from above we can write $C_{V_\alpha} = \frac{1}{M_{V_\alpha}} (aC_{V_S} + bC_{V_L})$ where C_{V_S} is the component of the centroid computed over spoofed nodes and C_{V_L} is the component of the centroid computed over legitimate nodes and M_{V_α} is defined shortly. We rewrite C_{V_S} as a perturbation of the centroid over legitimate nodes as $C_{V_S} = C_{V_L} + \vec{v}\|\vec{e}\|$ where \vec{v} is an arbitrary unit vector and the magnitude of \vec{e} can be as large as the length of the operative environment, $\|\vec{e}\| \leq L(\mathcal{Q})$. Let the total mass be $T = M_{V_S} + M_{V_L}$. We can write a similar expression for the mass M_{V_α} using the bounds a and b as

$M_{V_\alpha} = bT + (a-b)M_{V_L}$. Substituting these expressions into C_{V_α} and simplifying gives $C_{V_\alpha} = \frac{C_{V_L} + b\vec{e}\|\vec{e}\|}{bT + (a-b)M_{V_L}}$. Combining this expression with the server control input:

$$u_l = k \left([(a+b)C_{V_L} - p_l] + b\|\vec{e}\|\vec{v} \right) \quad (8)$$

Where $k = -2(bT + aM_{V_L})$. If $(a+b) = 1$, this control input drives the server robot l to a neighborhood of size $\epsilon = b\|\vec{e}\| \leq bL(\mathcal{Q})$ centered around the centroid C_L defined over the legitimate clients. So if $b = \max \left\{ [\sqrt{\hat{\sigma}_\theta \hat{\sigma}_\phi \kappa}]^m [2mk\sigma_\theta\sigma_\phi]^s, cm\hat{\sigma}_\theta\hat{\sigma}_\phi [\sqrt{2\sigma_\theta\sigma_\phi\kappa}]^{mk} \right\}$ from Theorem 5.4 Equation (5), then:

$$\epsilon = \max \left\{ [\sqrt{\hat{\sigma}_\theta \hat{\sigma}_\phi \kappa}]^m [2mk\sigma_\theta\sigma_\phi]^s, cm\hat{\sigma}_\theta\hat{\sigma}_\phi [\sqrt{2\sigma_\theta\sigma_\phi\kappa}]^{mk} \right\} L(\mathcal{Q})$$

then we have $(a+b) = 1$ as desired, proving the lemma. \square

VII. EXPERIMENTAL RESULTS

This section describes our results from an experimental evaluation of our theoretical claims. Our aerial servers were implemented on two AscTec Atomboard computing platforms equipped with Intel 5300 Wi-Fi cards with two antennas each, mounted on two AscTec Hummingbird quadrotors. Our clients were ten iRobot Create robots, each equipped with Asus EEPc netbooks and single-antenna Wi-Fi

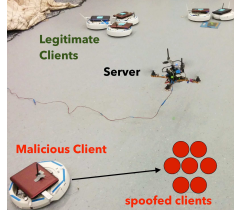


Fig. 10: Test system.

cards. An adversarial client forged multiple identities by spawning multiple packets containing different identities (up to 75% of the total number of legitimate clients in the system), and could use a different transmit power for each identity. The adversary advertised identities by modifying the Wi-Fi MAC field, a common technique for faking multiple identities [33]. **Evaluation:** We evaluate our system in two environments: 1) An indoor multipath-rich environment with walls and obstacles equipped with a Vicon motion capture system to aid quadrotor navigation; 2) An anechoic chamber to emulate a free-space setting that is particularly challenging to our system. We estimated the average theoretical expected standard deviation to be $\sigma_\theta, \sigma_\phi$ of 0.7° (Lemma 5.1). After including the standard deviation in reported location, based on the known errors of our localization framework, this increased the average $\hat{\sigma}_\theta, \hat{\sigma}_\phi$ by 2° (variances in each experiment depend on measured SNR) We compare our system against a baseline that uses a Received Signal Strength (RSSI) comparison (akin to [29]).

Roadmap: We conduct three classes of experiments: (1) Microbenchmarks to validate our client confidence metric, both in free-space and multipath indoor environments (Sec. §VII-A). (2) Experiments applying this confidence metric to quarantine adversaries (Sec. §VII-B). (3) Application of our system to secure the coverage problem against Sybil attacks (Sec. §VII-C).

A. Microbenchmarks on the Confidence Metric

This experiment studies the correctness of our system's confidence metric α . Recall from theory in §V that α 's measured by a server robot distinguish between unique clients based on their diverse physical directions and the presence

of multipath reflections. Thus, a free-space environment (i.e., with no multipath) is particularly challenging to our system.

Method: To approximate free-space, we measured α values in a radio-frequency anechoic chamber which attenuates reflected paths by about 60 dB, for a legitimate and malicious client from one server robot 12 m away. Next, in a 10 m x 8 m indoor room (a typical multipath case), we measured α 's from one server for up to ten legitimate clients and ten spoofed clients.

Results: In Fig. 7, the values of α in the anechoic chamber tightly follow our theoretical bounds in Theorem 5.4 (Fig. 8(c)). As expected, our results in indoor multipath environments exhibit a larger variance but follow the trend suggested by theory. Further, we stress our confidence metric by isolating the case of colinearity in both environments. In Fig. 8, we consider a spoofing adversary initially co-aligned with a legitimate client, and measure α as the angle of separation, ϕ , is increased from 0° to 20° relative to the server robot. In the anechoic chamber at ϕ close to 0° , the fingerprints of both the legitimate and adversarial nodes are virtually identical, each with precisely one peak at 0° . Consequently, α for the legitimate node is much below 1, indicating that is believed to be adversarial (i.e., the term $1-\gamma$ in α approaches 0 in Eqn. 2). However, α for the legitimate client quickly approaches 1, even if $\phi = 3^\circ$ in the anechoic chamber. In fact, α is virtually identical to 1 beyond 10° , indicating that a single server robot can distinguish closely aligned legitimate and adversarial clients even in free-space. Fig. 8b shows that multipath can distinguish clients even at $\phi = 0^\circ$, due to additional reflected paths that help disambiguate these clients.

B. Performance of Sybil Attack Detection

In this experiment, we measure our system's classification performance on legitimate and spoofed clients, in the presence of static, mobile, and power-scaling adversaries.

Method: This experiment was performed in the multipath-rich indoor testbed with walls and obstacles. Each run consisted of one quadrotor server, and (randomly positioned) ten control clients, or nine legitimate clients with an adversary reporting two to nine spoofed clients. Each Sybil attack was performed under three modalities: (1) a stationary attacker with a fixed transmission power, (2) a mobile attacker (random-walk and linear movements), and (3) an attacker scaling the per-packet power by a different amount for each spoofed client, from 1 to 31 mW. The quadrotor server classifies clients with an $\alpha < 0.5$ as spoofed (see Fig. 7). The baseline RSSI classifier uses a 2 dB thresholded minimum dissimilarity, a technique previously applied in static networks [29, 37].

Results: For each modality, our performance against an RSSI baseline over multiple network topologies is summarized here as true positive rates (TPR) and false positive rates (FPR). In particular, our classifier is robust to power-scaling Sybil attacks (where RSSI performs poorly) since we use the ratio of

	Our System		RSSI	
	TPR	FPR	TPR	FPR
Static	96.3	3.0	81.5	9.1
Mobile	96.3	6.1	85.2	6.1
Δ mW	100.0	3.0	74.1	27.3

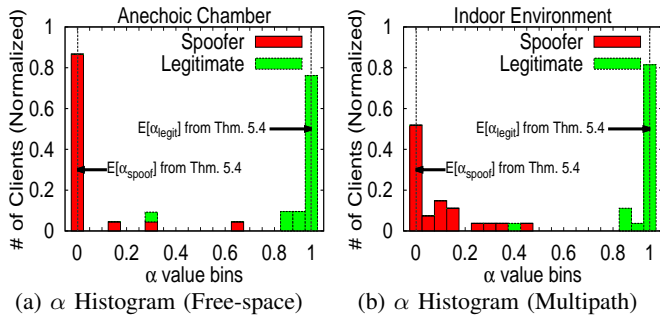


Fig. 7: **Experimental Evaluation of α** : (a) In an anechoic chamber approximating our assumptions A.1-A.3 (§5.4), α largely agrees with theory. (b) In a typical multipath environment, experimental results closely follow theoretical predictions. Data shows that $\alpha = 0.5$ is a good threshold value.

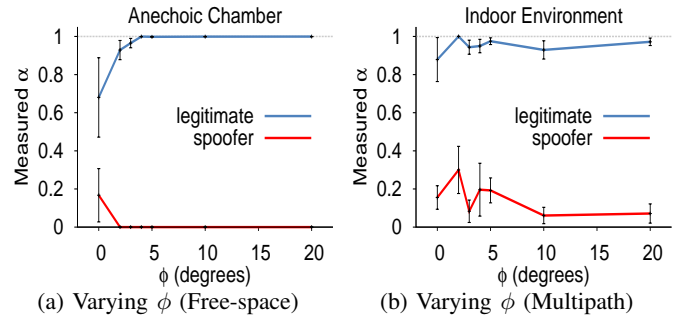


Fig. 8: **Co-Aligned Clients**: We vary the angle ϕ between a legitimate and malicious client, relative to a single server, and plot α in (a) an anechoic chamber and (b) an indoor environment. The minimum ϕ needed to distinguish the clients is only: (a) 3° in freespace, (b) 0° in multipath settings.

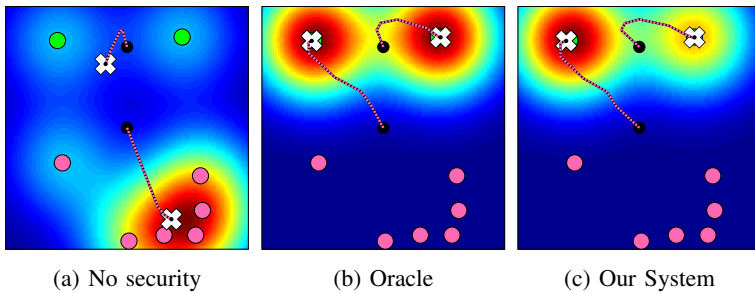
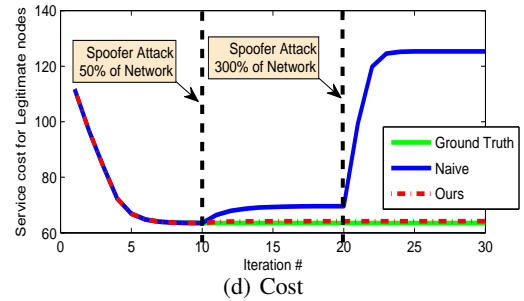


Fig. 9: **Experimental Results for Sybil Attack in Multi-Agent Coverage**: Depicts the total distance of converged quadrotor server positions (white \times) to two legitimate clients (\bullet) and six spoofed clients (\bullet). We consider: (a) an insecure system where each spoofed client creates a false peak in the importance function, (b) a ground truth importance function, and (c) our system where applying α weights from Algorithm 1 recovers the true importance function. (d) Depicts a ground-truth cost computed with respect to legitimate clients as Sybil nodes dynamically enter the network. Our system (red dotted line) performs near-optimal even when spoofed clients comprise more than twice the network.



wireless channels in computing α (Sec. §IV). Our client classifier exhibits consistent performance in both power-scaling and mobile scenarios with a TPR $\approx 96\%$ and FPR $\approx 4\%$.

C. Application to Multi-Agent Coverage

We implement the multi-agent coverage problem from [5], where a team of aerial servers position themselves to minimize their distance to client robots at reported positions $p_i, i \in [c]$. We use an importance function $\rho(q, P) = \rho_1(q) + \dots + \rho_c(q)$ defined in Sec. §VI where each client term is a Gaussian-shaped function $\rho_i(q) = \exp(-\frac{1}{2}(q - p_i)^T (q - p_i))$ (Fig. 9b). An α -modified importance function is implemented as $\rho(q, P)_\alpha = \alpha_1 \rho_1(q) + \dots + \alpha_c \rho_c(q)$ where the α terms are computed using Algorithm 1 (Fig. 9c).

Method. This experiment was performed in the multipath-rich indoor testbed. For each experiment we randomly place three clients in an 8 m x 10 m room with two AscTec quadrotor servers. Fig. 9(a)-(c) shows one client-server topology where an adversary spoofs six Sybil clients. Upon convergence, we measure the distance of each server from an optimal location in 3 scenarios: 1) a naive system with no security, 2) an oracle which discards Sybil clients *a priori*, and 3) our system.

Results: Fig. 9(a)-(c) depicts the converged locations for a candidate topology in the above three scenarios. We observe that by incorporating α weights in our controller, our system approximates oracle performance. Fig. 9d demonstrates the

ability of our system to bound the service cost to near optimal even as spoofers enter the network (comprising up to 300%). **Aggregate Results:** Across multiple topologies and 12 runs, with no security the maximum distance from each quadrotor to an oracle solution is on average 3.77 m (stdev: 0.86). Our system achieves a 0.02 m (stdev: 0.02) average from oracle.

VIII. CONCLUSION

In this paper, we develop a new system to guard against the Sybil attack in multi-robot networks. We derive theoretical guarantees on the performance of our system, which are validated experimentally. While this paper has focused on coverage, it can be readily extended to secure other multi-robot controllers against Sybil attacks, e.g., unmanned delivery [20], search-and-rescue [21], and formation control [38]. We note for future work that our method of detecting spoofed clients is applicable to servers as well, since they also communicate wirelessly. Since our approach is based on the fundamental physics of wireless signals, we believe that it will easily generalize beyond Sybil attacks to other Wi-Fi based security issues in robot-swarms such as packet path validation [23] and detecting packet injection attacks to name a few.

Acknowledgement: This work was partially supported by the NSF and MAST project (ARL grant W911NF-08-2-0004). We thank members of the MIT Center for Wireless Networks and Mobile Computing: Amazon.com, Cisco, Google, Intel, MediaTek, Microsoft, and Telefonica for their interest and general support.

REFERENCES

- [1] Amazon prime air. URL <http://www.amazon.com/b?node=8037720011>.
- [2] Fadel Adib, Swarun Kumar, Omid Aryan, Shyamnath Gollakota, and Dina Katabi. Interference Alignment by Motion. MOBICOM, 2013.
- [3] R.W. Beard, T.W. McLain, D.B. Nelson, D. Kingston, and D. Johanson. Decentralized cooperative aerial surveillance using fixed-wing miniature uavs. *Proceedings of the IEEE*, 94(7):1306–1324, July 2006. ISSN 0018-9219. doi: 10.1109/JPROC.2006.876930.
- [4] Airlie Chapman, Marzieh Nabi-Abdolyousefi, and Mehran Mesbahi. Identification and infiltration in consensus-type networks. *1st IFAC Workshop on Estimation and Control of Networked Systems*, 2009.
- [5] J. Cortes, S. Martinez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *Robotics and Automation, IEEE Transactions on*, 20(2):243–255, April 2004. doi: 10.1109/TRA.2004.824698.
- [6] K. Daniel, B. Dusza, A. Lewandowski, and C. Wietfeld. Airshield: A system-of-systems muav remote sensing architecture for disaster response. In *Systems Conference, 2009 3rd Annual IEEE*, pages 196–200, March 2009. doi: 10.1109/SYSTEMS.2009.4815797.
- [7] JohnR. Douceur. The sybil attack. In Peter Druschel, Frans Kaashoek, and Antony Rowstron, editors, *Peer-to-Peer Systems*, volume 2429 of *Lecture Notes in Computer Science*, pages 251–260. Springer Berlin Heidelberg, 2002. ISBN 978-3-540-44179-3. doi: 10.1007/3-540-45748-8_24. URL http://dx.doi.org/10.1007/3-540-45748-8_24.
- [8] Patrick J. Fitch. *Synthetic Aperture Radar*. Springer, 1988.
- [9] H. Gazzah and S. Marcos. Directive antenna arrays for 3d source localization. In *Signal Processing Advances in Wireless Communications, 2003. SPAWC 2003. 4th IEEE Workshop on*, pages 619–623, June 2003. doi: 10.1109/SPAWC.2003.1319035.
- [10] Houcem Gazzah and Sylvie Marcos. Cramer-Rao bounds for antenna array design. *IEEE Transactions on Signal Processing*, 54:336–345, 2006. doi: 10.1109/TSP.2005.861091.
- [11] Stephanie Gil, Swarun Kumar, Mark Mazumder, Dina Katabi, and Daniela Rus. Guaranteeing spoof-resilient multi-robot networks. *Full paper version with supplementary material available as a TECH REPORT at MIT CSAIL Publications and Digital Archive (<http://publications.csail.mit.edu>)*.
- [12] Stephanie Gil, Swarun Kumar, Dina Katabi, and Daniela Rus. Adaptive Communication in Multi-Robot Systems Using Directionality of Signal Strength. ISRR, 2013.
- [13] Andrea Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.
- [14] Monson H. Hayes. *Statistical Digital Signal Processing and Modeling*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1996. ISBN 0471594318.
- [15] Fiona Higgins, Allan Tomlinson, and Keith M. Martin. Threats to the swarm: Security considerations for swarm robotics. *International Journal on Advances in Security*, 2, 2009.
- [16] Dongxu Jin and JooSeok Song. A traffic flow theory aided physical measurement-based sybil nodes detection mechanism in vehicular ad-hoc networks. In *Computer and Information Science (ICIS), 2014 IEEE/ACIS 13th International Conference on*, pages 281–286, June 2014. doi: 10.1109/ICIS.2014.6912147. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6912147&tag=1.
- [17] Helmut Klausning. Feasibility of a sar with rotating antennas (rosar). In *Microwave Conference, 1989*, 1989.
- [18] Swarun Kumar, Stephanie Gil, Dina Katabi, and Daniela Rus. Accurate indoor localization with zero start-up cost. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking, MobiCom '14*, pages 483–494, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2783-1. doi: 10.1145/2639108.2639142. URL <http://doi.acm.org/10.1145/2639108.2639142>.
- [19] Swarun Kumar, Ezzeldin Hamed, Dina Katabi, and Li Erran Li. Lte radio analytics made easy and accessible. In *Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14*, pages 211–222, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2836-4. doi: 10.1145/2619239.2626320. URL <http://doi.acm.org/10.1145/2619239.2626320>.
- [20] Aleksandr Kushleyev, Brian MacAllister, and M. Likhachev. Planning for landing site selection in the aerial supply delivery. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 1146–1153, Sept 2011. doi: 10.1109/IROS.2011.6094840.
- [21] Lanny Lin and Michael A Goodrich. Uav intelligent path planning for wilderness search and rescue. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 709–714. IEEE, 2009.
- [22] Hongbo Liu, Yan Wang, Jian Liu, Jie Yang, and Yingying Chen. Practical user authentication leveraging channel state information (csi). In *Proceedings of the 9th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '14*, pages 389–400, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2800-5. doi: 10.1145/2590296.2590321. URL <http://doi.acm.org/10.1145/2590296.2590321>.
- [23] Xin Liu, Ang Li, Xiaowei Yang, and David Wetherall. Passport: Secure and adoptable source authentication. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation, NSDI'08*, pages 365–378, Berkeley, CA, USA, 2008. USENIX Association. ISBN 111-999-5555-22-1. URL <http://dl.acm.org/citation.cfm?id=1387589.1387615>.
- [24] M. MalmirChegini and Y. Mostofi. On the spatial predictability of communication channels. *Wireless Com-*

- munications, IEEE Trans.*, 11(3), 2012.
- [25] Cherian P. Mathews and Michael D. Zoltowsk. Signal subspace techniques for source localization with circular sensor arrays. Purdue University TechReport, 1994.
- [26] J. Newsome, E. Shi, D. Song, and A. Perrig. The sybil attack in sensor networks: analysis defenses. In *Information Processing in Sensor Networks, 2004. IPSN 2004. Third International Symposium on*, pages 259–268, April 2004. doi: 10.1109/IPSN.2004.1307346.
- [27] R. Olfati-Saber and R.M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *Automatic Control, IEEE Transactions on*, 49(9): 1520–1533, Sept 2004. ISSN 0018-9286. doi: 10.1109/TAC.2004.834113.
- [28] Lynne E. Parker. Distributed algorithms for multi-robot observation of multiple moving targets. *Autonomous Robots*, 12, 2002.
- [29] Jr. Pires, W.R., T.H. de Paula Figueiredo, H.C. Wong, and A.A.F. Loureiro. Malicious node detection in wireless sensor networks. In *Parallel and Distributed Processing Symposium, 2004. Proceedings. 18th International*, pages 24–, April 2004. doi: 10.1109/IPDPS.2004.1302934.
- [30] I. Sargeant and A. Tomlinson. Modelling malicious entities in a robotic swarm. In *Digital Avionics Systems Conference (DASC), 2013 IEEE/AIAA 32nd*, Oct 2013.
- [31] M. Schwager, Brian J. Julian, and D. Rus. Optimal coverage for multiple hovering robots with downward facing cameras. In *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pages 3515–3522, May 2009. doi: 10.1109/ROBOT.2009.5152815.
- [32] Mac Schwager, Daniela Rus, and Jean-Jacques Slotine. Decentralized, adaptive coverage control for networked robots. *The International Journal of Robotics Research*, 28(3):357–375, 2009. URL <http://ijr.sagepub.com/content/28/3/357.abstract>.
- [33] Yong Sheng, K. Tan, Guanling Chen, D. Kotz, and A. Campbell. Detecting 802.11 mac layer spoofing using received signal strength. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages –, April 2008. doi: 10.1109/INFOCOM.2008.239. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4509834&tag=1.
- [34] Petre Stoica and Nehorai Arye. Music, maximum likelihood, and cramer-rao bound. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 37(5):720–741, May 1989. ISSN 0096-3518. doi: 10.1109/29.17564.
- [35] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [36] Jue Wang and Dina Katabi. Dude, where’s my card?: Rfid positioning that works with multipath and non-line of sight. SIGCOMM, 2013.
- [37] Ting Wang and Yaling Yang. Analysis on perfect location spoofing attacks using beamforming. In *INFOCOM, 2013 Proceedings IEEE*, pages 2778–2786, April 2013. doi: 10.1109/INFOCOM.2013.6567087. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6567087.
- [38] Xiaohua Wang, Vivek Yadav, and SN Balakrishnan. Cooperative uav formation flying with obstacle/collision avoidance. *Control Systems Technology, IEEE Transactions on*, 15(4):672–679, 2007.
- [39] Yong Wang, G. Attebury, and B. Ramamurthy. A survey of security issues in wireless sensor networks. *Communications Surveys Tutorials, IEEE*, 8(2):2–23, Second 2006. ISSN 1553-877X. doi: 10.1109/COMST.2006.315852.
- [40] Jie Xiong and Kyle Jamieson. Securearray: Improving wifi security with fine-grained physical-layer information. In *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking, MobiCom '13*, pages 441–452, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1999-7. doi: 10.1145/2500423.2500444. URL <http://doi.acm.org/10.1145/2500423.2500444>.
- [41] Jie Yang, Yingying Chen, W. Trappe, and J. Cheng. Detection and localization of multiple spoofing attackers in wireless networks. *Parallel and Distributed Systems, IEEE Transactions on*, 24(1):44–58, Jan 2013. ISSN 1045-9219. doi: 10.1109/TPDS.2012.104.
- [42] Zhimin Yang, E. Ekici, and Dong Xuan. A localization-based anti-sensor network system. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 2396–2400, May 2007. doi: 10.1109/INFOCOM.2007.288. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4215870.