Auton Robot (2017) 41:1383–1400
DOI 10.1007/s10514-017-9621-5

Guaranteeing spoof-resilient multi-robot networks

Stephanie Gil[1] · **Swarun Kumar**[2] · **Mark Mazumder**[3] · **Dina Katabi**[1] · **Daniela Rus**[1]

Received: 15 December 2015 / Accepted: 11 January 2017 / Published online: 28 February 2017

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 multi-robot problems, including locational coverage and unmanned delivery. We experimentally validate our claims using a team of AscTec quadrotor servers and iRobot Create ground clients, and demonstrate spoofer detection rates over 96%.

Keywords Multi-robot systems · Cybersecurity · Sybil attack · Wireless networks · Coordinated control · Anechoic chamber · Performance bounds

This is one of several papers published in *Autonomous Robots* comprising the "Special Issue on Robotics Science and Systems".

Stephanie Gil and Swarun Kumar: Co-primary authors.

✉ Stephanie Gil
sgil@mit.edu

Swarun Kumar
swarun@cmu.edu

Mark Mazumder
mazumder@ll.mit.edu

Dina Katabi
dk@mit.edu

Daniela Rus
rus@mit.edu

[1] Massachusetts Institute of Technology, Cambridge, MA, USA

[2] Carnegie Mellon University, Pittsburgh, PA, USA

[3] MIT Lincoln Laboratory, Lexington, MA, USA

1 Introduction

Multi-robot networks rely on wireless communication to enable a wide range of tasks and applications: coverage (Parker 2002; Cortes et al. 2004; Schwager et al. 2009a), disaster management (Daniel et al. 2009), surveillance (Beard et al. 2006), and consensus (Olfati-Saber and Murray 2004) 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 http://www.amazon.com/b?node=8037720011) or traffic rerouting algorithms (e.g., Google Maps Navigation) that will 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. These attacks are notoriously easy to implement (Sheng et al. 2008) 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 (Fig. 1). Although a vast body of literature is dedicated to cybersecurity in general multi-

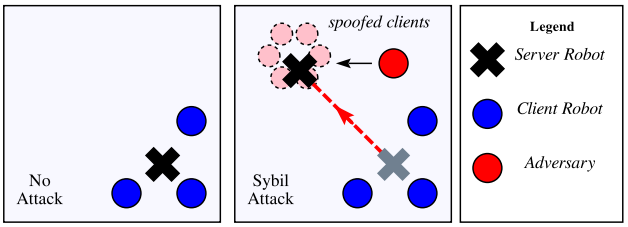

Fig. 1 Sybil attack on coverage: a server robot provides locational coverage to legitimate clients when no attack is present. In a Sybil attack, an adversary spoofs many fake clients to draw away coverage from the legitimate clients

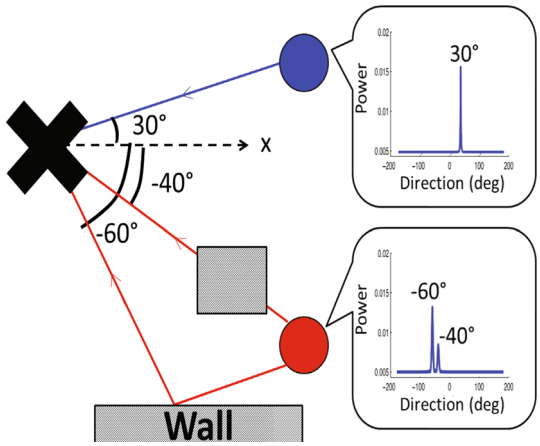

Fig. 2 Spatial fingerprints: a quadrotor server measures the directional signal strength of each client (here, simplified to 2-D). The *top-right* client has one line-of-sight peak; the other (*bottom-right* client) , 2 signal paths

node networks (e.g., a wired LAN), the same is not true for multi-robot networks (Higgins et al. 2009; Sargeant and Tomlinson 2013), 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 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 (Goldsmith 2005; MalmirChegini and Mostofi 2012).

Using these derived fingerprints, we show that a confidence metric, $\alpha \in (0, 1)$ can be obtained for each client in the network. We prove that these confidence metrics have a desirable property where legitimate clients have an expected confidence metric close to one, while spoofed clients will have an expected confidence metric close to zero. A particularly attractive feature of the confidence metric α is that it can be readily integrated into a wide variety of multi-robot controllers. In particular, we demonstrate two natural methods to integrate α into these controllers: either as a continuous per-client weighting function or as a means to classify clients discretely into two groups as either legitimate or spoofed. More importantly, we prove analytical bounds on α that provably limit the influence of adversarial clients on the performance of these controllers. We integrate our confidence metric with multi-robot controllers in the context of two well-known problems: locational coverage (Cortes et al. 2004; Schwager et al. 2009a) and unmanned delivery (http://www.amazon.com/b?node=8037720011, Laporte et al. 1988; Pavone et al. 2011).

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. Our results in both the coverage and unmanned delivery problems demonstrate a spoofer detection rate of 96%. In addition, 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.

1.1 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 two contexts: (1) locational coverage, where each robot providing coverage remains within a bounded radius of its optimal position even in the presence of adversarial clients. (2) unmanned delivery, where the total path length traversed by the service vehicle remains bounded relative to its value in the absence of

an attack. Our theoretical results are validated extensively through experiments in diverse settings.

2 Related work

2.1 Sybil attack

In a network comprised of multiple agents, a Sybil attack is one where a single agent can simultaneously forge multiple identities in order to gain a disproportionate influence in the network. The survey article (Newsome et al. 2004) describes in detail many different types of Sybil attacks and in Douceur (2002), Douceur proves several results showing that without a centralized authority, Sybil attacks are always possible for any practical distributed network. The survey paper in Levine et al. (2006) acknowledges the sparsity of solutions applicable to distributed systems that lack a centralized key distribution system. In contrast, our proposed method does not rely on centralized key distribution or verification of keys but rather verification using inter-agent communication signals themselves, already present in distributed systems.

2.2 Other approaches to security

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 Wang et al. (2006)]. These require trusted central authorities and computationally expensive distributed key management, to account for dynamic clients that enter and leave the network (Wang et al. 2006). (2) Non cryptographic techniques in the wireless networking community leverage wireless physical-layer information to detect spoofed client identities or falsified locations (Jin and Song 2014; Yang et al. 2007, 2013; Xiong and Jamieson 2013; Xiao et al. 2009). 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) (Liu et al. 2014; Wang and Yang 2013; Pires et al. 2004) and channel state information (Liu et al. 2014). Such techniques need clients to remain static, since mobility can cause wireless channels to fluctuate rapidly (Adib et al. 2013). 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 (Higgins et al. 2009; Sargeant and Tomlinson 2013; Chapman et al. 2009).

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 (Xiong and Jamieson 2013; Yang et al. 2007). (3) It is robust to client mobility and power-scaling attacks.

Finally, our system builds on Synthetic Aperture Radar (SAR) to construct signal fingerprints (Fitch 1988).

2.3 Synthetic Aperture Radar (SAR)

SAR has been widely used for radar imaging (Fitch 1988; Klausing 1989) and indoor positioning (Kumar et al. 2014a, b; Wang and Katabi 2013; Gil et al. 2013). 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 (Cortes et al. 2004; Schwager et al. 2009a) as shown in Sect. 6 and drone delivery controllers (Laporte et al. 1988; Pavone et al. 2011) as described in Sect. 7.

Our previous work in Gil et al. (2015b) provides a theoretical and experimental framework for using SAR in the context of cybersecurity for multi-agent networks, where the influence of spoofed nodes is considered to be a continuous function. As a result, the previous formulation would not be applicable to graph-based problems that require a binary classification for the spoofed nodes. This paper extends upon our previous work by (1) deriving theoretical results for multi-robot problems that require optimization over a graph, (2) giving explicit treatment to the unmanned delivery problem as an example in the graph-based problem space, and (3) presenting an experimental framework for binary classification of spoofed nodes using Wi-Fi fingerprints.

3 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 $P := \{p_1, \ldots, p_c\}$ denote the client positions in $\mathbb{R}^3$. Let $X := \{x_1, \ldots, x_m\}$ be the positions of the servers in $\mathbb{R}^3$ and the notation $[m] = \{1, \ldots, m\}$ denote their indices. We con-

sider the case where a subset of the clients, $S \subset P$ (with $s := |S|$) are "spoofed" clients.

Definition 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*.

3.1 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.[1] 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 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 the received signal from client i along that direction. Using this mapping, or "fingerprint", our first problem is to derive a *confidence metric* 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 (*Spoofer 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.[2] 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 metric $\alpha_i(\mathcal{F}_i) \in (0, 1)$ 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 as weights that can bound the influence of spoofers in multi-robot problems. Specifically, we consider the well-known coverage problem in Cortes et al. (2004), Schwager et al. (2009a). We show that by integrating the confidence metric 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 $C_V = \{x_1^*, \ldots, x_m^*\}$ 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., $\rho(q) = \rho_1(q) + \cdots + \rho_c(q)$. Let $C_V = \{x_1^*, \ldots, x_m^*\}$ 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 $C_{V_\alpha} = \{x_1, \ldots, x_m\}$ be the set of server positions returned by an optimization of an α-*modified* importance function $\rho(q) = \alpha_1\rho_1(q) + \cdots + \alpha_c\rho_c(q)$ 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.

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). Section 6 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.

Finally, Problem 3 below shows that the α weights can be used to derive discrete decision variables for selecting what

[1] The 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.

[2] Detecting 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 Sect. 5 and experimental results (Sect. 9.1) use as little as one server.

clients to service, for example, in a drone delivery context. Here, the goal is to bound the difference between the resulting expected path length and the expected path length in the optimal case of no spoofed clients. For consistency, we will refer to the delivery drone as a "server" throughout.

Problem 3 (*Sybil-resilience in Drone Delivery*) Consider the graph $G = (V, E)$ where vertices $V = P \cup x$ are client and depot positions P and x respectively, and edges $e_i \in E$ connect the vertex of every client $p_i \in P$ to the depot vertex x, i.e., a star graph where x is the inner vertex. Note that we consider the case for one server and several clients where the goal of the server is to serve each client, by iteratively picking up its package at the depot location x and transporting it to the client's location $p \in P$.

Let the path cost for each edge $d : (E) \rightarrow \mathbb{R}$ be the Euclidean distance of that edge in G. We wish to show that an indicator function I_{α_i} defined over the α_i from Problem 1 can be used as a decision variable to select a subset of clients $P^* \subset P$ to be serviced by the delivery vehicle. The resulting subset of clients P^* has the property that the expected path length computed over this subset of clients, $L = \sum_{p_i \in P^*} d(p_i, x)$, is the same *to within a computable bound*, as the expected path length computed over only legitimate clients $L_{\text{legit}} = \sum_{p_i \in P \setminus S} d(p_i, x)$. In other words, we wish to find a set of problem parameters $\mathcal{P}$ and a bound $\delta(\mathcal{P})$ such that $|E[L] - E[L_{\text{legit}}]| \leq \delta(\mathcal{P})$.

4 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)

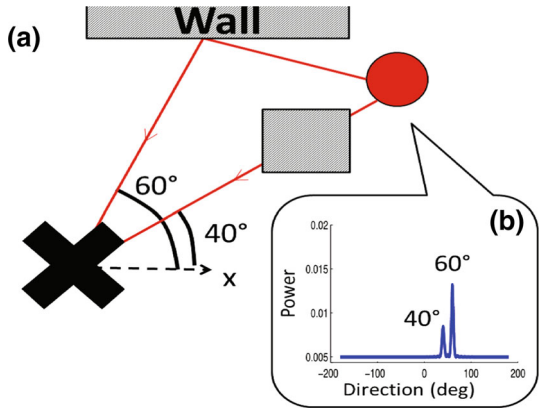

Fig. 3 Example signal fingerprint: **a** server (×) receives a client (*filled circle*) signal on 2 paths: direct along 40° attenuated by an obstacle (*shaded square*) and reflected by a wall along 60°. **b** is a corresponding fingerprint: peak heights at 40° and 60° correspond to their relative attenuations

several different paths (Tse and Vishwanath 2005). Figure 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 (Fitch 1988) (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 (Hayes 1996) 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:

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

where $\hat{\mathbf{h}}_{\mathbf{il}}$ is a vector of the ratio of wireless channel snapshots between two antennas mounted on the body of the server l and $\Psi_{\mathbf{il}}(\phi, \theta) = \frac{2\pi r}{\lambda} \cos(\phi - \mathbf{B_l}) \sin(\theta - \Gamma_\mathbf{l})$, λ is the wavelength of the signal and r is the distance between the antennas, $\mathbf{B_l}$, $\Gamma_\mathbf{l}$ are the server's angular orientation, $Eig_n(\cdot)$ are noise eigen-

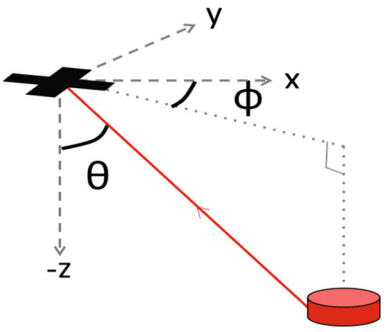

Fig. 4 3-D angles: this figure depicts the notation for the azimuthal angle ϕ and polar angle θ for the direct path from a ground client (*filled circle*) to aerial server robot ($\times$) in 3 dimensions. More generally, the set of all angles between client i and server l are denoted as Φ_{il}, Θ_{il}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

vectors, $(\cdot)^\dagger$ is conjugate transpose. We denote the number of signal eigenvectors, equal to the number of paths, by k.

While our above formulation is derived from MUSIC (Hayes 1996), it varies in one important way: while MUSIC uses a single-antenna channel snapshot h_{il}, we use the channel ratio $\hat{h}_{il} = h_{1_{il}}/h_{2_{il}}$ 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 $\hat{h}_{il} = \chi h_{1_{il}}/(\chi h_{2_{il}})$; a value unaffected by power scaling (Table 1).

5 Constructing a client confidence metric

Given a client fingerprint $F_{il}(\phi, \theta)$ for each client i relative to a robotic server l, we wish to generate a confidence metric $\alpha_i \in (0, 1)$ 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 confidence metric α_i for client i also approaches zero.

$$\begin{aligned}
\alpha_i &= \beta_i \prod_{j \neq i} (1 - \gamma_{ij}) \text{ where,} \\
\beta_i &= \prod_{l \in \mathcal{N}_i} \mathcal{L}(i \text{ is at } (\phi_{il}, \theta_{il}) | F_{il}) \\
\gamma_{ij} &= \prod_{l \in \mathcal{N}_i} \mathcal{L}(i \text{ spoofs } j | F_{il}, F_{jl})
\end{aligned} \tag{2}$$

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.

5.1 Defining honesty and similarity metrics

The honesty metric β_i and similarity metric γ_{ij} are derived using peak locations in client fingerprints. In practice however, peaks may have slight shifts owing to noise. Thus,

Table 1 Table of most common notations

Symbol	Meaning
m, c, s	No. of servers, clients, spoofers
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 = $\left(\left(\sqrt{2} + \sqrt{\pi}\right)/\pi\right)^2$
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 $\mathbf{e}$ within ϵ
$L(\mathcal{Q}), \rho(q)$	Footprint, mass function

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 (Tse and Vishwanath 2005). 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 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 $\theta \in (10°, 80°)$. Then $\Delta\theta_i$ and $\Delta\phi_i$ are normally distributed with a mean 0, and expected variance σ_ϕ^2 and σ_θ^2:*

$$\sigma_\theta^2 = \sigma_\phi^2 = 9\lambda^2/(8M\pi^2 r^2 SNR) \tag{3}$$

where, λ is the wavelength of the signal, SNR is the signal-to-noise ratio in the network,[3] M is the number of packets per-rotation, and r is the distance between the antennas. □

The above lemma follows from well-known Cramer-Rao bounds (Mathews and Zoltowsk 1994; Gazzah and Marcos 2006, 2003) shown previously for linear antenna movements in SAR (Stoica and Arye 1989) but readily extensible to circular rotations (proof in supplementary text Gil et al. 2015a). Note that from Eq. (3) the relationship between the antenna distance r and the resolution of the resulting fingerprint σ becomes apparent. The larger the distance between two mounted antennas used in the channel ratio (Eq. 1), the smaller the error variance and thus the finer the resolution of the fingerprint.

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 3 (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:

$$\beta_i = \prod_l g(\phi_{il} - \phi_{F_{il}}; 0, \hat{\sigma}_\phi^2) \times g(\theta_{il} - \theta_{F_{il}}; 0, \hat{\sigma}_\theta^2) \tag{4}$$

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. □

[3] For 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 1 we define the similarity metric γ_{ij} as the likelihood that two client fingerprints share identical peaks:

Definition 4 (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:

$$\gamma_{ij} = \prod_{\phi_i \in \Phi_{il}, \phi_j \in \Phi_{jl}} g(\phi_i - \phi_j; 0, 2\sigma_\phi^2) \times \prod_{\theta_i \in \Theta_{il}, \theta_j \in \Theta_{jl}} g(\theta_i - \theta_j; 0, 2\sigma_\theta^2) \tag{5}$$

where $g(\cdot; \mu, \sigma^2)$ is from Definition 3, and the factor of 2 in the variance accounts for computing the difference of two normally distributed values. □

5.2 Defining the confidence metric

We notice that Eqs. (2), (4) and (5) fully define α_i for each client i. In summary, the confidence metric 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 metric. Algorithm 1 below summarizes the steps to construct α_i for a given client i. The computational complexity of obtaining the confidence metric for each client i depends on the number of servers m and clients c in the neighborhood $\mathcal{N}_i$ of i as shown in Algorithm 1, and for each client-server pair, the dominant complexity is in computing the fingerprint which can be done using the well-known MUSIC algorithm in $O(M \log(d))$ where d is the desired fingerprint resolution.

We now present our main result that solves Problem 1 in the problem statement (Sect. 3). 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).[4]

[4] This 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.

Algorithm 1 Algorithm to Compute Client Confidence Metric

▷ Input: Ratio of Channels $\hat{\mathbf{h}}_{il}$ and SNR
▷ Output: Confidence Metric, α_i for client i
▷ Step (1): Measure fingerprints for client i
for $l = 1, \ldots, m$ **do**
  for $\phi \in \{0°, \ldots, 360°\}$; $\theta \in \{0°, \ldots, 90°\}$ **do**
    Find $F_{il}(\phi, \theta)$ using measured $\hat{\mathbf{h}}_{il}$ (Eq. 1)
  end for
end for
▷ Step (2): Measure variances in peak locations using SNR
$\sigma_\theta^2 = \sigma_\phi^2$ = Apply Lemma 1 SNR
▷ Step (3): Find honesty, similarity and confidence metric
β_i = Apply Definition 3 using σ_θ^2, σ_ϕ^2, peaks of F_{il}
for $j = \{1, \ldots, c\} \setminus \{i\}$ **do**
  γ_{ij} = Apply Definition 4 using σ_θ^2, σ_ϕ^2, peaks of F_{il}, F_{jl}
end for
$\alpha_i = \beta_i \prod_{j \neq i}(1 - \gamma_{ij})$

Theorem 1 *Consider a network with m servers and c clients. A new client i either:* (1) *spoofs s clients reporting a random location, potentially scaling power, or;* (2) *is a uniformly randomly located legitimate client. Let α_{spoof}, α_{legit} be the confidence metrics in either case. Assume that the client obtains its signals from servers along k paths (where the number of paths k is defined by Eq.* (1) *in Sect.* 4*). Under A.1–A.3, the expected α_{spoof}, α_{legit} are bounded by:*

$$E[\alpha_{spoof}] \leq \left[\sqrt{\hat{\sigma}_\theta \hat{\sigma}_\phi}\kappa\right]^m [2mk\sigma_\theta\sigma_\phi]^s$$
$$E[\alpha_{legit}] \geq 1 - cm\hat{\sigma}_\theta\hat{\sigma}_\phi \left[\sqrt{2\sigma_\theta\sigma_\phi}\kappa\right]^{mk} \quad (6)$$

where $\kappa = \left((\sqrt{2} + \sqrt{\pi})/\pi\right)^2$, σ_θ, σ_ϕ, $\hat{\sigma}_\theta$, $\hat{\sigma}_\phi$ *are the variances defined in Lemma* 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 Gil et al. 2015a). 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 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 v 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. $E[g(v)]$ and $E[g(u)]$ are as follows:

$$E[g(v)] = \int_{-S}^{S} f(x;0,\sigma^2)dx + \sqrt{8\pi}\int_{-\infty}^{-S}[f(x;0;\sigma^2)]^2 dx$$
$$\geq \int_{-S}^{S} f(x;0,\sigma^2)dx = \text{erf}\left(\frac{S}{\sigma\sqrt{2}}\right) \geq 1 - \sigma \quad (7)$$

where 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:

$$E[g(u)] = \int_{-S}^{S}\frac{1}{2\pi}dx + 2\sqrt{2\pi}\int_{-2\pi}^{-S}\frac{1}{2\pi}f(x;0;\sigma^2)dx$$
$$\leq \frac{S}{\pi} + \frac{1}{\sqrt{2\pi}}\left(1 - \text{erf}\left(\frac{S}{\sigma\sqrt{2}}\right)\right) \leq \sqrt{\sigma}\kappa \quad (8)$$

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:

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

Applying a similar argument, the similarity metric γ is:

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

Combining the above equations, we prove Eq. (6). □

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 Sect. 9 show that our bounds on α approximately predict performance in general environments. Further, Sect. 9.1 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 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 shows how α readily integrates with robotic coverage controllers, in particular.

6 Threat-resistant locational coverage

This section describes how our spoof detection method from Sect. 5 integrates with well-known coverage controllers from (Cortes et al. 2004; Schwager et al. 2009a, b). 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 function are determined by client positions P, e.g., $\rho(q, P) = \rho_1(q) + \cdots + \rho_c(q)$ where $\rho_i(q)$ quantifies the influence of client i's position on the importance function. Using (Cortes et al. 2004; Schwager et al. 2009a, b), 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: $\rho(q, P)_\alpha = \alpha_1\rho_1(q) + \cdots + \alpha_c\rho_c(q)$. Given the properties of these weights derived in Theorem 1, 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. 5); providing resilience to the spoofing attack.

For simplicity, we assume the importance function $\rho(q)$ is static (from Cortes et al. (2004)) 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 Schwager et al. (2009a, b) 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 problem parameters like signal-to-noise ratio. Theorem 2 below solves Problem 2 of our problem statement (Sect. 3).

Theorem 2 *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 $\{\alpha_1, \ldots, \alpha_c\}$ be a set of confidence weights satisfying Theorem 1 and assume a known importance function $\rho(q, P) = \rho_1(q) + \cdots + \rho_c(q)$ that is defined over the environment $\mathcal{Q} \subset \mathbb{R}^3$ of size $D(\mathcal{Q})$. Define*

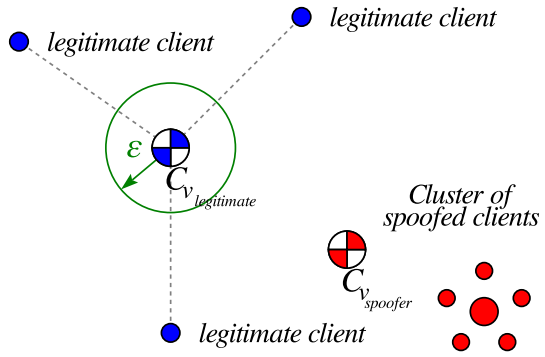

Fig. 5 Coverage guarantee: an ϵ ball around the ground-truth centroid, $C_{V_{\text{legitimate}}}$, is shown as a *circle* with radius epsilon. Theorem 2 finds $\epsilon(\mathcal{P})$ so that server positions remain in this ball in the presence of spoofed clients

$C_V = \{x_1^*, \ldots, x_m^*\}$ 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 $\rho(q, P)_\alpha = \alpha_1\rho_1(q) + \cdots + \alpha_c\rho_c(q)$ where there is at least one spoofed client, i.e., $|S| \geq 1$. If $\{\alpha_1, \ldots, \alpha_c\}$ satisfy Theorem 1, we have that $\forall x \in C_{V_\alpha}$ there exists a unique $y \in C_V$, where in the expected case dist$(x, y) \leq \epsilon(m, s, \sigma_\phi, \sigma_\theta, \kappa)$ where

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

and $m, s, \sigma_\phi, \sigma_\theta, \kappa$ are problem parameters as in Theorem 1.

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:

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

is the expected maximal effect that the presence of spoofed clients can have on the importance function. Intuitively, in the expected case, 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. 5). Specifically, recall from Cortes et al. (2004) 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 dist$(q, x_l) <$ 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 cen-

troid over legitimate nodes as $C_{V_S} = C_{V_L} + \mathbf{v}\|\mathbf{e}\|$ where $\mathbf{v}$ is an arbitrary unit vector and the magnitude of $\mathbf{e}$ can be as large as the length of the operative environment, $\|\mathbf{e}\| \leq D(\mathcal{Q})$. 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 $C_{V_\alpha} = \frac{C_{V_L} + b\mathbf{v}\|\mathbf{e}\|}{bT + (a-b)M_{V_L}}$. Combining this expression with the server control input:

$$u_l = k\left([(a+b)C_{V_L} - p_l] + b\|\mathbf{e}\|\mathbf{v} \right) \tag{9}$$

where $k = -2(bT + aM_{V_L})$. If $(a+b) = 1$, this control input drives the server robot l to a neighborhood of size $\epsilon = b\|\mathbf{e}\| \leq bD(\mathcal{Q})$ 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 1, Eq. (6), then:

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

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

7 Threat-resistant drone delivery

The previous section describes an application of α from Sect. 5 as continuous weights to bound the influence of adversarial clients. While this approach is useful for problems of a continuous nature like coverage, other problems in control require a more discrete approach. For example, in delivery problems a decision must be made whether to visit a client site or not since traversing a path some fraction of its length is equivalent to not visiting the client site at all. In other words, it is an inherently binary decision problem. This section shows how the α weights from Sect. 5 can be used as a *classifier* to select a subset of clients to be serviced, as in a drone delivery context. The drone delivery problem is described in Problem 3. The result below shows that the total path length traversed in the drone delivery problem is impacted by the presence of spoofed nodes to within a closed-form bound, that depends on problem parameters like the signal-to-noise ratio.

Theorem 3 *Let x be the server robot position 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*

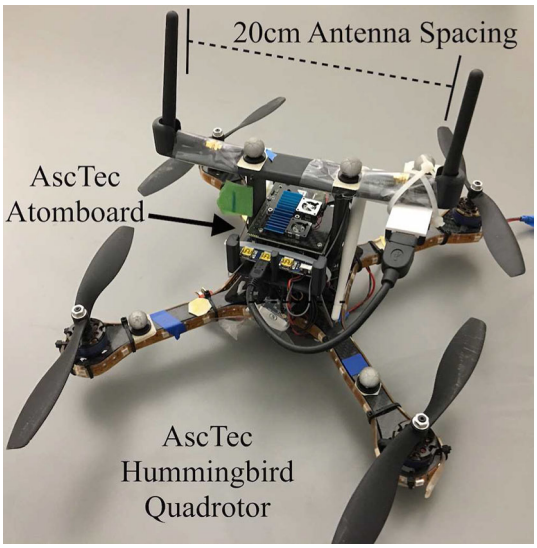

Fig. 6 Server platform: our aerial servers measure signal fingerprints for each client using the two antennas shown

$\{\alpha_1, \ldots, \alpha_c\}$ be a set of confidence weights satisfying Theorem 1 and environment size $D(\mathcal{Q})$. There exists a decision threshold $T > 0$ such that the indicator function defined as:

$$I_{\alpha_i} = \begin{cases} 1 & \alpha_i > T \\ 0 & otherwise \end{cases}$$

for each client $i \in \{1, \ldots, c\}$, can be derived to determine whether client i will be serviced by the delivery drone, i.e., $I_{\alpha_i} = 1$. Using this indicator function we define the total path length covered by the server to be $L = \sum_{p_i \in P} I_{\alpha_i} dist(p_i, x)$. Let $L_{legit} = \sum_{p_i \in \tilde{S}} dist(p_i, x)$ be the total path length covered by the server in the optimal case of no spoofed nodes. Then the difference in expectations is bounded such that:

$$|E[L] - E[L_{legit}]| \leq max(|S|, |\tilde{S}|)bD(\mathcal{Q}) \tag{10}$$

$$= max(|S|, |\tilde{S}|)\epsilon \tag{11}$$

where $\epsilon = bD(\mathcal{Q})$, $b = \max\left\{ [\sqrt{\hat{\sigma}_\theta \hat{\sigma}_\phi}\kappa]^m [2mk\sigma_\theta\sigma_\phi]^s, cm\hat{\sigma}_\theta\hat{\sigma}_\phi[\sqrt{2\sigma_\theta\sigma_\phi}\kappa]^{mk} \right\}$, and $m, s, \sigma_\phi, \sigma_\theta, \kappa$ are problem parameters as in Theorem 1.

Proof For each client $i \in 1, \ldots, c$, let us denote:

$$I_{\alpha_i} = \begin{cases} 1 & \alpha_i > T \\ 0 & \text{otherwise} \end{cases}$$

where T is a constant chosen so that:

$$E[\alpha_i] = \int_0^1 P(\alpha_i > x)dx \tag{12}$$

$$= P(\alpha_i > T) \quad \text{(using Mean Value Theorem)} \tag{13}$$

$$= E[I_{\alpha_i}] \tag{14}$$

The last equation holds from the fact that I_{α_i} is an indicator function for the event $\alpha_i > T$. Note that here we show the existence of such a T, but we do not find an analytical value for T. In Sect. 9 however, we show the empirical performance of the median threshold $T = 0.5$. We can then write the expected total path length of the delivery drone as:

$$E[L] = E\left[\sum_{p_i \in P} I_{\alpha_i} \text{dist}(p_i, x)\right] \tag{15}$$

$$= \sum_{p_i \in P} E[I_{\alpha_i}]\text{dist}(p_i, x) \tag{16}$$

$$= \sum_{p_i \in P} E[\alpha_i]\text{dist}(p_i, x) \text{ (using Eq. 14)} \tag{17}$$

$$= \sum_{p_i \in \tilde{S}} E[\alpha_i]\text{dist}(p_i, x) + \sum_{p_l \in S} E[\alpha_l]\text{dist}(p_l, x) \tag{18}$$

Recall from Theorems 1 and 2 that we can bound $E[\alpha_i]$ as:

$$E[\alpha_{i,spoof}] \leq \epsilon / D(\mathcal{Q}) \qquad E[\alpha_{i,legit}] \geq 1 - \epsilon / D(\mathcal{Q})$$

where

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

Applying the above bounds to Eq. (18), we have:

$$E[L] \leq \sum_{p_i \in \tilde{S}} \text{dist}(p_i, x) + \frac{\epsilon}{D(\mathcal{Q})} \sum_{p_l \in S} \text{dist}(p_l, x)$$

$$\leq E[L_{legit}] + |S|\epsilon$$

$$E[L] \geq \sum_{p_i \in \tilde{S}} \left(1 - \frac{\epsilon}{D(\mathcal{Q})}\right) \text{dist}(p_i, x) + \sum_{p_l \in S} 0$$

$$\geq E[L_{legit}] - |\tilde{S}|\epsilon$$

Combining the above two equations, we conclude that:

$$|E[L] - E[L_{legit}]| \leq \max(|S|, |\tilde{S}|)\epsilon$$

which proves the theorem. □

8 General multi-robot control problems

The above sections demonstrate two modalities of integrating the confidence metric α to secure multi-robot controllers: either as a continuous per-agent weight, or as a means to classify agents as legitimate or spoofed. Theorems 2 and 3 show theoretical bounds on the influence of adversaries to controllers in the coverage and unmanned delivery contexts. Further, empirical results in Sect. 9 demonstrate that α performs well when applied both in continuous and discrete settings. However, it is natural to ask which of these two modes ought be applied to secure any given multi-robot problem of interest, beyond coverage and unmanned delivery. In this regard, we make the following observations:

8.1 Applying α as continuous weights

For many control objectives, the contribution of each agent to the total optimization function is naturally expressed as a continuous quantity. In these contexts, a natural modality to integrate α is to incorporate it as a per-agent weight that directly reduces the contributions of spoofed agents to the optimization function. Doing so has two key advantages: (1) It enables provable bounds in expectation on the influence of spoofers to the multi-robot objective (akin to Theorem 2). (2) Per-client weighting limits the extent to which spoofed agents can influence the controller in the worst-case.

8.2 Applying α to decision-based problems

Unfortunately, many problems do not allow for a continuous weighting since their objectives are inherently discrete decisions on each agent in the network (e.g., unmanned delivery). In these cases, α can still be used to derive an indicator function that classifies agents as legitimate or spoofed. This modality still allows for obtaining bounds in expectation on the influence of spoofers (akin to Theorem 3). However, by the sheer nature of these problems, false positives or negatives have a greater impact on the objective function in the worst-case.

9 Experimental results

This section describes our results from an experimental evaluation of our theoretical claims. We implemented two aerial servers on AscTec Hummingbird quadrotors. Each server (Fig. 6) was comprised of an AscTec Atomboard onboard computer and two $2dBi$ antennas, spaced 20cm apart, attached to an Intel 5300 Wi-Fi card which estimates the wireless channels from each client to each antenna via the 802.11n CSI tool (Halperin et al. 2011). We note that each aerial server only needs to perform a single spin to simultaneously measure the fingerprints of all clients in our experimental setup (Step 1 of Algorithm 1). Our clients were ten iRobot Create robots, each equipped with Asus EEPC netbooks and single-antenna Wi-Fi cards (Fig. 7). An adversarial client forged multiple identities by spawning multiple

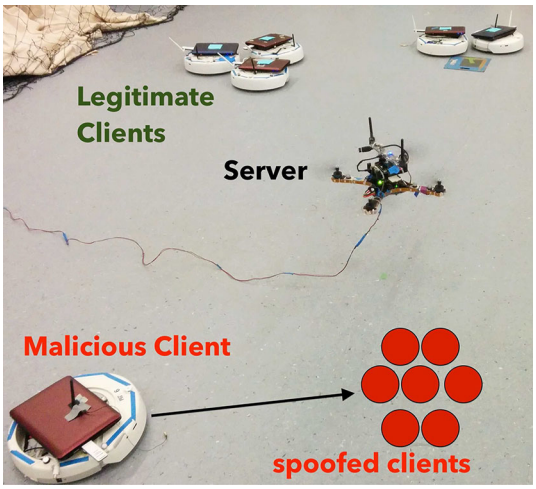

Fig. 7 Hardware evaluation: depicts an example robot network within our experimental setup with a quadrotor server and several mobile clients

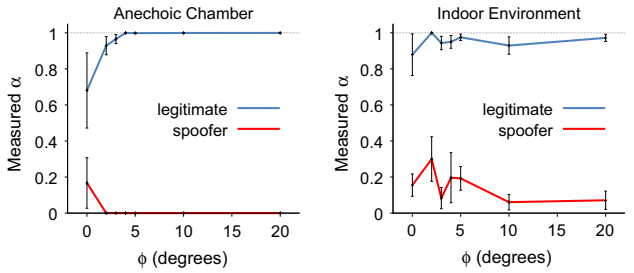

(a) Varying ϕ (Free-space) **(b)** Varying ϕ (Multipath)

Fig. 8 Co-aligned clients: we vary the angle ϕ between a legitimate and malicious client, relative to a single server (as shown in Fig. 9b), and plot α in **a** an anechoic chamber and **b** an indoor environment. The minimum ϕ needed to distinguish between clients is only: **a** 3° in freespace, **b** 0° in multipath settings

(a) Chamber **(b)** ϕ

Fig. 9 Microbenchmarks on α: **a** an anechoic chamber simulating freespace. **b** We measure α while varying the angle between a legitimate and malicious client, relative to the robotic server

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 (Sheng et al. 2008).

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°. This was calculated using Eq. (3) from Lemma 1 with $\lambda = 5.4$ cm, $r = 20$ cm and a worst case number of packets and SNR being $M = 20$ and 10 dB respectively. We note that our chosen antenna spacing of $r = 20$ cm is small enough to be accommodated on a quadrotor while still providing a high spatial resolution by Lemma 1 and as shown in our experimental results on the angular resolution of our confidence metric (Sect. 9.1; Fig. 8). After including the standard deviation in reported location, based on the known errors of our Vicon-based localization framework, this increased the average $\hat{\sigma}_\theta, \hat{\sigma}_\phi$ by 2°. We compare our system against a baseline that uses a Received Signal Strength (RSSI) comparison (akin to Pires et al. 2004).

Roadmap We conduct four classes of experiments: (1) Microbenchmarks to validate our client confidence metric, both in free-space and multipath indoor environments (Sect. 9.1). (2) Experiments applying this confidence metric to quarantine adversaries (Sect. 9.2). Application of our system to secure against Sybil attacks: (3) the coverage problem (Sect. 9.3); (4) the drone delivery problem (Sect. 9.4).

Table 2 Summarized classification performance: true positive rates (TPR) and false positive rates (FPR) for classifying clients as spoofed, when $\alpha < 0.5$ in our system, and with a 2 dB minimum dissimilarity for RSSI

	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

9.1 Microbenchmarks on the confidence metric

This experiment studies the correctness of our system's confidence metric α. Recall from theory in Sect. 5 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 (Fig. 9a) which attenuates reflected paths by about 60 dB, for a legitimate

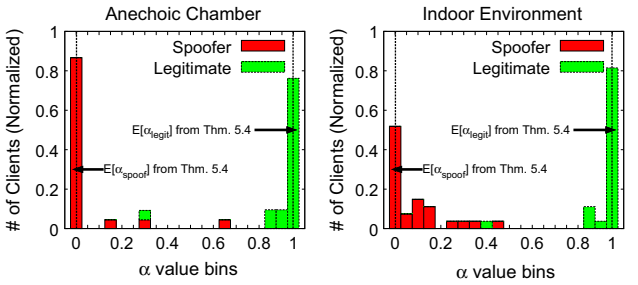

(a) α Histogram (Free-space) **(b)** α Histogram (Multipath)

Fig. 10 Experimental evaluation of α: **a** in an anechoic chamber approximating our assumptions A.1–A.3 (Sect. 1), 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

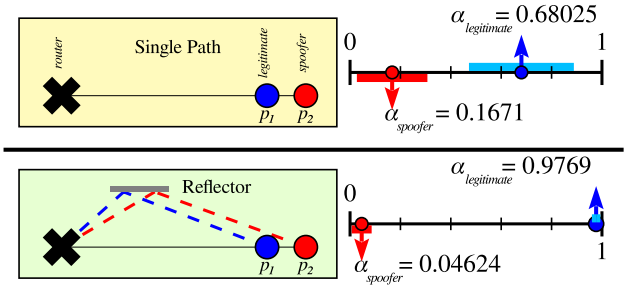

Fig. 11 Anechoic chamber multipath: we measure α for a spoofing client coaligned with a legitimate client ($\phi = 0°$) in the anechoic chamber before and after adding a reflector to introduce multipath. The increased separation of α and lower standard deviation (shown as *bars*) is depicted on the right

and malicious client from one server robot 12 m away. We also introduced a metallic reflector in this controlled setting, to measure the contribution of multipath to α. Next, in a 10 m $\times$ 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. 10, the values of α in the anechoic chamber tightly follow our theoretical bounds in Theorem 1. 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. We consider a spoofing adversary initially co-aligned with a legitimate client as the angle of separation, ϕ, increases from 0° to 20° relative to the server robot (Fig. 9b). Figure 8 depicts the measured α values for the legitimate and spoofed clients. In the anechoic chamber at ϕ close to 0°, the fingerprints of the legitimate and adversarial nodes are virtually identical: each has precisely one peak at 0°. Consequently, α for the legitimate node is much below 1, indicating that the legitimate client is believed to be adversarial (i.e., the term $(1-\gamma)$ in α approaches zero in Eq. 2). However, α for the legitimate client quickly approaches 1 at only $\phi = 3°$ in the anechoic chamber. In fact, α is virtually identical to 1 beyond 10°, indicating that a single server robot can distinguish between closely aligned legitimate and adversarial clients even in free-space. To evaluate the effects of multipath on the α values of coaligned clients in a controlled manner, we positioned a small metallic reflector several meters away from the two clients and server in the anechoic chamber when $\phi = 0°$. Figure 11 demonstrates that the the additional reflected signal paths strongly disambiguate the α values for coaligned clients. Specifically, the term $(1-\gamma)$ in Eq. (2) approaches zero only for the adversary. Figure 8b depicts the larger separation of α values for coaligned clients in a typical indoor setting compared to free-space. As expected, multipath reflections from walls and obstacles clearly distinguish spoofing clients from legitimate clients even at $\phi = 0°$.

9.2 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 clients—either ten legitimate clients, or nine legitimate clients and an adversary reporting two to nine additional 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. We compare our system to a baseline RSSI classifier using a thresholded minimum dissimilarity, a technique previously applied in static networks (Pires et al. 2004; Wang and Yang 2013). Measured signal-to-noise ratios for clients ranged from 5 to 25 dB. In our system, quadrotor servers performed classification by applying a threshold using the measured α values for each client.

Results In Fig. 12, we measure true-positives against false-positives collected over multiple network topologies, resulting in the well-known Receiver Operating Characteristics (ROC) curves (Fawcett 2004). Our theoretical results in Sect. 7 indicate that α measurements are suitable for use in a thresholding classification context. Empirically, Fig. 10 shows that a threshold of $\alpha < 0.5$ performs well to classify clients as spoofed. Table 2 summarizes our performance results when using this threshold for each of the three attack modalities, compared to RSSI-based classification where a 2 dB thresholded minimum dissimilarity performed best.

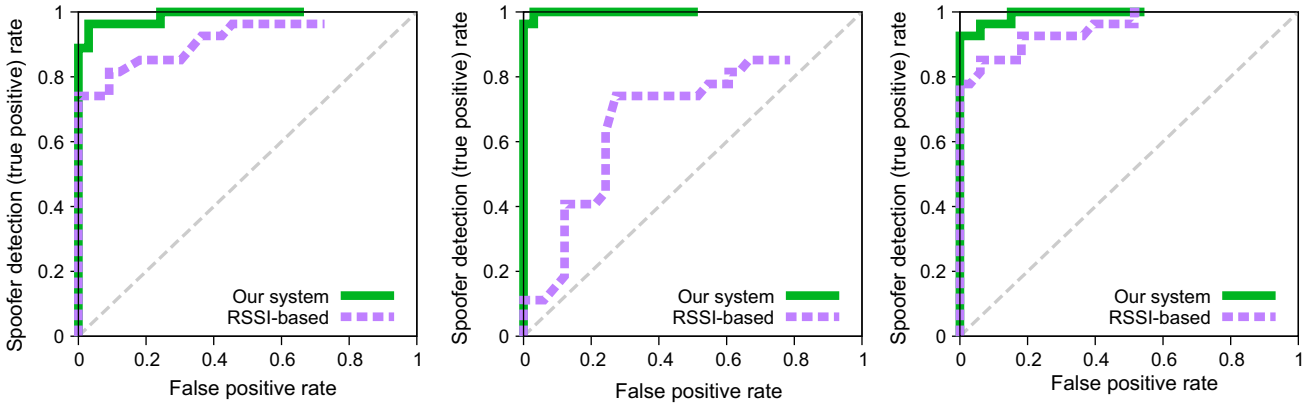

(a) Static adversaries **(b)** Power-Scaling adversaries **(c)** Mobile adversaries

Fig. 12 Receiver operating characteristics: we measure ROC *curves* for adversaries which **a** are static; **b** scale power differently while spoofing different clients; and **c** are mobile. We compare the performance of our system against a baseline using received signal power

In particular, our classifier is robust to power-scaling Sybil attacks (where RSSI performs poorly) since we use the ratio of wireless channels in computing α (Sect. 4). Our client classifier exhibits consistent performance in both power-scaling and mobile scenarios with a TPR $\approx$ 96% and FPR $\approx$ 4%.

9.3 Application to multi-agent coverage

We implement the multi-agent coverage problem from Cortes et al. (2004), 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 $\rho(q, P) = \rho_1(q) + \cdots + \rho_c(q)$ defined in Sect. 6 where each client term is a Gaussian-shaped function $\rho_i(q) = \exp(-\frac{1}{2}(q-p_i)^T(q-p_i))$ (Fig. 13b). An α-modified importance function is implemented as $\rho(q, P)_\alpha = \alpha_1\rho_1(q) + \cdots + \alpha_c\rho_c(q)$ where the α terms are computed using Algorithm 1 (Fig. 13c).

Method This experiment was performed in the multipath-rich indoor testbed. For each experiment we randomly place three clients in an 8 m $\times$ 10 m room with two AscTec quadrotor servers. Figure 13a–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 (s3) our system.

Results Figure 13a–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 an oracle's performance. Figure 13d demonstrates the ability of our system to bound the service cost to near optimal even as additional spoofers enter the network (comprising up to 300%).

Aggregate results Across multiple topologies and 12 runs, for a system with no security the maximum distance from each quadrotor to an oracle solution is on average 3.77 m (SD 0.86). In contrast, our system achieves a 0.02 m (SD 0.02) average distance from an oracle solution (Fig. 14).

9.4 Application to unmanned delivery

This experiment applies our Sybil attack detection algorithm in the context of unmanned delivery. Specifically, we consider a delivery quadrotor that iteratively visits multiple client locations from a depot to deliver packages, for instance delivering relief material in a disaster area. An adversarial power-scaling client spawning multiple non-existent client locations could readily disrupt such a system, drawing the delivery robot away to service regions where no clients exist. We study the effectiveness of our system in guarding against such attacks and compare it against the RSSI baseline (Sect. 9.2).

Method Multiple heuristics exist for approximating optimal solutions to unmanned delivery problems which minimize distance, payload, or fuel usage (Laporte et al. 1988; Pavone et al. 2011). We use a simple distance metric—the shortest quadrotor flight path which visits all client locations iteratively, returning to the depot each time—and deploy a system that uses our binary classifier based on signal fingerprints to filter malicious clients. We compare our results both against a baseline classifier based on RSSI as well as a naive system which visits every reported client location. We repeat the experiment across ten randomly chosen topologies. Figure 14a depicts a candidate topology where two legitimate clients report their positions p_1 and p_2 to a quadrotor beginning its delivery route at location x, while a malicious client at position p_3 reports six (inclusive) false client locations. The average minimum trajectory length for

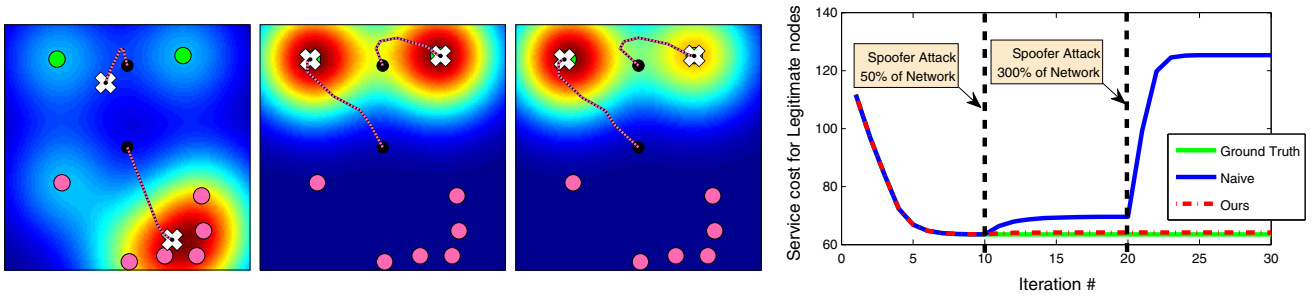

(a) No security **(b)** Oracle **(c)** Our System **(d)** Cost

Fig. 13 Experimental results for Sybil attack in multi-agent coverage: depicts the total distance of converged quadrotor server positions (white $\times$) to two legitimate clients *light filled circle* at top of figure and six spoofed clients *dark filled circles* at bottom right of figure. 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

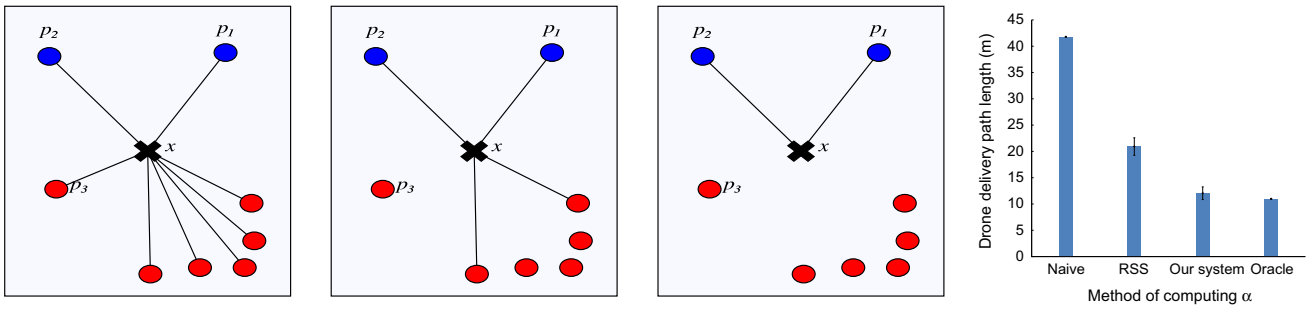

(a) Naive System **(b)** Received power baseline **(c)** Our System **(d)** Aggregate Results

Fig. 14 Path of delivery robot: depicts sample trajectories of a delivery robot iteratively visiting a reported client location and returning to a depot among two legitimate clients and an adversary spoofing six clients, for: **a** a naive system with no security; **b** a baseline classifier based on received signal power; and **c** our system. **d** Depicts the mean and standard deviation of total trajectory lengths across ten scenarios

the quadrotor to visit all 8 clients across our topologies is 41.78 m.

Results Figure 14a–c depicts candidate trajectories of the quadrotor in the three scenarios: (1) A naive system without cyber-security; (2) The RSSI-baseline; (3) Our system. In the RSSI baseline, the quadrotor compares the received power per packet for each client, but misclassifies a subset of the spoofed clients as legitimate (owing to noise), resulting in the quadrotor traveling a mean path length of 20.92 m. In contrast, our system benefits from the large margin of separation when classifying clients using their α value (as in Sect. 9.2), with the quadrotor's resultant mean path length of 12.05 m performing close to an oracle system's ground truth trajectory length of 10.91 m across topologies (see Fig. 14d).

10 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 and unmanned delivery, our approach can be readily extended to secure other multi-robot controllers against Sybil attacks, e.g., applications within the Vehicle Routing Problem (Laporte et al. 1988; Pavone et al. 2011), in search-and-rescue tasks (Lin et al. 2009), and in formation control (Wang et al. 2007). We note for future work that our method of detecting spoofed clients is applicable to servers as well, since they also communicate wirelessly. Additionally, while this paper addresses Sybil attacks in which spoofed clients assume unique identities, our approach generalizes to defense against replay attacks (Feng et al. 2011; Miao et al. 2013) where adversaries imitate existing legitimate clients in the network. Since our approach is based on the fundamental physics of wireless signals, we believe that it also applies to other Wi-Fi based security issues in robot swarms such as packet path validation (Liu et al. 2008) and detecting packet injection attacks to name a few.

Acknowledgements 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

Amazon prime air. http://www.amazon.com/b?node=8037720011.

Adib, F., Kumar, S., Aryan, O., Gollakota, S., & Katabi, D. (2013). Interference Alignment by Motion. In *MOBICOM*.

Beard, R., McLain, T., Nelson, D., Kingston, D., & Johanson, D. (2006). Decentralized cooperative aerial surveillance using fixed-wing miniature UAVs. *Proceedings of the IEEE*, *94*(7), 1306–1324. doi:10.1109/JPROC.2006.876930.

Chapman, A., Nabi-Abdolyousefi, M., & Mesbahi, M. (2009). Identification and infiltration in consensus-type networks. In *1st IFAC Workshop on Estimation and Control of Networked Systems*.

Cortes, J., Martinez, S., Karatas, T., & Bullo, F. (2004). Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, *20*(2), 243–255. doi:10.1109/TRA.2004.824698.

Daniel, K., Dusza, B., Lewandowski, A., & Wietfeld, C. (2009). AirShield: A system-of-systems MUAV remote sensing architecture for disaster response. In *Systems conference, 2009 3rd Annual IEEE* (pp. 196–200) doi:10.1109/SYSTEMS.2009.4815797.

Douceur, J. (2002). The sybil attack. In P. Druschel, F. Kaashoek, & A. Rowstron (Eds.), *Peer-to-Peer systems, Lecture Notes in Computer Science* (vol. 2429, pp. 251–260). Berlin: Springer. doi:10.1007/3-540-45748-8_24.

Fawcett, T. (2004). *ROC graphs: Notes and practical considerations for researchers*. Technical Report.

Feng, Z., Ning, J., Broustis, I., Pelechrinis, K., Krishnamurthy, S.V., & Faloutsos, M. (2011). Coping with packet replay attacks in wireless networks. In *Sensor, mesh and ad hoc communications and networks (SECON), 2011 8th Annual IEEE Communications Society Conference on* (pp. 368–376). IEEE.

Fitch, P. J. (1988). *Synthetic aperture radar*. Berlin: Springer.

Gazzah, H., & Marcos, S. (2003). Directive antenna arrays for 3D source localization. In *Signal processing advances in wireless communications, 2003. SPAWC 2003. 4th IEEE Workshop on* (pp. 619–623). doi:10.1109/SPAWC.2003.1319035.

Gazzah, H., & Marcos, S. (2006). Cramer-Rao bounds for antenna array design. *IEEE Transactions on Signal Processing*, *54*, 336–345. doi:10.1109/TSP.2005.861091.

Gil, S., Kumar, S., Katabi, D., & Rus, D. (2013). *Adaptive communication in multi-robot systems using directionality of signal strength*. ISRR.

Gil, S., Kumar, S., Mazumder, M., Katabi, D., & Rus, D. (2015a). Guaranteeing spoof-resilient multi-robot networks. In *Full paper version with supplementary material available as a TECH REPORT at MIT CSAIL Publications and Digital Archive*. http://publications.csail.mit.edu.

Gil, S., Kumar, S., Mazumder, M., Katabi, D., & Rus, D. (2015b). Guaranteeing spoof-resilient multi-robot networks. In *Proceedings of robotics: Science and systems*. Rome, Italy.

Goldsmith, A. (2005). *Wireless communications*. Cambridge: Cambridge University Press.

Halperin, D., Hu, W., Sheth, A., & Wetherall, D. (2011). Tool release: Gathering 802.11n traces with channel state information. ACM SIGCOMM. *Computer Communication Review*, *41*(1), 53.

Hayes, M. H. (1996). *Statistical digital signal processing and modeling* (1st ed.). New York, NY, USA: Wiley.

Higgins, F., Tomlinson, A., & Martin, K.M. (2009). Threats to the swarm: Security considerations for swarm robotics. *International Journal on Advances in Security*, *2*.

Jin, D., & Song, J. (2014). 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* (pp. 281–286). doi:10.1109/ICIS.2014.6912147. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6912147&tag=1.

Klausing, H. (1989). Feasibility of a SAR with rotating antennas (ROSAR). In *Microwave Conference, 1989*.

Kumar, S., Gil, S., Katabi, D., & Rus, D. (2014). Accurate indoor localization with zero start-up cost. In *Proceedings of the 20th annual international conference on mobile computing and networking, MobiCom '14* (pp. 483–494). New York, NY, USA: ACM. doi:10.1145/2639108.2639142.

Kumar, S., Hamed, E., Katabi, D., & Erran Li, L. (2014). LTE radio analytics made easy and accessible. In *Proceedings of the 2014 ACM conference on SIGCOMM, SIGCOMM '14* (pp. 211–222). New York, NY, USA: ACM. doi:10.1145/2619239.2626320.

Laporte, G., Nobert, Y., & Taillefer, S. (1988). Solving a family of multi-depot vehicle routing and location-routing problems. *Transportation Science*, *22*(3), 161–172.

Levine, B. N., Shields, C., & Margolin, N. B. (2006). *A survey of solutions to the sybil attack*. Techincal Report, Document Number 2006-052. Amherst: University of Massachusetts Amherst

Lin, L., & Goodrich, M. A. (2009). UAV intelligent path planning for wilderness search and rescue. In *Intelligent robots and systems, 2009. IROS 2009. IEEE/RSJ International Conference on* (pp. 709–714). IEEE.

Liu, H., Wang, Y., Liu, J., Yang, J., & Chen, Y. (2014). Practical user authentication leveraging channel state information (CSI). In *Proceedings of the 9th ACM symposium on information, computer and communications security, ASIA CCS '14* (pp. 389–400). New York, NY, USA: ACM. doi:10.1145/2590296.2590321.

Liu, X., Li, A., Yang, X., & Wetherall, D. (2008). Passport: Secure and adoptable source authentication. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation, NSDI'08* (pp. 365–378). Berkeley, CA, USA: USENIX Association. http://dl.acm.org/citation.cfm?id=1387589.1387615.

Liu, Y., Bild, D., Dick, R., Mao, Z.M., & Wallach, D. (2014). The Mason test: A defense against Sybil attacks in wireless networks without trusted authorities. *CORR, abs/1403.5871*. http://dblp.unitrier.de/rec/bib/journals/corr/LiuBDMW14.

Malmirchegini, M., & Mostofi, Y. (2012). On the spatial predictability of communication channels. *IEEE Transactions on Wireless Communications*, *11*(3), 964–978.

Mathews, C. P., & Zoltowsk, M. D. (1994). *Signal Subspace Techniques for Source Localization with Circular Sensor Arrays*. Purdue University TechReport. http://docs.lib.purdue.edu/cgi/viewcontent.cgi?article=1175&context=ecetr

Miao, F., Pajic, M., & Pappas, G.J. (2013). Stochastic game approach for replay attack detection. In *Decision and control (CDC), 2013 IEEE 52nd annual conference on* (pp. 1854–1859). IEEE.

Newsome, J., Shi, E., Song, D., & Perrig, A. (2004). The sybil attack in sensor networks: Analysis defenses. In *Information processing in sensor networks, 2004. IPSN 2004. Third international symposium on* (pp. 259–268). doi:10.1109/IPSN.2004.1307346.

Olfati-Saber, R., & Murray, R. (2004). Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, *49*(9), 1520–1533. doi:10.1109/TAC.2004.834113.

Parker, L. E. (2002). Distributed algorithms for multi-robot observation of multiple moving targets. *Autonomous Robots*, *12*, 231–255.

Pavone, M., Frazzoli, E., & Bullo, F. (2011). Adaptive and distributed algorithms for vehicle routing in a stochastic and dynamic

environment. *IEEE Transactions on Automatic Control*, *56*(6), 1259–1274.
Pires W. R., de Paula Figueiredo, T., Wong, H., & Loureiro, A. (2004). Malicious node detection in wireless sensor networks. In *Parallel and distributed processing symposium, 2004. Proceedings. 18th international* (p. 24). doi:10.1109/IPDPS.2004.1302934.
Sargeant, I., & Tomlinson, A. (2013). Modelling malicious entities in a robotic swarm. In *Digital avionics systems conference (DASC), 2013 IEEE/AIAA 32nd*.
Schwager, M., Julian, B. J., & Rus, D. (2009). Optimal coverage for multiple hovering robots with downward facing cameras. In *Robotics and automation, 2009. ICRA '09. IEEE international conference on* (pp. 3515–3522). doi:10.1109/ROBOT.2009.5152815.
Schwager, M., Rus, D., & Slotine, J. J. (2009). Decentralized, adaptive coverage control for networked robots. *The International Journal of Robotics Research*, *28*(3), 357–375. http://ijr.sagepub.com/content/28/3/357.abstract.
Sheng, Y., Tan, K., Chen, G., Kotz, D., & Campbell, A. (2008). Detecting 802.11 MAC layer spoofing using received signal strength. In *INFOCOM 2008. The 27th Conference on Computer Communications*. IEEE. doi:10.1109/INFOCOM.2008.239. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4509834&tag=1.
Stoica, P., & Arye, N. (1989). Music, maximum likelihood, and Cramer-Rao bound. *IEEE Transactions on Acoustics, Speech and Signal Processing*, *37*(5), 720–741. doi:10.1109/29.17564.
Tse, D., & Vishwanath, P. (2005). *Fundamentals of wireless communications*. Cambridge: Cambridge University Press.
Wang, J., & Katabi, D. (2013). Dude, Where's my card?: RFID positioning that works with multipath and non-line of sight. In *SIGCOMM*.
Wang, T., & Yang, Y. (2013). Analysis on perfect location spoofing attacks using beamforming. In *INFOCOM, 2013 Proceedings IEEE* (pp. 2778–2786). doi:10.1109/INFCOM.2013.6567087. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6567087.
Wang, X., Yadav, V., & Balakrishnan, S. (2007). Cooperative UAV formation flying with obstacle/collision avoidance. *IEEE Transactions on Control Systems Technology*, *15*(4), 672–679.
Wang, Y., Attebury, G., & Ramamurthy, B. (2006). A survey of security issues in wireless sensor networks. *Communications Surveys Tutorials, IEEE*, *8*(2), 2–23. doi:10.1109/COMST.2006.315852.
Xiao, L., Greenstein, L., Mandayam, N. B., & Trappe, W. (2009). Channel-based detection of sybil attacks in wireless networks. *IEEE Transactions on Information Forensics and Security*, *4*(3), 492–503. doi:10.1109/TIFS.2009.2026454.
Xiong, J., & Jamieson, K. (2013). SecureArray: Improving Wifi security with fine-grained physical-layer information. In *Proceedings of the 19th annual international conference on mobile computing & networking, MobiCom '13* (pp. 441–452). New York, NY, USA: ACM. doi:10.1145/2500423.2500444.
Yang, J., Chen, Y., Trappe, W., & Cheng, J. (2013). Detection and localization of multiple spoofing attackers in wireless networks. *IEEE Transactions on Parallel and Distributed Systems*, *24*(1), 44–58. doi:10.1109/TPDS.2012.104.
Yang, Z., Ekici, E., & Xuan, D. (2007). A localization-based anti-sensor network system. In *INFOCOM 2007. 26th IEEE international conference on computer communications* (pp. 2396–2400). IEEE, doi:10.1109/INFCOM.2007.288. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4215870.

Stephanie Gil is a postdoctoral associate in the Distributed Robotics Lab at MIT where she works on secure and reliable multirobot coordination and control. Her focus is on developing both theory and experimental frameworks for making multirobot teams robust in realworld implementations. Her work is highly interdisciplinary and has been published in both communications and robotics fields. She obtained her Ph.D. in Aerospace Engineering at MIT and her undergraduate degree from Cornell University in Mechanical Engineering.

Swarun Kumar is assistant professor position at CMU where he works on wireless networks and systems. He designs and builds new systems that leverage a deep understanding of the wireless physical layer to enable faster wireless networks and deliver new services. His work has been featured as research highlights in the Communications of the ACM (CACM) and the International Journal of Robotics Research (IJRR). He is a recipient of the George Sprowls Award for best Ph.D. thesis in Computer Science at MIT and the President of India gold medal at IIT Madras. He obtained his Ph.D. in Computer Science at MIT and undergraduate degree from IIT Madras.

Mark Mazumder is an Associate Staff member at MIT Lincoln Laboratory. He received his A.B. in Computer Science from Harvard University in 2009. His professional interests include dependently typed proof assistants, machine learning, and decentralized robotics.

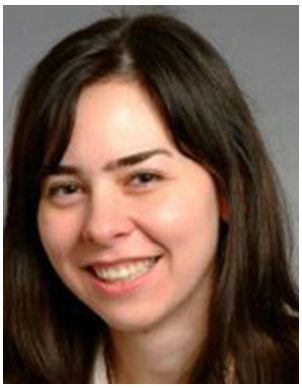

Dina Katabi is a Professor in the Department of Electrical Engineering and Computer Science and the director of MIT's research center, Wireless@MIT. Her work focuses on wireless networks, mobile applications, network security, and distributed resource management. She received her Bachelor of Science from Damascus University in 1995 and her M.S. and Ph.D. from MIT in 1999 and 2003. Her doctoral dissertation won an ACM Honorable Mention award and a Sprowls award for academic excellence. She has received best paper awards from ACM SIGCOMM and Usenix NSDI. She was awarded an NSF CAREER award in 2005, the NBX Career Development chair and a Sloan Fellowship in 2006, the IEEE William R. Bennett prize in 2009, a Faculty Research Innovation Fellowship in 2011, and the ACM Grace Murray Hopper Award and a MacArthur Foundation Fellowship in 2013.

Daniela Rus is a Professor in the EECS Department at MIT. She is the Director of the Computer Science and Artificial Intelligence Laboratory (CSAIL). Her research interests include distributed robotics, mobile computing and programmable matter. At CSAIL she has led numerous groundbreaking research projects in the areas of transportation, security, environmental modeling and monitoring, underwater exploration, and agriculture. She is the recipient of an NSF Career award and an Alfred P. Sloan Foundation fellowship. She is a class of 2002 MacArthur Fellow. She is a fellow of AAAI and IEEE. She earned her Ph.D. in Computer Science from Cornell University. Prior to coming to MIT, She was an assistant professor, associate professor, and professor in the Computer Science Department at Dartmouth College.