Spoof Resilient Coordination in Distributed and Robust Robotic Networks

Venkatraman Renganathan, *Graduate Student Member, IEEE*, Kaveh Fathian, *Member, IEEE*, Sleiman Safaoui, *Graduate Student Member, IEEE*, and Tyler Summers, *Member, IEEE*

***Abstract*— As cyber–physical networks become increasingly equipped with embedded capabilities, they are made vulnerable to malicious attacks with the increased number of access points available to attackers. A particularly pernicious attack is spoofing, in which a malicious agent spawns multiple identities and can compromise otherwise attack–resilient algorithms that rely on assumed network robustness structures. We generalize a class of resilient consensus strategies, known as weighted mean-subsequence-reduced (W-MSR) consensus, to further provide spoof resilience by incorporating a physical layer authentication. By comparing the physical fingerprints of received signals, legitimate agents can identify and isolate malicious agents that attempt spoofing attacks. A key technical contribution is to quantify worst case misclassification probability using distributionally robust Chebyshev bounds computed via semidefinite programming when the physical fingerprints of received signals are stochastic. Numerical simulations and experimental results illustrate the effectiveness of the proposed methods. Our framework is applicable to a variety of problems involving multirobot systems coordinating via wireless communication.**

***Index Terms*— Distributional robustness, graph robustness, network resiliency, robot coordination, spoof attack.**

I. Introduction

A LARGE and growing literature has emerged on security, resilience, and robustness of the cyber–physical systems in the presence of noncooperative and adversarial agents [1], [2]. Malicious agents under a distributed environment might gain an undesirable advantage by influencing neighboring legitimate agents in the network. One such pernicious attack is *spoofing*,[1] in which a malicious agent spawns multiple nonexistent identities or impersonates existing legitimate agents. Spoofing is not just an abstract concern; successful attacks have been realized in several critical networks, such as global navigation satellite systems [4] and antilock braking systems [5].

Autonomous multirobot systems are a rapidly emerging type of cyber–physical network, in which many tasks, including distributed estimation and cooperative manipulation, utilize distributed consensus protocols to coordinate an agreement on certain quantities of interest. Recently resilient consensus algorithms have been developed to prevent malicious agents from exerting undue influence, by designing sufficient redundancy in the algorithms and underlying network structures [6]–[9]. Such approaches have recently been applied to multirobot systems in [10] and [11]. Yet, spoofing attacks can easily compromise these otherwise attack resilient algorithms and network structures, which assumes an upper bound on the number of malicious agents in the network. Thus, malicious information easily propagates through the network leading to severe performance degradation or safety constraint violations.

In this brief, we address spoofing attacks in robust robotic networks. It was argued in [3] that any defense against a spoofing attack requires either a trusted central authority to certify (perhaps cryptographically) the identities of all legitimate agents in the network or a reliable method to distinguish physical fingerprints of signals received from neighboring agents. We focus here on the latter as the former is undesirable in distributed multirobot setting. Physical fingerprint analysis and discrimination have been used to detect spoofing in specific application contexts [12] but not in the context of generally distributed algorithms in cyber–physical networks or dynamic multirobot systems. Noise in the communication channel can result in a legitimate robot being wrongly classified as a spoofed neighbor, and vice versa, thereby emphasizing the need to quantify worst case misclassification probability of such scenarios.

Contributions: The current brief is a significant extension of our preliminary work in [13] where we generalize a class of resilient consensus strategies, known as weighted mean subsequence reduced (W-MSR) consensus, to provide spoof resilience by incorporating a physical fingerprint analysis of signals received from neighboring agents (see Algorithm 1). Performing physical layer authentication by comparing the physical fingerprints of received signals, legitimate agents can detect and isolate malicious agents that attempt spoofing attacks. Our algorithm achieves resilient consensus despite an arbitrary number of spoofed agents in the network. Our main contributions in this brief are as follows.

1) When fingerprint signals are stochastic, we quantify the worst case misclassification probability using distributionally robust Chebyshev bounds computed via

[1] Also known as a "Sybil" attack [3].

Manuscript received October 15, 2019; revised April 28, 2020 and September 17, 2020; accepted February 2, 2021. Manuscript received in final form March 2, 2021. This work was supported by the National Science Foundation under Grant CNS-1566127. Recommended by Associate Editor M. Oishi. *(Corresponding author: Tyler Summers.)*

Venkatraman Renganathan and Tyler Summers are with the Department of Mechanical Engineering, The University of Texas at Dallas, Richardson, TX 75080 USA (e-mail: vrengana@utdallas.edu; tyler.summers@utdallas.edu).

Kaveh Fathian is with the Aerospace Controls Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: kavehf@mit.edu).

Sleiman Safaoui is with the Department of Electrical Engineering, The University of Texas at Dallas, Richardson, TX 75080 USA (e-mail: sleiman.safaoui@utdallas.edu).

This article has supplementary material provided by the authors and color versions of one or more figures available at https://doi.org/10.1109/TCST.2021.3063924.

Digital Object Identifier 10.1109/TCST.2021.3063924

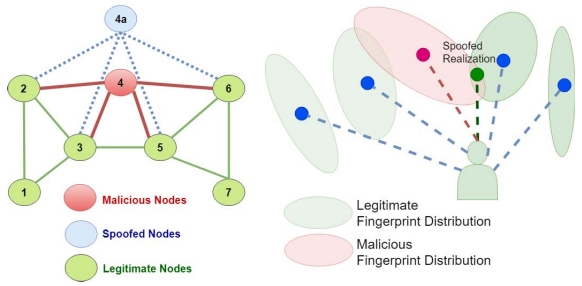

Fig. 1. Network under spoofing attack and a typical neighbors fingerprint distributions being compared by a legitimate robot.

semidefinite programming (see Theorem 2). We model true distributions of the fingerprints as unknown but belonging to a moment-based ambiguity sets.

2) We demonstrate that a malicious node spoofing at least F identities is sufficient to compromise a robust graph resulting in the failure of W-MSR strategy (see Theorem 1).
3) We perform numerical simulations and experiments using Sphero robots swarms to illustrate the effectiveness of the proposed methods.

Our framework is applicable to a variety of problems involving multirobot systems coordinating via wireless communication.

The rest of the brief is organized as follows. Section II formulates a model for spoofing attacks in multirobot networks and presents the attack detection technique using physical fingerprint analysis. Section III proposes the extension of W-MSR for spoof resiliency. The misclassification probability in a stochastic setting is addressed in Section IV, where a semidefinite programming formulation is proposed to arrive at distributionally robust Chebyshev bounds. Numerical simulations results are then presented in Section V followed by experimental results on a Sphero rolling robot swarm in Section VI. Finally, the brief is concluded in Section VII.

II. Spoofing Attacks in Multirobot Networks

A. Spoofing Attack—Network Model

Fig. 1 illustrates an example of a spoofing attack that we aim to address. We model the network with an undirected graph $\mathcal{G}$ comprising a node set $\mathcal{V}$ representing m agents and edge set $\mathcal{E}[t] \subset \mathcal{V} \times \mathcal{V}$ representing a set of (possibly time-varying) communication links amongst the agents. The node set is partitioned into two disjoint subsets $\mathcal{V} = S_l \cup S_a$. The set S_l represents the set of legitimate agents. A malicious agent attempts to disrupt the network by communicating subversive information to neighboring agents and *may*, in addition, attempt to perform a spoofing attack by creating multiple nonexistent identities. Thus, the set of adversaries S_a is composed of both malicious and spoofed agents so that $S_a = S_m \cup S_s$, where S_m denotes malicious agents and S_s denotes the agents spoofed by S_m. An upper bound of F number of malicious agents is assumed, whereas an arbitrary number of agents could be spoofed.

B. Consensus Dynamics—Update Model

We associate with each node $i \in \mathcal{V}$, a state $x_i[t] \in \mathbb{R}$ at time $t \in \mathbb{Z}_{\geq 0}$. The state may represent a position or some quantity to be estimated or optimized, depending on the application context. In order to achieve some objective, the nodes interact synchronously by exchanging their state value with neighbors in the network. Let the set of *inclusive neighbors* be defined as $\mathcal{J}_i[t] = \mathcal{N}_i[t] \cup \{i\}$, where $\mathcal{N}_i[t] = \{j \in \mathcal{V} : (j,i) \in \mathcal{E}[t]\}$ is the neighbor set of agent i at time t, whose states are available to agent i via communication links. Each legitimate node updates its own state over time based on its current state and the state of neighboring agents according to a prescribed rule of the form

$$x_i[t+1] = f_i\big(x_j[t]\big), \quad j \in \mathcal{J}_i[t], \ i \in S_l. \tag{1}$$

The degree of i is denoted as $d_i[t] = |\mathcal{N}_i[t]|$, and every node is assumed to have access to its own state at time t.

Definition 1: A node $i \in \mathcal{V}$ is said to be malicious if it sends $x_i[t]$ to all of its neighbors at each time step but applies some other function $f_i'(\cdot)$ at some time step [6].

Remark: We limit our discussion to malicious threat model, though our approach extends to the Byzantine model[2] as well.

C. Attack Detection Using Physical Fingerprint Analysis

We imagine a scenario in which the agents in the network communicate amongst themselves using a wireless communication protocol. We assume a complex multipath environment where a transmitted signal is scattered off of walls and objects, manifesting themselves as measurable peaks in fingerprints and, thereby, contributing significantly to its uniqueness. Physical properties of the received wireless signal profiles are leveraged to detect the spoofing attack. The physical fingerprint of an agent j received by agent i at time t, as described in [12], is modeled by a p-dimensional feature vector, $\mathcal{F}_i^j[t] \in \mathbb{R}^p$, containing physical signal properties, such as angle-of-arrival, time-of-arrival, and other features. These features can be measured using a synthetic aperture radar (SAR) and further can be processed using a well-studied signal processing algorithm called multiple signal classification (MUSIC) to generate spatial fingerprint corresponding to each neighboring agent. We assume that, when a malicious agent spoofs, it reports another identity through the noisy wireless communication channel where the noise can be modeled by associating a probability distribution with received signal fingerprints. Hence, fingerprints of the spoofed agents are realized from the same distribution of the spoofing agent. This will be discussed in detail in Section IV. Based on the received signal fingerprints of pairs of neighboring agents, we define a *similarity metric*

$$\gamma_i^{jk} = \frac{1}{1 + \left\| \mathcal{F}_i^j - \mathcal{F}_i^k \right\|}, \quad j,k \in \mathcal{N}_i \tag{2}$$

which quantifies how similar the fingerprint of neighboring agent j is to that of neighboring agent k, as received by agent i.

[2]A node $i \in \mathcal{V}$ is said to be **Byzantine** if it sends different values to different neighbors at some time step and if it applies some other update rule $f_i'(\cdot)$, at some time step.

Gill *et al.* [12] modeled the fingerprint by a directional signal strength profile that depends on wireless signal wavelengths, distances and relative angles between directional antennae, multiple possible signal paths, and random channel properties with the additive Gaussian noise. Our development is inspired by their approach. Agent i computes these similarity metrics for each neighbor pair. From these similarity metrics, a *confidence weight* $\alpha_i^j \in [0,1]$ can be associated with neighboring agent j, which should be close to 1 for legitimate neighbors and close to 0 for spoofed and spoofing neighbors. For example, in a deterministic setting, the confidence weights for neighbors j and k are 0 if the neighbor j has the same fingerprint as neighbor k and 1 otherwise. For the neighbor $j \in \mathcal{N}_i$, we define the confidence weight associated with agent i as

$$\alpha_i^j = \prod_{k \in \mathcal{N}_i,\, j \neq k} \left(1 - \gamma_i^{jk}\right). \tag{3}$$

D. Choosing a Spoofing Threshold

In a stochastic setting, we define a spoof detection threshold $\omega \in [0,1]$. Since robots have physical extent, there is a nonzero minimum distance between the sensors or receivers located in each robot that is used for discriminating received signals. We assume that this minimum distance translates to a corresponding minimum distance $\mathcal{F}_{\min}$ in feature vector space, where the pairwise comparisons are made, so that

$$\left\| \mathcal{F}_i^j - \mathcal{F}_i^k \right\| \geq \mathcal{F}_{\min} \quad \forall j,k \in \mathcal{N}_i, \quad j \neq k. \tag{4}$$

This suggests a threshold for robustly classifying neighboring agents whose fingerprints satisfy this bound as malicious or spoofed. Specifically, the similarity metrics and confidence weights for legitimate neighbors then satisfy

$$\gamma_i^{jk} \leq \frac{1}{\mathcal{F}_{\min}} \implies \alpha_i^j \geq \left(1 - \frac{1}{\mathcal{F}_{\min}}\right)^{|\mathcal{N}_i|}. \tag{5}$$

It follows that the threshold:

$$\omega = \left(1 - \frac{1}{\mathcal{F}_{\min}}\right)^{\xi}, \ \xi = \max\{|\mathcal{N}_i|\}, \quad i = 1,2,\ldots,|\mathcal{V}| \tag{6}$$

correctly discriminates between legitimate and malicious or spoofed neighbors, assuming that the fingerprints satisfy (4). However, fingerprints of received signals may be stochastic and not easily bounded resulting in a possibility of misclassifying a malicious neighbor as legitimate. If the likelihood that the physical fingerprints of two neighbors are different is below the threshold, the neighbors are classified as spoofed or spoofing agents, and otherwise, they are classified as legitimate; for example

$$\begin{aligned} g\left(\alpha_i^j\right) \leq \omega \quad &\Rightarrow \quad j \text{ is spoofed or spoofing} \\ g\left(\alpha_i^j\right) > \omega \quad &\Rightarrow \quad j \text{ is legitimate} \end{aligned} \tag{7}$$

where $g(\cdot)$ is a prescribed detection function. In the stochastic settings, $g(\cdot)$ and ω could be selected based on an assumed model for the probability distributions of the fingerprints and associated bounds on misclassification probability. To recover the deterministic case, we can set $\omega = 0$ and $g(x) = x$.

III. Design of a Spoof Resilient Coordination Algorithm

In this section, we describe a coordination algorithm that is resilient to anonymous malicious agents who share adversarial state values and *may* also attempt to spoof nonexistent agents who also share adversarial state values. Since malicious agents do not *all* necessarily attempt to spoof, we build upon recent work on resilient consensus algorithms that do not handle spoofing, and resiliency is achieved by effectively designing and exploiting redundancy in the communication graph.

A. Resilient Asymptotic Consensus

Let $x_M[t]$ and $x_m[t]$ denote the maximum and minimum values of the legitimate nodes at time t, respectively. The legitimate agents in the network are said to achieve *resilient asymptotic consensus* [7] in the presence of a particular threat model if for any initial conditions it holds.

1) $\exists\, \mathrm{L} \in \mathbb{R}$ such that $\lim_{t\to\infty} x_i[t] = L \forall i \in S_l$.
2) The interval $[x_m[0], x_M[0]]$ is an invariant set (i.e., the legitimate values remain in the interval $\forall t$).

Resilient asymptotic consensus implies that $L \in [x_m[0], x_M[0]]$, despite the presence of misbehaving nodes given a particular threat model and scope of the threat. We now review the resilient graph properties and an existing resilient consensus algorithm called W-MSR, as described in [6].

Definition 2: Given a graph $\mathcal{D}$ and a nonempty subset of nodes $\mathcal{S}$, we say that $\mathcal{S}$ is an (r,s)-reachable set if there are at least s nodes in $\mathcal{S}$, each of which has at least r neighbors outside of $\mathcal{S}$, where $r, s \in \mathbb{Z}_{\geq 0}$, i.e., given $\mathcal{X}_{\mathcal{S}}^r = \{i \in \mathcal{S} : |\mathcal{N}_i \backslash \mathcal{S}| \geq r\}$, then $|\mathcal{X}_{\mathcal{S}}^r| \geq s$.

Definition 3: A graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ on n nodes ($n \geq 2$) is *(r, s)-robust*; for nonnegative integers, $r \in \mathbb{Z}_{\geq 0}$, $1 \leq s \leq n$, if, for every pair of nonempty, disjoint subsets $\mathcal{S}_1$ and $\mathcal{S}_2$ of $\mathcal{V}$, at least one of the followings holds (recall $\mathcal{X}_{\mathcal{S}_k}^r = \{i \in \mathcal{S}_k : |\mathcal{N}_i \backslash \mathcal{S}_k| \geq r\}$ for k = {1, 2}).

1) $|\mathcal{X}_{\mathcal{S}_1}^r| = |\mathcal{S}_1|$.
2) $|\mathcal{X}_{\mathcal{S}_2}^r| = |\mathcal{S}_2|$.
3) $|\mathcal{X}_{\mathcal{S}_1}^r| + |\mathcal{X}_{\mathcal{S}_2}^r| \geq s$.

The (r,s)-robustness property introduces information redundancy by specifying a minimum number of nodes that are sufficiently influenced from outside of their set. Furthermore, (r,s)-robust graph is also $(\hat{r}, \hat{s})$-robust with $\hat{r} \leq r, \hat{s} \leq s$, but not vice versa [6].

B. W-MSR Algorithm

At every time t, each legitimate node i obtains the values of other nodes in its neighborhood. There are at most F total malicious nodes in the network, and some of node i's neighbors may misbehave; however, node i is unsure of which neighbors may be compromised. To ensure that node i updates its state in a safe manner, we consider a protocol where each node removes the extreme values with respect to its own value. We denote by $\mathcal{R}_i[t]$ the set of nodes whose values were removed by legitimate node i during the comparison process at

time t. Each legitimate node i then applies the W-MSR update

$$x_i[t+1] = \sum_{j \in \mathcal{J}_i[t] \backslash \mathcal{R}[t]} w_{ij}[t] x_j[t]$$
$$w_{ij}[t] = \frac{1}{1 + d_i[t] - |\mathcal{R}_i[t]|}, \quad j \in \mathcal{J}_i[t] \backslash \mathcal{R}[t] \tag{8}$$

where $w_{ij}[t]$[3] is the weight associated with node j's value by node i at time step t. Assuming that there are F total malicious agents and no spoofed agents, $(F+1, F+1)$-robustness is a necessary and sufficient condition for the normal nodes to achieve resilient asymptotic consensus [6] under the W-MSR update protocol [14] specified in (8). We first observe that the W-MSR protocol is easily compromised by a spoofing attack and subsequently present our spoof resilient adaptation of W-MSR. It is clear that any malicious node can create a sufficiently large number of spoofed nodes so that a neighboring legitimate node is forced to use the values of a spoofed node even after applying the W-MSR update rule. Here, we show that spoofing at least F identities is sufficient for a malicious node to cause W-MSR to fail (though fewer or even a single node can cause failure in certain cases).

Theorem 1: Consider an undirected network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ represents the set of m agents and $\mathcal{E}$ represents the set of communication links between them. Assume that the network is $(F+1, F+1)$-robust and not $(F+2, F+2)$-robust, where F is the upper bound on the total number of malicious agents in the network. Let all the legitimate agents use the W-MSR update rule. Then, there exists a location in the network for the malicious node where spoofing at least F number of identities is sufficient to cause the failure of the legitimate agents to achieve resilient asymptotic consensus.

Proof: Since $\mathcal{G}$ is not $(F+2, F+2)$-robust, then there are nonempty, disjoint subsets $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ such that none of the three conditions 1)–3) holds. Take one such pair $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ such that $|\mathcal{X}_{\mathcal{S}_1}^{F+2}| \leq |\mathcal{S}_1| - 1$, $|\mathcal{X}_{\mathcal{S}_2}^{F+2}| \leq |\mathcal{S}_2| - 1$, and $|\mathcal{X}_{\mathcal{S}_1}^{F+2}| + |\mathcal{X}_{\mathcal{S}_2}^{F+2}| \leq F+1$. Consider the case that $|\mathcal{X}_{\mathcal{S}_1}^{F+2} \cup \mathcal{X}_{\mathcal{S}_2}^{F+2}| = F+1$, and out of that, suppose that F nodes are malicious. This indicates that each of $\mathcal{S}_1$ and $\mathcal{S}_2$ has at least one legitimate node, each with no more than $F+1$ neighbors outside of its respective set. Denote the sets of such nodes as $\mathcal{V}_{\mathcal{S}_1}$ and $\mathcal{V}_{\mathcal{S}_2}$, respectively. Let one of the malicious nodes in $\mathcal{X}_{\mathcal{S}_1}^{F+2}$, say $\mathcal{V}_m$, spoof F nodes denoted by $\{\mathcal{V}_m^{s_l}\}_{l=1}^F$ with edges to the same neighbors of $\mathcal{V}_m$. Now, consider the new graph $\hat{\mathcal{G}} = (\hat{\mathcal{V}}, \hat{\mathcal{E}})$ with $\hat{\mathcal{V}} = \mathcal{V} \cup \{\{\mathcal{V}_m^{s_l}\}_{l=1}^F\}$ and $\hat{\mathcal{E}} = \mathcal{E} \cup \{(i,j) | i \in \mathcal{N}_{\mathcal{V}_m}, j \in \{\mathcal{V}_m^{s_l}\}_{l=1}^F\}$. Define the disjoint subsets $\hat{\mathcal{S}}_1, \hat{\mathcal{S}}_2 \subset \hat{\mathcal{V}}$ for the new graph as $\hat{\mathcal{S}}_1 = \mathcal{S}_1 \cup \{\{\mathcal{V}_m^{s_l}\}_{l=1}^F\}$ and $\hat{\mathcal{S}}_2 = \mathcal{S}_2$, respectively, so that there are at least $F+1$ adversarial nodes inside $\hat{\mathcal{S}}_1$: the F spoofed nodes, the malicious spoofing node $\mathcal{V}_m$, and other malicious nodes if any. Denote the sets $\mathcal{V}_{\mathcal{S}_1}$ and $\mathcal{V}_{\mathcal{S}_2}$ after spoofing as $\mathcal{V}_{\hat{\mathcal{S}}_1}$ and $\mathcal{V}_{\hat{\mathcal{S}}_2}$, respectively.

Let $0 < a < b < c < d < e$, the initial values of legitimate nodes in $\hat{\mathcal{S}}_1$ and $\hat{\mathcal{S}}_2$, be b and d, respectively, and the initial values of the rest of the legitimate nodes in the network be c. Furthermore, let all malicious and spoofed nodes in $\hat{\mathcal{S}}_1, \hat{\mathcal{S}}_2$ have constant values a and e, respectively, across all time steps. The W-MSR algorithm with parameter F discards at most the F highest and F lowest values and then applies a convex combination update rule with each weight lower bounded by κ. Applying the convex combination of values, nodes in $\mathcal{V}_{\hat{\mathcal{S}}_2}$ with values $d = x_M[t]$ will be monotonic and bounded functions of time t (by Lemma 1 of [7]). However, when nodes in $\mathcal{V}_{\hat{\mathcal{S}}_1}$ with values $b = x_m[t]$ apply W-MSR, they discard at most the highest F values out of $F+1$ neighbor values from outside $\hat{\mathcal{S}}_1$, and similar to discarding the lowest F values, they remove F of the $F+1$ or more a values in $\hat{\mathcal{S}}_1$ leaving at least one a value in update consideration. The convex combination-based W-MSR updates with the remaining neighboring values, which includes at least one a value from the adversarial node; thus, $x_m[t]$ will no longer be a monotone but bounded function of time t. Therefore, $\mathcal{V}_{\hat{\mathcal{S}}_1}$ and $\mathcal{V}_{\hat{\mathcal{S}}_2}$, and hence, all legitimate nodes reach the consensus value of $a \notin [b, d]$, thereby resulting in failure of the W-MSR protocol. Analogously, if the spoofing is happened in $\mathcal{S}_2$, the attained consensus would be $e \notin [b, d]$. □

C. Spoof Resilient W-MSR Algorithm

Our spoof resilient adaptation of the W-MSR algorithm is summarized in Algorithm 1. Based on a pairwise comparison of physical fingerprints of signals received from neighboring agents and associated confidence weights, spoofed agents in the network are identified. Achieving resiliency then involves removing the identified spoofed and spoofing agents from the state update if their confidence weight is at most equal to the spoofing threshold ω. Thus, in a stochastic setting, a spoofing threshold can be employed, as explained in Algorithm 1. To conclude this section, we show that Algorithm 1 exactly achieves resiliency to an arbitrary number of spoofed agents in the deterministic case where fingerprint noise is zero, with an appropriate selection of the spoofing threshold.

Lemma 1: Consider an undirected network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ represents the set of agents and $\mathcal{E}$ represents the set of communication links between them, and consider the Spoof-Resilient W-MSR described in Algorithm 1. Suppose that the physical fingerprints of each agent are deterministic, the spoofing threshold ω is set to zero, and the detection function is set to $g(x) = x$. Suppose further that the network is $(F+1, F+1)$-robust, assuming an upper bound of F total malicious agents in the network, *some* of which may spoof. Then, the network achieves resilient asymptotic consensus under Algorithm 1 in the presence of any spoofing attack.

Proof: Since the physical fingerprints are deterministic and the malicious agent reports multiple identities from the same location, then all fingerprints of signals of spoofed agents are identical to that of the spoofing agent. Thus, any spoofed and spoofing agents (say $\hat{F}$ in total) are exactly detected using zero thresholds and identity detection function and ignored from the state update in line 14 of the Algorithm 1. Furthermore, applying W-MSR update (with parameter $F_{\text{new}} = F - \hat{F}$) after the identification of spoofing agents guarantees resiliency in the presence of remaining F_{new} malicious nodes

[3]For all $t \in \mathbb{Z}_{\geq 0}$, weights satisfy $w_{ij}[t] \geq \kappa$, where $0 < \kappa < 1, \kappa \in \mathbb{R}$, $\sum_{j=1}^n w_{ij}[t] = 1 \forall i \in S_l$, and $w_{ij}[t] = 0 \forall j \notin \mathcal{J}_i[t], i \in S_l$ [7].

Algorithm 1 Spoof Resilient W-MSR (SR-W-MSR)

# Input: spoofing threshold ω, convergence threshold ϵ, initial time t = 0, initial states x[0], received signal fingerprints $\mathcal{F}_i^j$ for each agent at each time
while $\|x[t+1] - x[t]\| \geq \epsilon$ **do**
  # Iterate through all legitimate nodes
  for each $i \in S_l$ **do**
    for each $j \in \mathcal{N}_i$ **do**
      $\alpha_i^j \leftarrow 1$
      for each $k \in \mathcal{N}_i \backslash \{j\}$ **do**
        # Pairwise comparison of neighbors
        $\gamma_i^{jk} \leftarrow \frac{1}{1+\|\mathcal{F}_i^j - \mathcal{F}_i^k\|}$
        $\alpha_i^j \leftarrow \alpha_i^j(1-\gamma_i^{jk})$
      if $g(\alpha_i^j) \leq \omega$ **then**
        # Spoof attack is detected
        $\mathcal{R}[t] \leftarrow \mathcal{R}[t] \cup \{j\}$
    Update $x_i[t+1]$ using W-MSR rule (8)
  $t \leftarrow t+1$

who do not spoof but may apply adversarial updates to their state. □

IV. Quantifying Misclassification Probabilities in Spoofing Detection

In our proposed physical layer authentication technique, each legitimate robot associates a fingerprint with each wireless communication signal received from neighboring agents. Following the fingerprint construction procedure explained in Section IV of [12], we interpret the fingerprint as a vector of signal parameters, such as received signal strength and relative bearing. Since communication channels are inherently noisy, the pairwise comparison might result in misclassification, as depicted in Fig. 1. By associating probability distributions with fingerprints of received signals, misclassification probabilities can be estimated based on assumptions about the distributions. In practice, the true fingerprint distributions are generally unknown, may be difficult to estimate, and depend on the environment configuration, obstacle positions, battery power levels, and even the robot positions. However, there are many applications where knowledge about the environment, onboard sensor characteristics and experimental data allow quantification of expected variations in fingerprint distributions. For example, the spacing of vehicle-to-vehicle communication equipment in cars driving along a highway is constrained by the physical extent of actual cars. Furthermore, inherent statistical variations in the transmitters and receivers can be estimated from received signal data.

These difficulties in estimating fingerprint distributions also motivate a distributionally robust approach [15], where, rather than assuming or estimating a particular fingerprint distribution, we instead assume knowledge or estimation of only the first two moments. Interpreting $\mathcal{F}_{\min}$ as a threshold to make classification decisions may result in misclassification due to possibly unbounded fingerprint distributions. Here, we compute the worst case probabilities of misclassification over the set of fingerprint distributions with given mean and covariance. These worst case estimates can be interpreted as generalized Chebyshev bounds [16], which extends classical moment-based Chebyshev inequalities to vector-valued random variables.

Let there be n neighbors in total for a given signal receiving robot that consists of both legitimate and bad neighbors. The set of bad neighbors consists of both malicious neighbors, and any arbitrary number of spoofed identities emulated by each malicious neighbor. Let us denote the set of legitimate neighbors by $\mathcal{L} = \{1, 2, \ldots, n_l\}$ and the set of malicious neighbors by $\mathcal{M} = \{1, 2, \ldots, n_m\}$. Also, we refer by X_m^s the spoofed agent emulated by malicious agent m. Finally, we collectively refer to the set of all spoofed entities X_m^s emulated by agents $m \in \mathcal{M}$ by $\mathcal{S} = \{1, 2, \ldots, n_s\}$, where $n_s = \sum_{m\in\mathcal{M}} n_m^s$, and n_m^s refers to the number of spoofed entities emulated by the malicious agent, $m \in \mathcal{M}$. For every $l \in \mathcal{L}$ and $m \in \mathcal{M}$, we denote by $X_l \sim \mathbf{P}_l(\mu_l, \Sigma_l)$ and $X_m \sim \mathbf{P}_m(\mu_m, \Sigma_m)$ the fingerprints and associated probability distributions of the neighbors received by a robot, where X_l corresponds to a legitimate robot and X_m to a malicious one. We assume that the means $\mu_l, \mu_m \in \mathbb{R}^p$ and covariance matrices $\Sigma_l, \Sigma_m \in \mathbf{S}^p$ are known or can be estimated from received signal data or sensor hardware datasheets (with $\mathbf{S}^p$ denoting the set of symmetric $p \times p$ matrices) but the true distributions $\mathbf{P}_l$ and $\mathbf{P}_m$ are unknown and belong to a moment-based ambiguity sets $\mathcal{P}_l$ and $\mathcal{P}_m$, respectively, which are defined as follows:

$$\mathcal{P}_l = \left\{\mathbf{P}_l \,|\, \mathbf{E}[X_l] = \mu_l, \mathbf{E}\left[(X_l - \mu_l)(X_l - \mu_l)^\top\right] = \Sigma_l\right\} \quad (9)$$

$$\mathcal{P}_m = \{\mathbf{P}_m \,|\, \mathbf{E}[X_m] = \mu_m, \mathbf{E}\left[(X_m - \mu_m)(X_m - \mu_m)^\top\right] = \Sigma_m\}. \quad (10)$$

When a malicious agent $m \in \mathcal{M}$ attempts a spoofing attack, all true fingerprints distributions of X_m and that of the spoofed identity X_m^s, respectively, fall under the same ambiguity set $\mathcal{P}_m$ sharing the same moments data. That is, $X_m \sim \mathbf{P}_{m1}(\mu_m, \Sigma_m)$, $X_m^s \sim \mathbf{P}_{m2}(\mu_m, \Sigma_m)$ with $\mathbf{P}_{m1}, \mathbf{P}_{m2} \in \mathcal{P}_m$, and also $\mathbf{P}_{m1}, \mathbf{P}_{m2}$ need not denote the same distribution. Hence, every legitimate agent, $l \in \mathcal{L}$, and malicious agent, $m \in \mathcal{M}$, will have its own moment-based ambiguity set $\mathcal{P}_l$ and $\mathcal{P}_m$, respectively, where its true distribution lies satisfying the corresponding moments data. Out of the $\binom{n}{2}$ pairwise comparisons by the signal receiving robot, we enumerate the possible events that will result in misclassification. A misclassification event occurs when the following holds.

1) A legitimate neighbor's fingerprint realization is closer than $\mathcal{F}_{\min}$ to any malicious neighbor's fingerprint realization.
2) A legitimate neighbor's fingerprint realization is closer than $\mathcal{F}_{\min}$ to a spoofed fingerprint realization of any malicious neighbor.
3) A legitimate neighbor's fingerprint realization is closer to a malicious neighbor's fingerprint realization than the corresponding malicious neighbor's realization is to one of its own spoofed realizations.

4) A legitimate neighbor's fingerprint realization is closer to a spoofed realization than the corresponding malicious neighbor's realization that emulated the spoofed entity.

To model the above scenarios for particular pairs or triples of legitimate, malicious, and spoofed agents $l \in \mathcal{L}, m \in \mathcal{M}$, and $s \in \mathcal{S}$, we define the events

$$\begin{aligned}
\mathcal{C}_1^{l,m} &= \left\{X_l, X_m \,:\, \|X_l - X_m\|^2 \le \mathcal{F}_{min}\right\} \\
\mathcal{C}_2^{l,s} &= \left\{X_l, X_m^s \,:\, \|X_l - X_m^s\|^2 \le \mathcal{F}_{min}\right\} \\
\mathcal{C}_3^{l,m,s} &= \left\{X_l, X_m, X_m^s \,:\, \right. \\
&\qquad \left. \|X_l - X_m\|^2 \le \|X_m - X_m^s\|^2\right\} \\
\mathcal{C}_4^{l,m,s} &= \left\{X_l, X_m, X_m^s \,:\, \right. \\
&\qquad \left. \|X_l - X_m^s\|^2 \le \|X_m - X_m^s\|^2\right\} \qquad (11)
\end{aligned}$$

and denote corresponding families of these events as

$$\begin{aligned}
\mathcal{C}_1 &= \left\{\mathcal{C}_1^{l,m} \,:\, l \in \mathcal{L}, m \in \mathcal{M}\right\} \\
\mathcal{C}_2 &= \left\{\mathcal{C}_2^{l,s} \,:\, l \in \mathcal{L}, s \in \mathcal{S}\right\} \\
\mathcal{C}_3 &= \left\{\mathcal{C}_3^{l,m,s} \,:\, l \in \mathcal{L}, m \in \mathcal{M}, s \in \mathcal{S}\right\} \\
\mathcal{C}_4 &= \left\{\mathcal{C}_4^{l,m,s} \,:\, l \in \mathcal{L}, m \in \mathcal{M}, s \in \mathcal{S}\right\}. \qquad (12)
\end{aligned}$$

Since the fingerprints are independent, the joint random variable denoted by $X = [X_l, X_m, X_m^s]$ has mean and covariance

$$\mu = \begin{pmatrix} \mu_l \\ \mu_m \\ \mu_s \end{pmatrix}, \quad \Sigma = \begin{pmatrix} \Sigma_l & 0 & 0 \\ 0 & \Sigma_m & 0 \\ 0 & 0 & \Sigma_s \end{pmatrix}. \qquad (13)$$

For $k = 1, \ldots, 4$ and $i = 1, \ldots, |\mathcal{C}_k|$, respectively, the probability of misclassification event i is given by

$$\sup\left\{\mathbf{P}\left(X \in \mathcal{C}_k^i\right)\right\} = 1 - \inf\left\{\mathbf{P}\left(X \notin \mathcal{C}_k^i\right)\right\} \qquad (14)$$

where the supremum and infimum are taken over the set of probability distributions on $\mathbb{R}^p$ with the given mean and covariance. Utilizing generalized Chebyshev bounds from [16], the misclassification probabilities can be readily computed by solving a (convex) semidefinite programming problem. Though the events are generally independent within the same family, they are generally dependent between families. Thus, we address them separately here rather than jointly.

A. Computing Misclassification Probability via Semidefinite Programming

We emphasize that the proposed spoof resilient W-MSR algorithm (see Algorithm 1) can be implemented without knowledge of the fingerprint distributions. Knowledge of fingerprint distributions is only required to estimate bounds on misclassification probabilities, which then informs the choice of threshold in the algorithm. The first step in computing the misclassification probability for a particular pair of legitimate and malicious agents, or triple of legitimate, malicious, and spoofed agents, is to form the neighbor pair combination matrix $A_{\mathcal{C}_k^i}$ corresponding to an event i in the family $\mathcal{C}_k$. Assume a total of $n = 3$ neighbors with $n_l = n_m = n_s = 1$, and thus, $X = [X_l, X_m, X_m^s]^\top$. Now, for $i = 1, \ldots, |\mathcal{C}_k|$, we have

$$A_{\mathcal{C}_1^i} = \begin{pmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \otimes I_p, A_{\mathcal{C}_2^i} = \begin{pmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{pmatrix} \otimes I_p \qquad (15)$$

$$A_{\mathcal{C}_3^i} = \begin{pmatrix} -1 & 1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix} \otimes I_p,$$

$$A_{\mathcal{C}_4^i} = \begin{pmatrix} 1 & 0 & -1 \\ 0 & -1 & 1 \\ -1 & 1 & 0 \end{pmatrix} \otimes I_p. \qquad (16)$$

We now have the following result for computing distributionally robust misclassification probabilities using semidefinite programming.

Theorem 2: Consider the joint random fingerprint variable $X = [X_l, X_m, X_s]^\top$ with mean $\mathbf{E}[X] = \mu$ and covariance $\mathbf{E}[XX^\top] = \Sigma$. Now, consider an event i in the family $\mathcal{C}_k$, $k \in \{1, 2, 3, 4\}$, with $A_{\mathcal{C}_k^i}$ formed using (15) or (16). Let

$$\kappa = \begin{cases} \lambda \mathcal{F}_{\min}, & \text{if } k = 1, 2 \\ 0, & \text{otherwise.} \end{cases}$$

The worst case probability of misclassification, $\sup\{\mathbf{P}(X \in \mathcal{C}_k^i)\}$ over the set of all distributions of X with mean μ and covariance Σ is given by the optimal value of the following semidefinite program:

$$\begin{aligned}
&\max \quad \lambda \\
&\text{s.t.} \quad \mathbf{tr}\left(A_{\mathcal{C}_k^i} Z\right) \le \kappa \\
&\qquad \begin{pmatrix} Z & z \\ z^\top & \lambda \end{pmatrix} \preceq \begin{pmatrix} \Sigma + \mu^\top \mu & \mu \\ \mu^\top & 1 \end{pmatrix} \\
&\qquad \begin{pmatrix} Z & z \\ z^\top & \lambda \end{pmatrix} \succeq 0 \qquad (17)
\end{aligned}$$

with variables $\lambda \in \mathbb{R}$, $Z \in \mathbf{S}^{3p}$, and $z \in \mathbb{R}^{3p}$.

Proof: Here, we demonstrate a proof for a misclassification event i from the family $\mathcal{C}_3$. The proofs corresponding to events in the other families are entirely analogous. In this case, the event involves one legitimate neighbor X_l, one malicious neighbor X_m, and one spoofed identity X_m^s. Let us denote the joint random variable as $X = [X_l, X_m, X_m^s]^\top$. Let $A = (X_l - X_m)$, $B = (X_m - X_m^s)$. Now, the smallest probability of *not* misclassifying $\inf \mathbf{P}(X \notin \mathcal{C}_3^i)$ is given by

$$\inf \mathbf{P}\left(\|A\|^2 \ge \|B\|^2\right) = \inf \mathbf{P}\left(X^\top A_{\mathcal{C}_3^i} X \le 0\right).$$

Then, applying a generalized Chebyshev inequality (from the main result in Section II of [16]) and combining with (14) yields the result. □

B. Illustrative Example

For illustration purposes, consider the robots move in a highly noisy setting in $\mathbb{R}^2$. The receiving robot has sensors with a positive variance that allows discrimination of received signal strength in the radial direction and bearing measurement in the orthogonal direction, which are used to estimate the

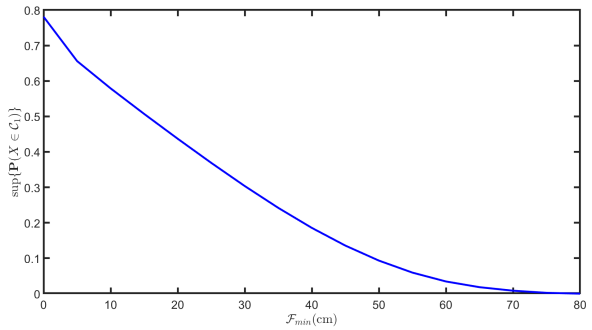

Fig. 2. Misclassification probability with respect to events in the family $\mathcal{C}_1$ decreases as $\mathcal{F}_{\min}$ is increased.

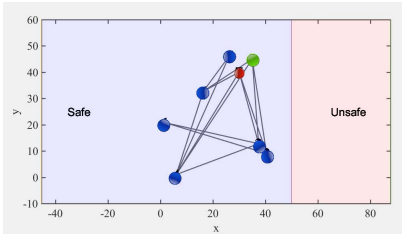

Fig. 3. (2, 2)-robust graph with one robot being malicious (red) and it spoofs a nonexistent robot (green).

positions of neighboring robots, which then serves as the fingerprint estimate. Suppose that the receiving robot is located at the origin, and the neighbors are located randomly around it. The semidefinite program (17) was solved for varying threshold values $\mathcal{F}_{\min}$ to compute the worst case misclassification probability with respect to the events in the family $\mathcal{C}_1$. As the robots get more physically separated than the minimum separation ($\mathcal{F}_{\min}$), the misclassification probability decreases, as shown in Fig. 2.

V. Numerical Simulations

We now illustrate our spoof resilient W-MSR algorithm using the 7-node (2, 2)-robust network with six legitimate agents and one malicious agent who spoofs one agent and sends messages to all his neighbors using both identities with an intention to move the robot formation to unsafe region, as shown in Fig. 3. The objective of the network is to form a hexagonal formation and remain in the safe region, which can be expressed by introducing a constant bias in the consensus update equations. Specifically, at every iteration, the desired position of robot, $i = 1, \ldots, 6$, is computed via

$$x_i[t+1] = \sum_{j \in \mathcal{J}_i[t] \setminus \mathcal{R}[t]} w_{ij}[t]\big(x_j[t] - \bar{x}_j\big) + \bar{x}_i \qquad (18)$$

where $\bar{x}_i := [\sin(\theta_i), \cos(\theta_i)]^\top \in \mathbb{R}^2$ and $\theta_i := (2\pi(i-1))/6$ is a constant bias vector that is used to position each robot at a vertex of the hexagon. Note that, due to the addition and subtraction of the same bias vector for each agent, by defining $\tilde{x}_i := x_i - \bar{x}_i$, one can see from (18) that $\tilde{x}_i$ has the same dynamics as of (8). The six legitimate agents were given random initial states, as shown in Fig. 3. A constant bias of $\bar{x} = 5$ cm is added to the malicious robot's position in each dimension to obtain the spoofed robot's position. Now, consider the spoofed agent not being present. When a standard linear consensus protocol is employed to obtain a robot formation, it fails, as shown in Fig. 4(a), depicting that the protocol is not resilient against malicious attacks. When the W-MSR algorithm is used to obtain a robot formation, the legitimate robots are resilient against the malicious robot and achieve the desired formation in the safe region, as shown in Fig. 4(b), thus proving resiliency against malicious attacks, but without spoofing. Now, consider that agent 7 is malicious. When the malicious agent spoofs a single additional agent identity, the legitimate agents fail to achieve resiliency, and hence, they are pulled into an unsafe region by malicious robots, as shown in Fig. 4(c). This shows that spoofing attacks are capable of compromising graph robustness properties and, thereby, the network resiliency. Now, consider the same setting where the spoofing attack is simulated for the first ten time steps, and then, the algorithm switches to spoof resilient version guaranteeing spoof resilient formation in the safe region, as shown in Fig. 4(c), thereby emphasizing the need for earlier detection.

VI. Experimental Results

To demonstrate our approach, experiments are performed on a robotic platform using the Sphero rolling robot swarm. The communication among the robots is shown in Fig. 3, where each robot represents a node in the (2, 2)-robust network with six legitimate and one malicious agent.

A. Implementation Details

Our experimental setup consists of seven Sphero 2.0 robots, a Logitech C950 webcam, a Bluetooth-enabled smartphone, and a computer system that was equipped with Intel Core i7-6600U (four CPUs) at a 2.60-GHz processor. All routines executed during the experiment are implemented in MATLAB. We emphasize that, although this experimental implementation is centralized, the information that is made available to each robot is restricted according to the communication graph shown in Fig. 3. The webcam is set up to overlook a confined area, in which the Sphero robots are placed and allowed to move. The LEDs of the legitimate and malicious robots are set to emit blue and red colors, respectively. A color-based image segmentation routine, as explained in [17], is used to detect and track the robots in real time from the 640×480 images that are fetched from the webcam. The recovered coordinates from the camera data are used in a consensus-based formation control strategy to bring the robots to a hexagon formation. At every iteration, the desired position of legitimate robots is computed via (18). For demonstration purposes, a constant bias of $\bar{x} = 50$ cm is added to the malicious robot's position in each dimension to obtain the spoofed robot's position. Given the desired position $x_i[t+1]$, a low-level PID controller computes the required linear and angular velocity control commands that guide the robot to this location at the next iteration. These control commands are communicated to the robots via Bluetooth, as explained in [18]. The motion of the legitimate robots is controlled by the computer, and the malicious robot is controlled manually via a smartphone. A video

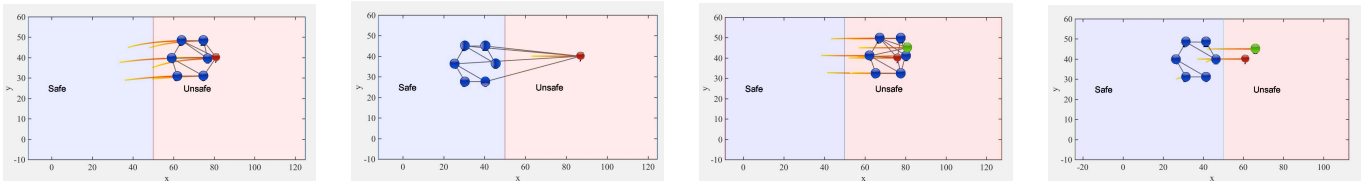

(a) (b) (c) (d)

Fig. 4. With the robots given random initial positions as in 3, the linear consensus protocol is not resilient against malicious attack, as shown in (a). The W-MSR protocol guarantees a resilient formation against malicious robots in a safe region, as shown in (b). Malicious robot spoofs nonexistent robot resulting in the robot formation being pulled into the unsafe region, as shown in (c). When the attack is detected, the malicious and spoofed robots are removed from the network, and spoof resilient W-MSR guarantees resiliency against spoofing attack, as shown in (d). Blue, red, and green represent legitimate, malicious, and spoofed robots, respectively. (a) Linear consensus protocol fails with one malicious robot. (b) W-MSR succeeds in keeping the formation in the safe region. (c) W-MSR algorithm fails under the spoofing attack. (d) Proposed SR-W-MSR algorithm achieves spoof resiliency.

supplement available at https://youtu.be/dcd0EexNnzE shows the experimental results, where we observe that the spoof resilient W-MSR algorithm successfully detects and isolates a spoofing attack in the experimental conditions.

VII. Conclusion

We proposed a spoof resilient consensus algorithm by extending a class of resilient consensus strategies, known as the W-MSR consensus, to provide resilience to malicious agents that may both adversely update state values and spoof nonexistent agent identities. The proposed algorithm using the physical fingerprint approach guarantees resiliency despite the presence of a certain number of malicious agents and an arbitrary number of spoofed agents in the network. As a novel contribution, a probabilistic spoof detection analysis is presented using a semidefinite programming technique to arrive at distributionally robust Chebyshev bounds for the probability of misclassification of robots. Experimental results using Sphero robot swarms and numerical simulations demonstrate the effectiveness of the proposed algorithm. Future research involves quantifying the worst case probability of misclassifications persisting overextended periods of time by considering fingerprints over multiple time periods.

References

[1] A. Cardenas, S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. Sastry, "Challenges for securing cyber physical systems," in *Proc. Workshop Future Directions Cyber-Phys. Syst. Secur.*, 2009, pp. 1–7.

[2] A. Banerjee, K. K. Venkatasubramanian, T. Mukherjee, and S. K. S. Gupta, "Ensuring safety, security, and sustainability of mission-critical cyber-physical systems," *Proc. IEEE*, vol. 100, no. 1, pp. 283–299, Jan. 2011.

[3] J. R. Douceur, "The sybil attack," in *Proc. Revised Papers 1st Int. Workshop Peer-to-Peer Syst.*, London, U.K.: Springer, 2002, pp. 251–260.

[4] M. L. Psiaki and T. E. Humphreys, "GNSS spoofing and detection," *Proc. IEEE*, vol. 104, no. 6, pp. 1258–1270, Jun. 2016.

[5] Y. Shoukry, P. Martin, P. Tabuada, and M. Srivastava, "Non-invasive spoofing attacks for anti-lock braking systems," in *Proc. 15th Int. Workshop Cryptograph. Hardw. Embedded Syst.*, Santa Barbara, CA, USA, 2013, pp. 55–72.

[6] H. Zhang, E. Fata, and S. Sundaram, "A notion of robustness in complex networks," *IEEE Trans. Control Netw. Syst.*, vol. 2, no. 3, pp. 310–320, Sep. 2015.

[7] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 4, pp. 766–781, Apr. 2013.

[8] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation via linear iterative strategies in the presence of malicious agents," *IEEE Trans. Autom. Control*, vol. 56, no. 7, pp. 1495–1508, Jul. 2011.

[9] F. Pasqualetti, A. Bicchi, and F. Bullo, "Consensus computation in unreliable networks: A system theoretic approach," *IEEE Trans. Autom. Control*, vol. 57, no. 1, pp. 90–104, Jan. 2012.

[10] L. Guerrero-Bonilla, A. Prorok, and V. Kumar, "Formations for resilient robot teams," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*. Singapore: IEEE-RAS, Jan. 2017, pp. 1–8.

[11] D. Saldana, A. Prorok, M. F. M. Campos, and V. Kumar, "Triangular networks for resilient formations," in *Proc. 13th Int. Symp. Distrib. Auton. Robotic Syst. (DARS)*, London, U.K.: IEEE-RAS, Mar. 2016, pp. 147–159.

[12] S. Gil, S. Kumar, M. Mazumder, D. Katabi, and D. Rus, "Guaranteeing spoof-resilient multi-robot networks," *Auton. Robots*, vol. 41, no. 6, pp. 1383–1400, 2017.

[13] V. Renganathan and T. Summers, "Spoof resilient coordination for distributed multi-robot systems," in *Proc. Int. Symp. Multi-Robot Multi-Agent Syst. (MRS)*, Dec. 2017, pp. 135–141.

[14] J. Usevitch and D. Panagou, "R-robustness and (R, S)-robustness of circulant graphs," in *Proc. IEEE 56th Annu. Conf. Decis. Control (CDC)*, Dec. 2017, pp. 4416–4421.

[15] W. Wiesemann, D. Kuhn, and M. Sim, "Distributionally robust convex optimization," *Oper. Res.*, vol. 62, no. 6, pp. 1358–1376, Dec. 2014.

[16] L. Vandenberghe, S. Boyd, and K. Comanor, "Generalized chebyshev bounds via semidefinite programming," *SIAM Rev.*, vol. 49, no. 1, pp. 52–64, Jan. 2007.

[17] P. Corke, *Robotics, Vision and Control: Fundamental Algorithms in MATLAB*, vol. 118. Berlin, Germany: Springer, 2017.

[18] D. Sethi and G. Campa. (2018). *Sphero Connectivity Package*. [Online]. Available: http://www.mathworks.com/matlabcentral/fileexchange/52481-spheroconnectivity-package