Title

The Impact of Complex and Informed Adversarial Behavior in Graphical Coordination Games

Authors

Paarporn, Keith
Canty, Brian
Brown, Philip N.
et al.

Publication Date

2020-10-03

The Impact of Complex and Informed Adversarial Behavior in Graphical Coordination Games

Keith Paarporn*, Brian Canty*, Philip N. Brown, Mahnoosh Alizadeh, and Jason R. Marden

***Abstract*—How does system-level information impact the ability of an adversary to degrade performance in a networked control system? How does the complexity of an adversary's strategy affect its ability to degrade performance? This paper focuses on these questions in the context of graphical coordination games where an adversary can influence a given fraction of the agents in the system, and the agents follow log-linear learning, a well-known distributed learning algorithm. Focusing on a class of homogeneous ring graphs of various connectivity, we begin by demonstrating that minimally connected ring graphs are the most susceptible to adversarial influence. We then proceed to characterize how both (i) the sophistication of the attack strategies (static vs dynamic) and (ii) the informational awareness about the network structure can be leveraged by an adversary to degrade system performance. Focusing on the set of adversarial policies that induce stochastically stable states, our findings demonstrate that the relative importance between sophistication and information changes depending on the the influencing power of the adversary. In particular, sophistication far outweighs informational awareness with regards to degrading system-level damage when the adversary's influence power is relatively weak. However, the opposite is true when an adversary's influence power is more substantial.**

I. Introduction

A networked system can be viewed as a collection of subsystems, each required to make local and independent decisions in response to available information. The information available to each subsystem could pertain to local environmental conditions or the behavior of a selected group of neighboring agents in the system; hence, the information available to one subsystem could be vastly different than the information available to other subsystems. Regardless of the specific problem domain and informational characteristics, the underlying goal is to derive agent control policies that ensure the emergent collective behavior is desirable with respect to a system-level performance metric.

A central focus of such systems is the design of networked control algorithms that provide strong guarantees on the quality of emergent outcomes. A networked control algorithm can be viewed as a decision-making rule that specifies how subsystems respond to local conditions. There are several noteworthy results in this domain ranging from consensus and flocking [2], [3], sensor allocation [4], [5], coordination of unmanned vehicles [6], and many others. A common theme in all of these works is the following: If all agents follow the prescribed decision-making rules, then the emergent behavior is both stable and desirable. In contrast to this work, here we seek to address whether such decision-making rules are robust to adversarial interventions.

While the decentralization associated with distributed architectures is undoubtedly appealing for a host of reasons, it is important to highlight that this also introduces vulnerabilities. In particular, the decision-making process of individual subsystems can potentially be influenced by adversarial actors in the system through corrupting or augmenting the information to the subsystems. Accordingly, in this paper we ask whether an adversary can exploit these interconnections to negatively influence the quality of the emergent collective behavior. Formal analysis of this interplay has emerged in recent years, often in the context of robust consensus, distributed optimization, and cyber-physical system security [7]–[9].

The focus of this paper is the susceptibility of a distributed algorithm known as *log-linear learning* in networked control systems [10]–[12]. Log-linear learning has received significant attention recently in the area of distributed control, as it can often be employed to ensure that the resulting behavior is near optimal. A representative set of examples range from control of wind farms [13], sensor networks [4], [5], [14]–[16], task assignment [17], among others [18]. However, the susceptibility of this approach to adversarial interventions is generally unknown.

The goal of this paper is to shed light on the susceptibility of log-linear learning to adversarial interventions. To that end, we focus on a well-studied class of systems known as *graphical coordination games* [19], [20]. Graphical coordination games model strategic scenarios where agents are tasked with adopting conventions and derive benefits from coordinating with the choices of their neighbors, e.g., adoption of technology or conventions [20], [21]. Regardless of the specifics of the graphical coordination game, log-linear learning is known to asymptotically achieve optimal system-level behavior. In this work, we focus on characterizing the degree to which the performance guarantees of log-linear learning algorithms can be undermined by adversarial manipulations.

In particular, our goal is to evaluate how different adversarial features can inflict harm on the system. How much more of a threat is an adversary that knows the underlying network structure versus one that does not? An adversary that can dynamically alter its strategy versus one that can not? We begin by stating our model to ensure that our contributions are clear.

This research was supported by UCOP grant LFR-18-548175, ONR grant #N00014-17-1-2060 and NSF grant #ECCS-1638214. The material in this paper substantially extends the conference paper [1] by providing complete proofs and novel results on dynamic policies. The current paper extends the results by fully characterizing dynamic adversarial influence.

K. Paarporn, M. Alizadeh, and J.R. Marden are with the Department of Electrical and Computer Engineering, University of California, Santa Barbara. B. Canty is with CACI International. P. N. Brown is with the Department of Computer Science at the University of Colorado, Colorado Springs. Contact: kpaarporn@ucsb.edu, {alizadeh,jrmarden}@ece.ucsb.edu, brian.canty@caci.com, philip.brown@uccs.edu.

*These authors contributed equally to this work.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2020.3038842, IEEE Transactions on Control of Network Systems

A. Model: Graphical Coordination Games

We consider the framework of graphical coordination games where there is a collection of agents $\mathcal{N} = \{1, 2, \ldots, n\}$ enmeshed in an underlying undirected network $G = (\mathcal{N}, \mathcal{E})$ where $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ defines the inter-agent interconnections. There are two different conventions, denoted by x and y, and each agent $i \in N$ must decide between a set of conventions $\mathcal{A}_i \subseteq \{x, y\}$. Note that if $\mathcal{A}_i = \{y\}$, this means that agent i is required to select convention y. The benefit agent i associates with a choice x or y depends on how many of its network neighbors $\mathcal{N}_i = \{j \in \mathcal{N} : (i, j) \in \mathcal{E}\}$ have selected the same convention. More formally, given a joint action profile $a = (a_1, \ldots, a_n) \in \mathcal{A} := \mathcal{A}_1 \times \ldots \mathcal{A}_n$, the total benefit agent i experiences is given by

$$U_i(a) := \sum_{j \in \mathcal{N}_i} V(a_i, a_j). \tag{1}$$

where $V : \{x, y\}^2 \to \mathbb{R}$ defines the per agent benefit of coordinating with a neighboring agent on a given convention. Throughout, we consider V of the following form

	x	y
x	$1+\alpha, 1+\alpha$	$0, 0$
y	$0, 0$	$1, 1$

where $\alpha > 0$. The *system welfare* associated with the action profile $a \in \mathcal{A}$ is given by

$$W(a) := \sum_{i \in \mathcal{N}} U_i(a). \tag{2}$$

The goal of a system operator is to assign decision-making rules for the agents such that their emergent collective behavior optimizes the system welfare, i.e., the emergent action profile is of the form

$$a^{\text{opt}} \in \arg\max_{a \in \mathcal{A}} W(a). \tag{3}$$

One such algorithm that achieves this objective is *log-linear learning* [11], [14], [22], [23]. Log-linear learning is a stochastic distributed algorithm that governs the evolution of agents' decisions over time. More formally, log-linear learning produces a sequence of joint action profiles $\{a(t)\}_{t=0}^{\infty}$, which we also call states, determined by the following process:

Definition 1 (Log-Linear Learning)**.** *Let $a(0) \in \mathcal{A}$ be any action profile. At each time $t \geq 1$, one agent $i \in \mathcal{N}$ is selected uniformly at random and allowed to alter its action choice. All other agents are required to repeat their previous action, i.e., $a_{-i}(t) = a_{-i}(t-1)$ where $a_{-i} = \{a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n\}$ captures the action choice of all agents $\neq i$. The updating agent i selects any action $a_i \in \mathcal{A}_i$ at time t with probability*

$$\frac{e^{\beta U_i(a_i, a_{-i}(t-1))}}{\sum_{\tilde{a}_i \in \mathcal{A}_i} e^{\beta U_i(\tilde{a}_i, a_{-i}(t-1))}} \tag{4}$$

where $\beta > 0$ is a given algorithm parameter. Once agent i selects her action, the process is repeated.

Log-linear learning induces in an ergodic process over the joint action profiles $\mathcal{A}$ in any graphical coordination game of the above form. The stochastically stable states, which we express by $\text{LLL}(G, \mathcal{A}, \alpha) \subseteq \mathcal{A}$ is defined as the support of the limiting distribution as $\beta \to \infty$. In the context of graphical coordination games, log-linear learning ensures that

$$\text{LLL}(G, \mathcal{A}, \alpha) = \arg\max_{a \in \mathcal{A}} W(a). \tag{5}$$

Note that log-linear learning guarantees that the emergent behavior optimizes the system-level objective irrespective of the graph G, the convention choices available to the agents $\mathcal{A}$, and the value of α. Note that in the special case when $\mathcal{A}_i = \{x, y\}$ for all $i \in N$, then $\text{LLL}(G, \mathcal{A}, \alpha) = \{\vec{x} = \{x, \ldots, x\}\}$ is the all x convention. For alternative choices of $\mathcal{A}$, the action profiles that optimize system welfare is not as straightforward.

B. Models of Adversarial Interventions

In this paper we consider an adversary seeking to influence the decision-making process of log-linear learning by strategically integrating $S = \{1, \ldots, |S|\}$ adversarial nodes into the system. Each of the adversarial nodes $s \in S$ will be integrated into the network through a connection to a unique single agent $i \in \mathcal{N}$ that the adversarial node is tasked with influencing though a choice $a_s = \{x\}$ or $a_s = \{y\}$. Let $S_x, S_y \subseteq \mathcal{N}$, $|S_x| + |S_y| \leq |S|$, denote the set of agents that are being influenced by an adversary promoting $\{x\}$ and $\{y\}$ respectively. Given S_x and S_y, the influenced utility of an agent $i \in \mathcal{N}$ is of the form

$$\tilde{U}_i(a; S_x, S_y) := \begin{cases} U_i(a) + V(a_i, x) & \text{if } i \in S_x \\ U_i(a) + V(a_i, y) & \text{if } i \in S_y \\ U_i(a) & \text{else} \end{cases} \tag{6}$$

In words, an agent $i \in S_y$ (resp. $i \in S_x$) experiences the usual benefits from its neighbors in $\mathcal{N}_i$, plus an additional utility of 1 if $a_i = y$ (resp. $1 + \alpha$ if $a_i = x$). While the adversarial nodes S do not directly contribute to the system-level objective as defined in (2), they modify the network agents' utility functions, which invariably influence the resulting asymptotic behavior associated with log-linear learning. We now denote by $\text{LLL}(G, \mathcal{A}, \alpha, \pi)$ the (possibly modified) stochastically stable states, where π defines the process, or *policy*, through which S_x and S_y are chosen. Technically speaking, the sets $S_x(t), S_y(t)$ are drawn from the distribution $\pi(t)$. The performance degradation associated with the adversarial policy π is measured by

$$\eta(G, \mathcal{A}, \alpha, \pi) := \min_{a \in \text{LLL}(G, \mathcal{A}, \alpha, \pi)} \left\{ \frac{W(a)}{W(a^{\text{opt}})} \right\} \geq 0. \tag{7}$$

We will focus on graphical coordination games where the agents have full choice of conventions, i.e., $\mathcal{A}_i = \{x, y\}$ for all $i \in N$. For that setting, we will omit highlighting the dependence of $\mathcal{A}$ in the definition of $\eta(\cdot)$ and $\text{LLL}(\cdot)$, i.e. we will instead write $\eta(G, \alpha, \pi)$ and $\text{LLL}(G, \alpha, \pi)$.

C. Summary of Contributions

The focus of this manuscript is on characterizing the susceptibility of log-learning learning to adversarial interventions in networked coordination games. In particular, our goal is

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2020.3038842, IEEE Transactions on Control of Network Systems

to identify the salient features of the worst-case adversarial policies. Specifically, we focus on identifying the importance of the following two attributes:

- Informational Awareness: Does the adversary know the network structure?

- Strategic Sophistication: Can the adversarial nodes dynamically alter their location and convention choice over time?

The above attributes define four classes of adversarial policies, which we represent by $\{\Pi_{\text{I,D}}, \Pi_{\text{I}}, \Pi_{\text{D}}, \Pi\}$, where the subscript I denotes informationally aware and the subscript D denotes dynamic adversarial policies. The absence of a subscript distinction means the negation. For example, $\Pi_{\text{I,D}}$ denotes the set of adversarial policies that are dynamic and can utilize information about the network structure. On the other hand, Π denotes the set of adversarial policies that are static and agnostic to network structure. By dynamic, we mean that the adversary can alter its behavior based on the current network state. That is, we consider stationary policies[1] $\{S_x(a(t)), S_y(a(t))\}_{t=1,2,\ldots}$. A static policy does not allow this flexibility: $S_x(a(t)) = S_x$ and $S_y(a(t)) = S_y$ $\forall t$.

Our first set of main results identify the most vulnerable graph structures. Focusing on a class of homogeneous ring graphs where an adversary can influence at most $\gamma \cdot n$ agents, where $\gamma \in [0, 1]$, we demonstrate that the most susceptible, i.e., graphs that lead to the lowest efficiency as defined in (7), are minimally connected ring graphs. We demonstrate this over the set of adversarial policies Π and Π_{D} (Theorems 2.1 and 2.2). This matches intuition as the graph with the fewest internal edges are in fact the most susceptible to adversarial interference.

Our second set of results focus exclusively on these ring graphs and seek to identify how information regarding the network structure can be exploited by the adversary. In doing so, we characterize the tight worst-case performance guarantees as in (7) over policies belonging to Π_{I} and $\Pi_{\text{I,D}}$ (Theorems 3.1 and 3.2). Figure 1 highlights an instance of worst-case performance guarantees for all four types of policies $\{\Pi_{\text{I,D}}, \Pi_{\text{I}}, \Pi_{\text{D}}, \Pi\}$ when $\alpha = 0.5$. As expected, the adversary leverages information and sophistication to most effectively degrade performance guarantees. However, the regimes where each of these attributes is most valuable is not so predictable. When an adversary has limited strength, i.e., $\gamma < 0.5$, sophistication is far more valuable than informational awareness to the adversary. That is, the best adversarial policy in Π_D significantly outperforms the best adversarial policy in Π_I. When an adversary has more substantial strength, i.e., $\gamma > 0.5$, the opposite is true. We formalize these conclusions in Theorem 3.3. Theorem 3.4 highlights the performance differences between static and dynamic policies. In particular, dynamic policies can achieve the same performance as static ones using fewer adversarial nodes, but such performance saturates above a threshold budget. We provide proofs in Sections IV (static adversaries) and V (dynamic adversaries).

[1]In this paper, we restrict attention to adversarial policies that induce stochastically stable states – in particular, static policies and dynamic policies that are stationary. It will be of interest in future work to investigate other types of dynamic policies that may not guarantee a SSS is induced.

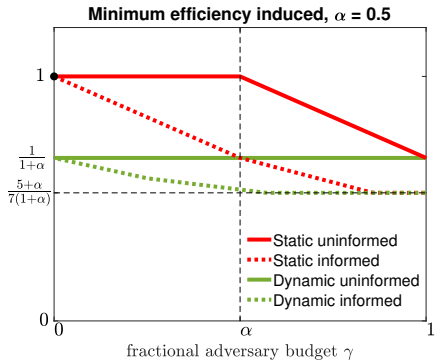

Fig. 1: This figure highlights the interplay between an adversary's informational awareness (informed vs uninformed), strategic sophistication (static vs dynamic), budget, and the minimum efficiency it can induce on the system. The green and red lines characterize minimum efficiencies induced from four different adversarial models on ring networks of sufficiently large size, as a function of fractional budget $\gamma \in [0, 1]$ (the fraction of agents the adversary can influence). At $\gamma = 0$, neither adversarial model can induce any damage on efficiency (black circle). For low budgets (i.e. $\gamma < 0.5$), strategic sophistication is more valuable than having system-level information about the network. The converse holds true for higher budgets (i.e. $\gamma > 0.5$): system-level information is more valuable than the ability to implement dynamic policies.

For space concerns, proofs of omitted results can be found in an online version https://arxiv.org/pdf/1909.02671.pdf.

D. Related Work

Previous work has studied to what extent networked distributed algorithms, designed to operate in the absence of adversarial interference, are susceptible to such influence [9], [24]–[28]. For example, distributed multi-agent optimization algorithms are shown to easily be compromised by adversarial behaviors [9], [28]. Indeed, there are fundamental limitations to these algorithms and their variants. Such algorithms cannot perform optimally in the absence of adversaries as well as be resilient to adversarial attacks at the same time [9], [28].

How emergent behavior associated with game-theoretic learning algorithms, such as log-linear learning, could be influenced by adversarial nodes was initially studied in [29]. The focus centered on how easily adversarial nodes could steer agents towards an inefficient Nash equilibrium. In this paper, we instead focus on an adversary seeking to minimize system-level performance.

II. Analysis of Susceptible Graphs

This section focuses on identifying which graph structures are most susceptible to adversarial influence. To that end, we will focus on a class of graphs that we term k-connected ring graphs, for $k \in \{1, \ldots, \lfloor n/2 \rfloor\}$. A graph $G = (\mathcal{N}, \mathcal{E})$ is a k-connected ring graph if $\mathcal{N}_i = \{i-k, \ldots, i-1, i+1, \ldots, i+k\}$ for each agent $i \in \mathcal{N}$, where addition and subtraction are both modulo n. Note that when $k = 1$ we have the usual ring graph and when $k = \lfloor n/2 \rfloor$ we have the complete graph. Let us denote $\mathcal{G}_n^k$ as the set of all k-connected ring graphs of size n. The following Theorems outline the degradation in

performance attainable through admissible adversarial policies belonging to Π and Π_{D} – static and dynamic uninformed adversaries, respectively.

Theorem 2.1. *Consider the class of network coordination games where (i)*[2] $\alpha \in [0,1)$*, and (ii) an admissible adversarial policy can influence at most a fraction* $\gamma \in [0,1]$ *of agents in the network. Recall* $\Pi(G,\gamma)$ *is the set of admissible static adversarial policies that are agnostic about the network structure. Then,*

$$\lim_{n\to\infty} \inf_{\substack{\pi\in\Pi(G,\gamma)\\ G\in\mathcal{G}_n^k}} \eta(G,\alpha,\pi) = \begin{cases} 1, & \text{if } \gamma < k\alpha \\ \frac{(1-(k-1)\alpha)-\alpha\gamma}{(1+\alpha)(1-k\alpha)}, & \text{if } \gamma \geq k\alpha \end{cases}. \tag{8}$$

Theorem 2.2. *Consider* $\Pi_D(G,\gamma)$*, the set of admissible dynamic adversarial policies that are agnostic about the network structure on any graph* $G \in \mathcal{G}_n^k$*. Then for* $\alpha \in [0,1)$ *and* $\gamma \in [0,1]$*,*

$$\inf_{\pi\in\Pi_D(G,\gamma)} \eta(G,\alpha,\pi) \geq \begin{cases} \frac{1}{1+\alpha}, & \text{if } \alpha < \frac{1}{k}, \gamma \neq 0 \\ 1, & \text{if } \alpha \geq \frac{1}{k} \text{ or } \gamma = 0 \end{cases}. \tag{9}$$

Furthermore, the limit of efficiency as the size of G *grows* $(n\to\infty)$ *equals the lower bound.*

There are several interesting things to note from Theorems 2.1 and 2.2. If $\alpha \geq 1/k$, neither classes of adversarial policies $\Pi_{\mathrm{D}}(G,\gamma)$ nor $\Pi(G,\gamma)$ can inflict any damage on the system regardless of the budget γ. Second, the achievable efficiency of a dynamic uninformed adversary, i.e., restriction to $\Pi_{\mathrm{D}}(\gamma)$, is constant for $\gamma \in (0,1]$. Third, the induced efficiency from a static uninformed adversary, i.e., restriction to $\Pi(G,\gamma)$, is decreasing in γ. Lastly, by tightness we know that for any $k \geq 1$ and $\gamma \in (0,1]$ we have

$$\inf_{G^1\in\mathcal{G}^1,\pi\in\Pi} \eta(G^1,\alpha,\pi) \leq \inf_{G^k\in\mathcal{G}^k,\pi\in\Pi} \eta(G^k,\alpha,\pi),$$

and an identical relation holds for policies in Π_{D}. Here, we omit highlighting the dependence on $\Pi(\cdot)$ for brevity. Hence, ring graphs ($k=1$) are the graphs that are most susceptible to adversarial interference.

III. The Impact of Information on Ring Graphs

The previous section demonstrated that ring graphs are the most susceptible to adversarial influence. In this section we explicitly characterize the impact informational awareness has on the potential degradation by admissible adversarial policies. We focus this analysis exclusively on ring graphs.

A. Static Informed Adversarial Policies

This section focuses on the potential degradation of the adversarial policies in the set Π_{I}. By knowing the graph's structure, the adversary can explicitly target specific agents $S_x, S_y \subseteq \mathcal{N}$ in the network. An adversarial policy $\pi \in \Pi_{\mathrm{I}}$ defines the process by which these agents are selected. The resulting policies is static is the sense that for all times $t \geq 1$, $S_x(t), S_y(t) = S_x, S_y$. The following Theorem characterizes the potential degradation caused by such adversarial policies.

Theorem 3.1. *Consider the class of network coordination games where (i)* $\alpha \in [0,1)$ *and (ii)* $G \in \mathcal{G}_n^1$*. Given a fractional adversarial budget consider* $\Pi_I(G,\gamma)$*, the set of admissible adversarial policies that are static, but can depend on the network structure. Then for* $\gamma \in (0,1]$*,*

$$\inf_{\pi\in\Pi_I(G,\gamma)} \eta(G,\alpha,\pi) \geq \inf_{\ell_{x_1},\ell_{x_2},\ell_{y_1},\ell_{y_2}\in\mathbb{Z}_{\geq 0}} \frac{1}{1+\alpha}\left(1+\frac{(2+\alpha)(\frac{s_1}{s_2}-1)+\alpha(\ell_{x_1}-\frac{s_1}{s_2}\ell_{x_2})}{\ell_{x_1}+\ell_{y_1}-\frac{s_1}{s_2}(\ell_{x_2}+\ell_{y_2})}\right),$$

subject to: for $j=1,2$*,*

$$\ell_{x_j} \geq 2, \quad \ell_{y_j} \geq \left\lceil \frac{2+\alpha}{1-\alpha} \right\rceil$$

$$s_j = \gamma(\ell_{x_j}+\ell_{y_j}) - \left\lceil \alpha(\ell_{y_j}+1) \right\rceil - 2 - \left\lceil \frac{\left[2-\alpha(\ell_{x_j}-1)\right]_+}{1+\alpha} \right\rceil$$

$$s_1 = 0 \text{ with } \ell_{x_2}, \ell_{y_2} = 0, \quad \text{or } s_1 > 0 \text{ and } s_2 < 0 \tag{SI-OPT}$$

For $\gamma = 0$*, the efficiency for any graph is 1. Here, we denote* $[z]_+ = \max\{z, 0\}$ *for any* $z \in \mathbb{R}$*. Furthermore, the limit of efficiency as the size of* G *grows* $(n\to\infty)$ *equals the lower bound* (SI-OPT).

There are several interesting things to note from Theorem 3.1, which characterizes the greatest damage that an adversary can inflict upon the system when relying on static policies that can depend on the graph structure. This theorem informs the structure of the worst-case attack, which involves the adversary attempting to stabilize alternating x, y sequences of four distinct lengths. While the structure of this adversarial attack is not necessarily fundamental, the interesting part of the theorem centers on tightness. That is, the adversary can never inflict more damage than the bounds given in Theorem 3.1, and the best adversarial strategy approaches this bound as the size of the ring graph in consideration gets larger.

B. Dynamic Informed Adversarial Policies

This section focuses on the potential degradation of the adversarial policies in the set Π_{DI}. Here, the adversary can target specific agents $S_x(a(t)), S_y(a(t)) \subseteq \mathcal{N}$ using knowledge of the graph structure G and the sequence of action profiles $\{a(t)\}_{t\in\mathbb{Z}_{\geq 0}}$. An adversarial policy $\pi \in \Pi_{DI}$ defines the process by which these agents are selected. The following Theorem characterizes the maximum potential degradation caused by such adversarial policies.

Theorem 3.2. *Consider* $\Pi_{DI}(G,\gamma)$*, the set of admissible adversarial policies that are dynamic and can depend on network structure, and* $\alpha \in [0,1)$*. Then the fundamental lower bound for* $\inf_{\pi\in\Pi_{DI}(G,\gamma)} \eta(G,\alpha,\pi)$ *is given by the RHS of* (SI-OPT)*, where the* s_j *variables are instead*

$$s_j = \begin{cases} \gamma(\ell_{x_j}+\ell_{y_j}) - 4 & \text{if } \alpha < \frac{1}{2} \text{ and } \ell_{x_j} \leq 1 + \left\lfloor \frac{1-\alpha}{\alpha} \right\rfloor \\ \gamma(\ell_{x_j}+\ell_{y_j}) - 2 & \text{else} \end{cases} \tag{10}$$

[2]Values of $\alpha \geq 1$ are not considered here. If this is the case, no y agents can be induced in the stochastically stable state under any adversarial policy on ring graphs, and no damage can be inflicted.

for $j = 1, 2$. Furthermore, the limit of efficiency as the size of G grows ($n \to \infty$) equals the lower bound.

Similar to (SI-OPT), the lower bound of Theorem 3.2 takes the form of an integer programming problem. While the structure of this adversarial attack is not necessarily fundamental, the interesting part of the theorem centers on tightness. That is, the adversary can never inflict more damage than the bound described in Theorem 3.2 and the best adversarial strategy approaches this bound as the size of the ring graph gets larger.

C. Comparison Between Information and Sophistication

Here, we emphasize the qualitative differences between information and sophistication. The Theorem below asserts that sophistication, i.e. the ability to implement a dynamic policy, is a more desirable attribute for the adversary if its budget is relatively low, while information is more valuable if its budget is high.

Theorem 3.3. *Suppose $\alpha \in [0, 1)$. For budgets $\gamma \in (0, \alpha)$ (empty interval if $\alpha = 0$), we have*

$$\lim_{n\to\infty} \inf_{\substack{G\in\mathcal{G}_n^1 \\ \pi\in\Pi_D(G,\gamma)}} \eta(G, \alpha, \pi) < \lim_{n\to\infty} \inf_{\substack{G\in\mathcal{G}_n^1 \\ \pi\in\Pi_I(G,\gamma)}} \eta(G, \alpha, \pi). \tag{11}$$

For budgets $\gamma \in (\alpha, 1]$, the opposite (strict) inequality holds. They are equal if $\gamma = \alpha$.

Hence, in the low budget regime $\gamma < \alpha$, the adversary prefers to be uninformed and dynamic over being informed but static. The opposite conclusion holds in the high budget regime $\gamma > \alpha$. This characterization allows us to explicitly identify the importance of information and sophistication in adversarial policies as highlighted in Figure 1[3].

The next result provides a comparison between static and dynamic informed adversaries. It states that given a sufficiently large adversarial budget, an optimal static informed policy can do just as much damage as an optimal dynamic informed policy.

Theorem 3.4. *The fundamental lower bound on performance for static informed policies is*

$$\left(\frac{1}{1+\alpha}\right)\frac{\ell^* + \alpha}{\ell^* + 2} \tag{12}$$

if and only if it has a budget $\gamma \geq \gamma_{sat}^{SI} := \frac{\ell^* + \left\lceil \frac{2-\alpha}{1+\alpha} \right\rceil}{\ell^*+2}$, where $\ell^* := \left\lceil \frac{2+\alpha}{1-\alpha} \right\rceil$. Furthermore, the fundamental lower bound on performance for dynamic informed policies coincides with (12) for budgets $\gamma \geq \gamma_{\text{sat}}^{DI} := \frac{2+2\cdot\mathbb{1}\left(\alpha<\frac{1}{2}\right)}{\ell^*+2}$, where $\mathbb{1}(\cdot)$ is the indicator function.

In other words, there are saturation levels on budget for both types of adversaries (DI and SI), where influencing more than $\gamma_{\text{sat}}^{\text{DI}}$ ($\gamma_{\text{sat}}^{\text{SI}}$) fraction of agents does not offer any additional performance gains. However, a static adversary will not exhibit saturation if $\alpha < \frac{1}{2}$. That is, the static adversary achieves performance level (12) if and only if it has a full budget $\gamma = 1$. It is interesting to note from the above Theorem that the dynamic informed adversary can maintain the performance level (12) for a wider range of budgets $\gamma \in [\gamma_{\text{sat}}^{\text{DI}}, 1] \supseteq [\gamma_{\text{sat}}^{\text{SI}}, 1]$ than the static informed adversary can. Here, the range is the same (no saturation exhibited for either) if and only if $\alpha = 0$. Essentially, dynamic policies can inflict the same level of damage with fewer adversaries than a static policy.

[3]Figure 1 plots the bounds that the four main results, Theorems 2.1, 2.2, 3.1, and 3.2, characterize. While the bounds are analytically derived for Theorems 2.1, 2.2 (uninformed adversaries), the plots for informed adversaries resemble a closely approximated value by solving their respective integer optimization problems with a finite upper bound of 100 on the decision variables.

IV. Proofs: Performance of static policies

In this section, we provide proofs for the minimum efficiency a static adversary can induce. We will first prove Theorem 3.1, the case of a static informed adversary. As discussed, we limit our attention here to ring graphs $G \in \mathcal{G}^1$. We then give a proof of Theorem 2.1, the case of a static uninformed adversary. This result relies on extending an intermediate step from the proof of Theorem 3.1 to k-connected ring graphs.

The adversary's objective is to steer the system to a stochastically stable state of minimal efficiency. We will refer to action profiles that can be stabilized through some static policy as the set of *target profiles* a static uninformed and static informed adversary can induce, respectively. Indeed, we would like to characterize the target profile of minimal efficiency an adversary can achieve over any ring graph, i.e.

$$\inf_{G\in\mathcal{G}^1, \pi\in\Pi_{\text{I}}(G,\gamma)} \eta(G, \alpha, \pi). \tag{13}$$

Our approach is to view any action profile a (and hence any target profile) as composed of alternating x and y segments. A y segment L_y is any subset $\{j, j+1, \ldots, j+|L_y|-1\} \subseteq \mathcal{N}$ such that $a_i = y \ \forall i \in L_y$ and $a_{j-1} = a_{j+|L_y|} = x$ (modulo n arithmetic). Similarly, L_x describes any such segment of x agents.

A. Proof of Theorem 3.1

To begin, we start with a general outline of the forthcoming proof, which we break up into three steps. Following the outline, we give proofs for each of the individual steps.

Step 1: Necessary and sufficient budget conditions to stabilize target profiles

We derive the minimum number of adversarial nodes that is necessary and sufficient to stabilize a given action profile a. Indeed, suppose S is an allocation of adversarial nodes. Then a is stochastically stable if and only if for every y segment L_y and x segment L_x contained in a,

$$|S_y \cap L_y| \geq \lceil \alpha(|L_y| + 1) \rceil + 2, \tag{14}$$

$$|S_x \cap L_x| \geq \left\lceil \frac{[2 - \alpha(|L_x| - 1)]_+}{1+\alpha} \right\rceil, \tag{15}$$

and the spacing from two sequential y adversarial nodes within L_y is no more than $\left\lceil \frac{1}{\alpha} \right\rceil$. We observe that segment lengths must satisfy $|L_y| \geq \left\lceil \frac{2+\alpha}{1-\alpha} \right\rceil$ and $|L_x| \geq 2$.

Step 2: Characterizing minimal efficiency target profiles

Having established the number of adversarial nodes needed to stabilize target profiles, we identify structural properties of

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2020.3038842, IEEE Transactions on Control of Network Systems

minimal efficiency target profiles that are stabilizable within the budget $\gamma \in (0, 1]$. In particular, we show that

(2A) Among adversarial policies that induce maximal damage, there is at least one that utilizes its full budget. Specifically, if the policy $\pi_S \in \Pi_{\mathrm{I}}(G, \gamma)$ with $|S| < \lfloor \gamma \cdot n \rfloor$ stabilizes profile a, then one can always use a policy $\pi_{S'}$ with $|S| = \lfloor \gamma \cdot n \rfloor$ that also stabilizes a.

(2B) The target profile of minimal efficiency contains at most two unique x y segment patterns.

Step 3: Optimization over worst-case target profiles

We formulate an integer optimization problem whose solution gives (13). The decision variables are the lengths of the two unique x y segment patterns, subject to necessity constraints derived from (14) and (15), as well as constraints given by the structural properties (2A) and (2B) of minimal efficiency target profiles. This formulation yields (SI-OPT), and thus the proof of Theorem 3.1.

Before getting into the proofs of the claims given in the outline, we first present preliminary analytical tools for characterizing the emergent behavior when an adversarial policy $\pi_S \in \Pi_{\mathrm{I}}$ interferes with the agents' log-linear learning dynamics. Specifically, we seek to compute the stochastically stable states $\mathrm{LLL}(G, \alpha, \pi_S)$. To do this, we can rely on the fact the graphical coordination game with static adversarial influence has a potential game structure [30]. In potential games, the stochastically stable states associated with log-linear learning are the action profiles that maximize the potential function [31], [32]. One can show that $\phi(a; S) := \frac{W(a)}{2} + \sum_{i \in S_x} V(a_i, x) + \sum_{i \in S_y} V(a_i, y)$. is a potential function for this game. Here, ϕ simply measures the number of coordinating links, including those induced from adversaries, weighted by their payoffs (i.e. x or y links). Hence, for any graph G and static policy π_S, we have $\mathrm{LLL}(G, \alpha, \pi_S) = \arg\max_{a \in \mathcal{A}} \phi(a; S)$.

Proof of Step 1

We present the proof only for y segments, as the arguments for x segments are analogous. Suppose a is stochastically stable, and contains a y segment $L_y = \{j, j+1, \ldots, j+|L_y|-1\}$. That is, $a_i = y$ for $i \in L_y$, and $a_{j-1} = a_{j+|L_y|} = x$. Consider any deviation a' from a that differs only within the segment L_y. Then it holds that $\phi(a'; S) \leq \phi(a; S)$. In particular, if a' is the profile where all agents in L_y deviate to x, then $(1+\alpha)(|L_y|+1) \leq |L_y|-1+|S_y \cap L_y|$ must hold. Rearranging, we obtain $|S_y \cap L_y| \geq \alpha(|L_y|+1)+2$. Since $|S_y \cap L_y|$ is a non-negative integer, it must hold that $|S_y \cap L_y| \geq 2+\lceil \alpha(|L_y|+1) \rceil$.

To prove sufficiency, we need to construct an allocation $S_y \cap L_y$ of $\lceil \alpha(|L_y|+1) \rceil + 2$ y adversarial nodes such that $\phi(a; S) \geq \phi(a'; S)$, where $a_i = y \ \forall i \in L_y$ and for any a' deviating from a in agents only in L_y. We first assume that $|L_y| \geq \lceil \alpha(|L_y|+1) \rceil + 2$, i.e. the length of the segment itself is greater or equal to the necessary number of adversaries needed. Indeed, let us define the sets W_1 and W_2 as follows:

$$W_1 = \{i \in L_y : \lfloor \alpha(i-j+1) \rfloor - \lfloor \alpha(i-j) \rfloor > 0\}, \quad (16)$$

$$W_2 = \{j, w, j+|L_y|-1\}, \quad (17)$$

where $w = \max\{i : i \in L_y \setminus (W_1 \cup \{j+|L_y|-1\})\}$, i.e. the largest index that is neither in W_1 nor is the endpoint

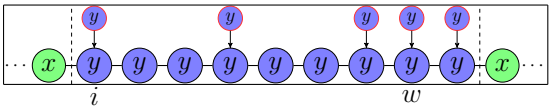

Fig. 2: An illustration of the constructed influence set given by (16), (17) to stabilize an isolated y segment. The y adversaries belonging to S_y are depicted as the smaller circles attaching to agents (larger circles) in the network. In this example, $\alpha = \frac{1}{4}$ and $|L_y| = 9$. The necessary and sufficient number of adversaries to stabilize the segment is 5.

$j + |L_k| - 1$. Then, set $S_y \cap L_y = W_1 \cup W_2$. An illustration of this influence set is depicted in Figure 2. Such a placement "spreads out" adversaries along L_y at a spacing of $\lceil \frac{1}{\alpha} \rceil$ nodes, and additionally places adversaries at the endpoints. This placement ensures the sufficiency condition. Proof of this claim can be found in the online version. ■

Proof of property (2A)

Suppose the minimum efficiency profile a^* on the graph $G \in \mathcal{G}^1$ is stabilized by the policy $\pi_S \in \Pi_{\mathrm{I}}(G, \gamma)$, where all adversaries are not utilized: $|S| < \lfloor \gamma \cdot n \rfloor$. The conditions (14) and (15) are met for all x and y segments, respectively. One can always add in remaining available x adversaries (y) to the existing x (y) segments while retaining stability of a^*. Therefore, there exists a policy π_S with $|S| = \lfloor \gamma \cdot n \rfloor$ that also stabilizes a^*. ■

Proof of property (2B)

Before proving this property explicitly, we first define some relevant notations. We can describe a profile a as a sequence of alternating segments $L_x^1 L_y^1 L_x^2 L_y^2 \cdots$. For each unique segment pair pattern that appears in a, i.e. $|L_x|$ x agents followed by $|L_y|$ y agents, let us define the (column) vector $\boldsymbol{\ell}_x(a)$ whose elements are the lengths $|L_x|$ among the unique patterns. We define $\boldsymbol{\ell}_y(a)$ similarly for the corresponding lengths $|L_y|$. Let us also define the vector $\boldsymbol{r}(a)$ whose elements are the number of times each unique pattern appears in a. We will drop the dependencies on a when the context is clear. We refer to $\boldsymbol{r}$ as the *repetition* vector. The efficiency of a can be rewritten in the following suggestive form:

$$\eta(a) = \frac{\boldsymbol{r}^\top ((1+\alpha)\boldsymbol{\ell}_x + \boldsymbol{\ell}_y) - (2+\alpha)||\boldsymbol{r}||_1}{(1+\alpha)\boldsymbol{r}^\top(\boldsymbol{\ell}_y + \boldsymbol{\ell}_x)}. \quad (18)$$

We note that the denominator of (18) is simply the number of links in the ring, n, multiplied by $1+\alpha$. This indicates the optimal welfare $\frac{1}{2}W(a^{\mathrm{opt}})$. The numerator of (18) counts (and weights with associated payoff) the number of coordinating x and y links given the description vectors $\boldsymbol{\ell}_x$, $\boldsymbol{\ell}_y$, and $\boldsymbol{r}$. For a profile a and its associated description vectors $\boldsymbol{\ell}_x$ and $\boldsymbol{\ell}_y$, let us define the vector $\boldsymbol{s}(a)$ of identical length, whose components are given by $s_j(a) := \gamma(\ell_{x,j} + \ell_{y,j}) - \lceil \alpha(\ell_{y,j} + 1) \rceil - 2 - \left\lceil (1+\alpha)^{-1} [2 - \alpha(\ell_{x,j} - 1)]_+ \right\rceil$. The number s_j is the difference between the adversaries available to a particular segment pattern (given budget γ) whose length is given by $\ell_{x,j}$ and $\ell_{y,j}$, and the minimum number of adversaries needed to ensure its stability (given by (14), (15)). We refer to $\boldsymbol{s}$ as the *surplus* vector. The quantity $\boldsymbol{r}^\top \boldsymbol{s}$ is the excess budget after using the minimum required number of adversaries to stabilize a. Property (2A) asserts that a target profile of minimum

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2020.3038842, IEEE Transactions on Control of Network Systems

efficiency satisfies $\boldsymbol{r}^\top \boldsymbol{s} = 0$.

Now, consider an action profile a^1 with $\boldsymbol{\ell}_x^1 = (\ell_{x,1}, \ell_{x,2}, \ell_{x,3})$, $\boldsymbol{\ell}_y^1 = (\ell_{y,1}, \ell_{y,2}, \ell_{y,3})$ and $\boldsymbol{s} = (s_1, s_2, s_3)$ with $s_1 > 0$ and s_2, $s_3 < 0$. Hence, we can find $\boldsymbol{r}^1$ such that $(\boldsymbol{r}^1)^\top \boldsymbol{s}^1 = 0$. Thus, a^1 is a candidate for a minimum efficiency stable state. Furthermore, consider the profiles a^2 and a^3 (possibly defined on different ring graphs), where a^2 is associated with $\boldsymbol{\ell}_x^2 = (\ell_{x,1}, \ell_{x,2})$ and $\boldsymbol{\ell}_y^2 = (\ell_{y,1}, \ell_{y,2})$, and a^3 is associated with $\boldsymbol{\ell}_x^3 = (\ell_{x,1}, \ell_{x,3})$ and $\boldsymbol{\ell}_y^3 = (\ell_{y,1}, \ell_{y,3})$. One can find repetition vectors $\boldsymbol{r}^2$, $\boldsymbol{r}^3$ that satisfy $(\boldsymbol{r}^2)^\top \boldsymbol{s} = 0$ and $(\boldsymbol{r}^3)^\top \boldsymbol{s} = 0$.

Define $g_i = \ell_{y,i} + (1+\alpha)\ell_{x,i} - (2+\alpha)$ and $\ell_i = \ell_{x,i} + \ell_{y,i}$ for each $i = 1, 2, 3$. We can express efficiency of a^1 as $\eta(a^1) = \frac{r_2^1(g_2 - \frac{s_2}{s_1} g_1) + r_3^1(g_3 - \frac{s_3}{s_1} g_1)}{(1+\alpha)(r_2^1(\ell_2 - \frac{s_2}{s_1}\ell_1) + r_3^1(\ell_3 - \frac{s_3}{s_1}\ell_1))}$. One can write the efficiencies of a^2, a^3 as $\eta(a^2) = \frac{g_2 - \frac{s_2}{s_1} g_1}{(1+\alpha)(\ell_2 - \frac{s_2}{s_1}\ell_1)}$ and $\eta(a^3) = \frac{g_3 - \frac{s_3}{s_1} g_1}{(1+\alpha)(\ell_3 - \frac{s_3}{s_1}\ell_1)}$. Observe that $\eta(a^1)$ is a mediant sum of weighted values $\eta(a^2)$ and $\eta(a^3)$. Hence, either $\eta(a^2)$ or $\eta(a^3)$ is less than or equal to $\eta(a^1)$. This result can be extended in a similar way to show that for any profile consisting of multiple segment patterns, one can construct another profile of lower efficiency using up to two unique segment patterns from the original action profile. ■

Proof of Step 3 (Theorem 3.1)

Using the collection of results we have obtained in Steps 1-3, we can now prove Theorem 3.1. From property (2B), the search for a minimal efficiency stable state, i.e., one that gives the efficiency (13), reduces to finding four lengths: $\boldsymbol{\ell}_x = (\ell_{x,1}, \ell_{x,2})$ and $\boldsymbol{\ell}_y = (\ell_{y,1}, \ell_{y,2})$. The form of the objective function in the integer program of (SI-OPT) thus coincides with the expression for $\eta(a^2)$ in property (2B). Each $\ell_{z,i}$, $z \in x, y$ and $i \in \{1, 2\}$, must satisfy the length criterion $\ell_{y,i} \geq \left\lceil \frac{2+\alpha}{1-\alpha} \right\rceil$ and $\ell_{x,i} \geq 2$ (3 if $\alpha = 0$). These length conditions are consequences of the stabilizability conditions (14) and (15). Lastly, one can find a repetition vector $\boldsymbol{r}$ that satisfies $\boldsymbol{r}^\top \boldsymbol{s} = 0$, as long as $s_1 > 0$ and $s_2 < 0$, or $s_1 = 0$ with $\ell_{x,1}, \ell_{x,2} = 0$. ■

B. Proof of Theorem 2.1

Here, we provide a proof of Theorem 2.1, which characterizes the minimal efficiency a static uninformed adversary can induce on a k-connected ring graph. The arguments rely on an extension of intermediate step 1 from the proof of Theorem 3.1 to k-connected ring graphs.

In particular, the necessary and sufficient condition to stabilize a y segment in a k-connected ring graph is

$$|S_y \cap L_y| \geq \left\lceil \alpha \left(k|L_y| + \frac{k(k+1)}{2} \right) \right\rceil + k(k+1), \quad (19)$$

and the spacing between two sequential y adversaries within L_y is no more than $\lceil \frac{1}{k\alpha} \rceil$. Note that according to this condition, the segment length must also satisfy $|L_y| \geq \max\left\{1, \lceil \frac{k(k+1)(1+\alpha/2)}{1-k\alpha} \rceil \right\}$. A derivation of the condition is as follows. There are $\sum_{j=1}^{k} (|L_y| - j)$ links between agents in L_y. Assuming $|L_y|$ satisfies the length requirement, there are $2\sum_{j=1}^{k} j$ links from L_y to outside L_y. The potential of y_{L_y} (all agents in L_y play y) exceeds that of x_{L_y} (all play x) if $|S_y \cap L_y| + \sum_{j=1}^{k}(|L_y| - j) \geq (1+\alpha)\left[\sum_{j=1}^{k}(|L_y| - j) + 2\sum_{j=1}^{k} j\right]$, which reduces to (19). One can prove sufficiency in a similar manner as step 1 from the previous section – by allocating the y adversaries with a spacing of $\lceil \frac{1}{k\alpha} \rceil$ apart, the potential of y_{L_y} exceeds that of any other $a_{L_y} \neq \{y_{L_y}, x_{L_y}\}$.

We are now ready to prove Theorem 2.1. A static and uninformed policy cannot strategically place adversarial nodes in the network. It can only specify the numbers of x and y adversaries. Its baseline performance is given by the minimal damage that can be inflicted over all possible allocations of these adversaries. Hence to characterize (8), we seek the allocation of adversaries that ensures the best-case efficiency for the network.

First, we consider the case $\gamma < k\alpha$. The adversarial nodes can be allocated sparsely enough across the entire network such that the condition (19) is violated. Consequently, the all x profile is the unique stochastically stable state. Therefore, no damage can be inflicted on the system in this regime. Note that if $k\alpha > 1$, no damage is possible regardless of the budget.

Now, consider $\gamma > k\alpha$. If $k\alpha < 1$ and $G \in \mathcal{G}^k$ is sufficiently large, an allocation of y adversaries according to (19) would ensure conversion of the entire network to y, giving an efficiency of $\frac{1}{1+\alpha}$. However, let us consider a re-allocation of these adversaries that maximally mitigates such damage. The idea is to only allow a minimal fraction f of the network to be converted to y, while the rest of the network plays x.

Suppose y adversaries are allocated to every agent in a contiguous segment, whose length is a fraction f of the entire network. Suppose this segment is sufficiently long such that (19) is satisfied. Now, the remaining $\gamma - f$ adversaries should be allocated to the rest of the network such that the remaining fraction $1 - f$ of the network (another contiguous segment) is still stable to x. Indeed, an adversarial agent density of up to $k\alpha$ in the remaining network fails to induce any y agents. The smallest f that satisfies these conditions is given by $f = \frac{\gamma - k\alpha}{1 - k\alpha}$. This establishes (8). Note in this analysis, the adversary exclusively chooses to implement y adversaries. Based on the above arguments, an optimal static uninformed policy never chooses to use x adversaries.

V. Proofs: Performance of Dynamic Policies

In this section, we give proofs for the minimum efficiency dynamic adversaries can induce. Similar to Section IV, we will first prove Theorem 3.2, the case of a dynamic informed adversary. We then give the proof of Theorem 2.2.

Due to the time-dependent nature of dynamic policies, we cannot rely on potential game arguments to compute stochastically stable states as we did in Section IV. One must instead leverage the theory of regularly perturbed Markov processes and resistance trees. Before delving into the proof of Theorem 3.2, we provide a brief overview of this theory below. More detailed treatments can be found in [21], [33].

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2020.3038842, IEEE Transactions on Control of Network Systems

A. Preliminary: Regularly perturbed Markov processes and resistance trees

Definition 2. *A Markov process with transition matrix P^ϵ defined over state space $\mathcal{A}$ and parameterized by a perturbation $\epsilon \in (0, \bar{\epsilon}]$ for some $\bar{\epsilon} > 0$ is a* regular perturbation of the process P^0 *if it satisfies:*

1) *P^ϵ is aperiodic and irreducible for all $\epsilon \in (0, \bar{\epsilon}]$.*
2) *$\lim_{\epsilon \to 0^+} P^\epsilon(a, a') \to P^0(a, a')$ for all $a, a' \in \mathcal{A}$.*
3) *If $P^\epsilon(a, a') > 0$ for some $\epsilon \in (0, \bar{\epsilon}]$ then there exists $r(a, a') \geq 0$ such that $0 < \lim_{\epsilon \to 0^+} \frac{P^\epsilon(a,a')}{\epsilon^{r(a,a')}} < \infty$. We call $r(a, a')$ the* resistance *of transition $a \to a'$.*

The log-linear learning process is a regularly perturbed process with error parameter $\epsilon = e^{-\beta}$. The transition graph of P^ϵ is a directed graph whose nodes are the action profiles $\mathcal{A}$ and the edge (a, a') exists if and only if $P^\epsilon(a, a') > 0$. The weights of such edges are given by the resistances $r(a, a')$. The resistance of a path of length m, $\zeta = (z_1 \to z_2 \to \cdots \to z_m)$, is the sum of resistances along the state transitions: $r(\zeta) := \sum_{k=1}^{m-1} r(z_k, z_{k+1})$. Let us denote the *recurrent classes* of the unperturbed process P^0 as $E_1, E_2, \ldots, E_N$ with $N \geq 1$ where each class $E_k \subset \mathcal{A}$. A recurrent class satisfies the following.

1) For all $a \in \mathcal{A}$, there is a zero resistance path from a to E_k for some $k \in \{1, \ldots, N\}$.
2) For all $k \in \{1, \ldots, N\}$, and all $a, a' \in E_k$, there exists a zero resistance path from a to a' and from a' to a.
3) For all a, a' with $a \in E_k$ for some $k \in \{1, \ldots, N\}$ and $a' \notin E_k$, $r(a, a') > 0$.

One can also consider another directed transition graph whose nodes are the N recurrent classes. In this graph, all edges exist. Edge (E_i, E_j) is weighted by ρ_{ij}, defined as the minimum resistance among paths in the action profile transition graph starting from E_i and ending in E_j: $\rho_{ij} := \min_{a \in E_i, a' \in E_j} \min_{\zeta \in \mathcal{P}(a \to a')} r(\zeta)$. where $\mathcal{P}(a \to a')$ denotes the set of all paths starting at a and ending at a'. Let $\mathcal{T}_k$ be the set of all spanning trees rooted in the class E_k. That is, an element of $T \in \mathcal{T}_k$ is a directed graph with $N - 1$ edges such that there is a unique path from E_j to E_k, for every $j \neq k$. The resistance $R(T)$ of the rooted tree T is the sum of resistances ρ_{ij} on the $N - 1$ edges that compose it. Now, define $\psi_k := \min_{T \in \mathcal{T}_k} R(T)$ as the *stochastic potential* of recurrent class E_k. We will use the following result to identify stochastically stable states.

Lemma 5.1 (from [33])**.** *The state $a \in \mathcal{A}$ is stochastically stable if and only if $a \in E_k$, where $k \in \arg\min_{j \in \{1, \ldots, N\}} \psi_j$. That is, it belongs to a recurrent class with minimum stochastic potential. It is the unique stochastically stable state if and only if $E_k = \{a\}$ and $\psi_k < \psi_j$, $\forall j \neq k$.*

B. Proof of Theorem 3.2

The logic of the proof follows the same three-step structure as the proof of Theorem 3.1. The only component that differs are the necessary and sufficient budget conditions to stabilize x and y segments. Indeed, we expect a dynamic informed adversary to need fewer adversaries than its static counterpart.

An outline of the proof is as follows. We show a particularly defined dynamic policy, which we term an *aggressive policy*, is sufficient to stabilize a given target profile. This entails proving that a is the recurrent class of minimum stochastic potential (see Lemma 5.1). To do so, we demonstrate that minimum resistance paths leaving each class leads to another that is more "similar" to the target profile a. We then prove necessity – any other dynamic policy utilizing strictly fewer adversarial nodes than the aggressive policy cannot stabilize a. We can then formulate an integer optimization problem similar to Theorem 3.1, but with different constraints on the number of adversaries needed for each x and y segment. To begin, we formally define the aggressive policy based on profile $a \in \mathcal{A}$.

Definition 3. *(Aggressive policy targeting a). An* aggressive policy targeting $a \in \mathcal{A}$ *is a state-dependent policy with adversarial placements $\{S_x(a(t)), S_y(a(t))\}_{t \geq 0}$ satisfying the following properties.*

1) (Defensive y strategy) *For each y-segment L_y contained in a, suppose $[p, q] = \{p, p+1, \ldots, q\} \subseteq L_y$, with $p \neq q$, is the longest segment of agents within L_y playing y in $a(t)$. Then,*

$$S_y(a(t)) \cap [p, q] = \begin{cases} \{p, q\} & \text{if } a_{p-1}(t) = a_{q+1}(t) = x \\ p & \text{if } a_{p-1}(t) = x,\ a_{q+1}(t) = y \\ q & \text{if } a_{p-1}(t) = y,\ a_{q+1}(t) = x \end{cases} \tag{20}$$

If the length of $[p, q]$ is one ($p = q$), then $p \notin S_y(a(t))$.

2) (Defensive x strategy) *For each x-segment L_x contained in a, suppose $[p, q]$, $p \neq q$, is the longest segment of agents within L_x playing x in $a(t)$.*

(a) If $\alpha < \frac{1}{2}$ and $|L_x| \leq 1 + \left\lfloor \frac{1-\alpha}{\alpha} \right\rfloor$, then

$$S_x(t) \cap [p, q] = \begin{cases} \{p, q\} & \text{if } q - p = 1 \\ \varnothing & \text{otherwise} \end{cases} \tag{21}$$

If $|L_x| \geq 2 + \left\lfloor \frac{1-\alpha}{\alpha} \right\rfloor$, then $S_x(t) \cap L_x = \varnothing$.

(b) Suppose $\alpha \geq \frac{1}{2}$. Then $S_x(t) \cap L_x = \varnothing$.

3) (Offensive strategies) *Consider the segment of lowest index that is not aligned, i.e. $a_i(t) \neq y$ for at least one $i \in L_y = [u, v]$, for a y-segment. Then the following properties hold for $S_y(t)$. A similar implementation holds if it is an x-segment.*

(a) Denote $[p, q] \subset L_y$ as the longest segment of agents within L_y playing y in $a(t)$. Then $S_y(t) \cap L_y$ contains either $p - 1$ or $q + 1$, but not both.

(b) If $a_i(t) = x$ for all $i \in L_y$, then $S_y(t) \cap L_y$ contains

- *u or v (but not both), when $a_{u-1}(t) = a_{v+1}(t)$.*
- *u, when $a_{u-1}(t) = y$ and $a_{v+1}(t) = x$.*
- *v, when $a_{u-1}(t) = x$ and $a_{v+1}(t) = y$.*

Properties 1 and 2 describe "defensive y (resp. x)" strategies to maintain y (x) segments over time. A defensive y strategy is implemented on all y-segments at any given time. In property 2, defensive x strategies are applied only if $\alpha < \frac{1}{2}$, and to segments that are shorter than a threshold length. Furthermore, the strategy does not "activate" until there are at most two consecutive agents playing x in the segment. Property 3 describes "offensive" strategies that are intended to convert segments back to their original type x or y. Note that the aggressive policy applies an offensive strategy to only

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2020.3038842, IEEE Transactions on Control of Network Systems

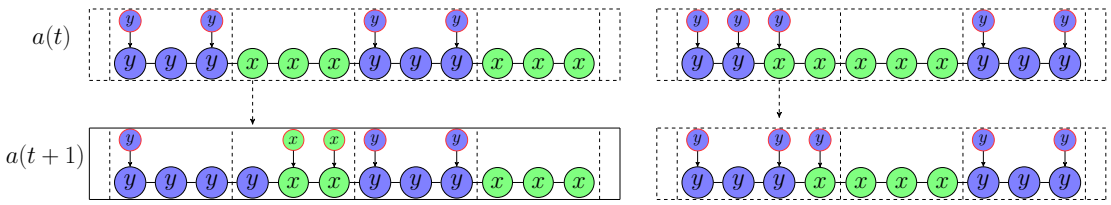

Fig. 3: We illustrate here both defensive and offensive strategies in an aggressive policy. (Left) Defensive y strategies are applied the first and third segments from the left. The fourth agent from the left transitioning from x to y at time $t+1$ activates a defensive x strategy in the second. No adversaries are deployed to x segments until only two neighboring agents playing x remain. (Right) A defensive and offensive y strategy are applied simultaneously to the first segment. The offensive strategy attaches a y adversary to the x agent that has a y neighbor.

a single segment at any time t, if needed. Figure 3 depicts an illustration of allocations of adversarial nodes to segments according to an aggressive policy.

Let n_y be the number of y segments and n_x the number of x segments of length at most $1+\left\lfloor \frac{1-\alpha}{\alpha} \right\rfloor$ in profile a. Then, the minimum number of adversarial nodes needed to implement an aggressive policy targeting a is

$$\begin{cases} 2(n_y+n_x)+1, & \text{if } \alpha < \frac{1}{2} \\ 2n_y+1, & \text{if } \alpha \geq \frac{1}{2} \end{cases} \tag{22}$$

Here, the additional $+1$ adversary is needed to implement an offensive strategy. Under log-linear learning, states transition via unilateral agent deviations. If two profiles a^1 and a^2 differ only by agent i's deviation, the resistance is $r(a^1 \to a^2) = \left[\tilde{U}_i(a_i^1, a_{-i}^1; S(a^1)) - \tilde{U}_i(a_i^2, a_{-i}^1; S(a^1))\right]_+$, where recall $\tilde{U}_i$ is agent i's perceived utility (6). The following result characterizes the minimum required lengths for x and y segments in a target profile.

Lemma 5.2. *If $a \in \mathcal{A}$ is stochastically stable under the aggressive policy targeting a, then*

- *all x segments of a are of length 2 or greater.*
- *all y segments of a are of length $\lceil \frac{2+\alpha}{1-\alpha} \rceil$ or greater.*

Proof. Suppose $(a_{i-1}, a_i, a_{i+1}) = (y, x, y)$ for some agent i. Regardless of what adversarial policy is applied, there is a zero resistance path out of a. Specifically, $r(a \to a') = 0$ where $a_i = y$ and $a'_{-i} = a_{-i}$. The least resistant path from a' to a is $1-\alpha > 0$, possible if and only if $i \in S_x(a')$. Hence, a is not a recurrent class and therefore is not stochastically stable.

Suppose a has a y segment L_y and $|L_y| \leq \left\lceil \frac{1+2\alpha}{1-\alpha} \right\rceil$. Let a' be the similar profile with $a'_{L_y} = x_{L_y}$. Note that a' is a recurrent class. When the aggressive policy applies an offensive strategy on L_y, the minimum resistance path starting from a' and ending in a is given by a border agent's $x \to y$ transition (having resistance $1+2\alpha$), followed by each subsequent neighbor's $x \to y$ transition (each having resistance 0). The resistance of this path is $\rho_{a',a} = 1+2\alpha$.

The minimum resistance path starting from a and ending in a' consists of $|L_y|-1$ transitions of resistance $1-\alpha$. Hence, $\rho_{a,a'} = (1-\alpha)(|L_y|-1) < (1-\alpha)\frac{1+2\alpha}{1-\alpha} = 1+2\alpha$. Let T be the minimum resistance tree rooted in a, and note that the edge (a', a) is necessarily part of T. Consider the tree T' rooted in a' by replacing the edge (a', a) from T with (a, a'). T' has lower stochastic potential than T and therefore a is not stochastically stable. ■

Henceforth, we only consider target profiles a with properties given by Lemma 5.2. Note that these minimum length requirements coincide with those in the static informed case. In the forthcoming analysis, we characterize the set of recurrent classes induced by the aggressive policy. We first define terminology to describe any profile a' relative to $a \in \mathcal{A}$. Let L_y be a y-segment of a. We say the segment L_y is *homogeneous in a'* if $a'_i = a'_j$ for all $i, j \in L_y$. We say L_y is *heterogeneous* if it is not homogeneous. Similar terminology applies for x-segments of a. We say the profile a' is homogeneous if every x and y-segment is homogeneous in a', and it is heterogeneous if it is not homogeneous. We will denote particular portions of a homogeneous action profile a' with brackets $|_x$ and $|_y$ that separate the segments based on the target profile a. For instance, $|X|_x X|_y$ refers to the actions of agents in an action profile for two consecutive segments with all agents playing x in the first as well as the second. The subscripts convey that agents play x in the target profile a in the first segment and y in the second segment.

The following result characterizes the set of all recurrent classes induced by the aggressive policy. In particular, each recurrent class consists of a single homogeneous action profile.

Lemma 5.3. *The recurrent classes associated with the aggressive policy targeting a are the homogeneous action profiles that do not contain an instance of $|X|_y Y|_x$.*

Let us denote $\mathcal{A}_{\mathsf{R}} \subset \mathcal{A}$ the set of recurrent classes. We can thus focus our attention to homogeneous action profiles described in Lemma 5.3 as candidates for stochastically stable states, which includes the target profile a itself.

Every $a' \in \mathcal{A}_{\mathsf{R}}$ can be assigned a level of "disagreement" $d(a') := |\{L_z | a'_{L_z} \neq a_{L_z}, z \in \{x, y\}\}|$, corresponding to the number of homogeneous segments that differ relative to their counterparts in target profile a. The next result demonstrates that disagreement decreases along minimum resistance paths between recurrent classes.

Lemma 5.4. *Consider the directed graph $\Sigma = (\mathcal{A}_R, \mathcal{E})$ in which the edges $\mathcal{E}$ are formed by connecting recurrent classes through the minimum resistance edge leaving each class. Then Σ is composed of a collection of disconnected subgraphs $\Sigma_u = (\mathcal{A}_u, \mathcal{E}_u)$, each one corresponding to a particular recurrent class u. Each subgraph Σ_u has the following properties.*

- *The class u belongs to Σ_u, and for every node $v \in \mathcal{A}_u$, $v \neq u$, there is a unique path from v to u.*
- *There exists a class $v \in \mathcal{A}_u$ s.t. $(u, v), (v, u) \in \mathcal{E}_u$.*

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2020.3038842, IEEE Transactions on Control of Network Systems

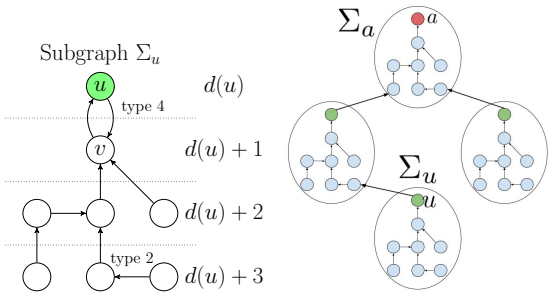

Fig. 4: (Left) The graph formed by connecting recurrent classes through the minimum resistance edge leaving each class, is composed of disconnected subgraphs Σ_u with the structure illustrated above (Lemma 5.4). Disagreement decreases along any path. (Right) The resistance tree rooted at a of minimum stochastic potential.

- $u = \arg\min_{v \in \mathcal{A}_u} d(v)$.

We refer to the class u as the head *of the subgraph Σ_u.*

An illustration of a subgraph Σ_u is shown in Figure 4 (left). Note that the rooted tree on a subgraph Σ_u with minimal stochastic potential is given by Σ_u without the edge (u, v). The next result asserts that the minimal resistance edge that leaves a subgraph leads to another subgraph whose head node has lower disagreement.

Lemma 5.5. *Consider the subgraph $\Sigma_u = (\mathcal{A}_u, \mathcal{E}_u)$ where $u \neq a$. Then there is an edge (u, v) starting at the head u leading to a class v of another subgraph $\Sigma_{u'}$ satisfying $(u, v) \in \arg\min_{a^1 \in \mathcal{A}_u, a^2 \notin \mathcal{A}_u} \rho_{a^1, a^2}$. Furthermore. $d(u') < d(u)$.*

These results give us enough structure about the resistance trees to deduce that a is the unique stochastically stable state.

Proposition 5.1. *The profile a is the unique stochastically stable state under an aggressive policy targeting a.*

The structure of the minimum potential rooted tree T is illustrated in Figure 4 (right). Proposition 5.1 is a sufficiency result – the aggressive policy is a dynamic policy that stabilizes the profile a. Our next result asserts necessity – the number of adversarial nodes employed by the aggressive policy is the minimum required budget to stabilize a.

Proposition 5.2. *A policy using fewer adversarial nodes than the aggressive policy targeting $a \in \mathcal{A}$ cannot stabilize a.*

Proof of Theorem 3.2. We have just established that an aggressive policy targeting a is the dynamic policy that stabilizes a with the fewest number of adversaries. By Lemma 5.2, a target profile must have x segments of length 2 or greater and y segments of length $\left\lceil \frac{2+\alpha}{1-\alpha} \right\rceil$ or greater. To implement the aggressive policy, there needs to be two adversaries for each y segment of any length, and two adversaries for each x segment of length no greater than $1+\left\lfloor \frac{1-\alpha}{\alpha} \right\rfloor$ when $\alpha < \frac{1}{2}$. When $\alpha \geq \frac{1}{2}$, no adversaries are needed on any x segment. Propositions 5.1 and 5.2 assert this amount of adversarial nodes are necessary and sufficient to stabilize the target profile. By Step (2B) from Theorem 3.1, the minimum efficiency target profile has at most two segment patterns. We thus obtain an integer optimization problem similar to (SI-OPT), except with the above constraints taken into account instead. ■

C. Proof of Theorem 2.2: Dynamic uninformed adversary

In this subsection, we derive the minimum efficiency a dynamic uninformed adversary can induce on k-connected ring graphs. Similar to the static uninformed case (Theorem 2.1), the dynamic uninformed adversary effectively can only select how many x and y adversaries to allocate at each time step, and cannot place them in a strategic manner. The major difference here is it can also randomize among influence sets by selecting different distributions at each time step. The idea of the proof is by being able to probabilistically attach an adversary to each agent in the network, the all y profile can be stabilized independently of the budget γ.

Let us define $\Pi^* \subset \Pi_{\text{D}}$ as a set of dynamic uninformed policies that have the following properties. Suppose $\pi \in \Pi^*$. Then

(a) $|S_x(t)| = 0$ for all $t = 0, 1, 2, \ldots$.
(b) for any $i \in \mathcal{N}$ and $t = 0, 1, 2, \ldots$, there exists a $0 \leq \tau < \infty$ such that $\Pr(i \in S_y(t+\tau)) > 0$.
(c) there exists a subset $T \subset \mathcal{N}$ of agents satisfying $\frac{|T|}{n} = \gamma' < \gamma$ for all n, such that $\Pr(i \in S_y(t)) = 1$ for all $i \in T$, $t = 0, 1, 2, \ldots$.

Property (a) asserts x adversaries are never utilized. Property (b) ensures any given agent is influenced by a y adversary infinitely often. Property (c) says the adversary determistically influences a fixed fraction of agents in the network. In a sense, policies belonging to Π^* are "partially static".

Lemma 5.6. *Under a policy $\pi \in \Pi^*$, the all-y and all-x profiles are the only two recurrent classes.*

To determine whether $\vec{x}$ or $\vec{y}$ is stochastically stable, we now calculate the minimum resistance path between $\vec{x}$ and $\vec{y}$ ($\rho_{\vec{x}\vec{y}}$). Suppose we are in $\vec{x}$. Consider a path in which k consecutive agents switch to y, and each agent that unilaterally switches is influenced by a y-adversary. This occurs with a non-zero probability due to property (b). Each unilateral switch in this sequence has a non-zero resistance. However after this sequence of k switches, every subsequent switch of neighboring agents has zero resistance. The total resistance (as long as $\alpha < \frac{1}{k}$) is therefore $\rho_{\vec{x}\vec{y}} = \sum_{i=1}^{k} [(1+\alpha)(2k-(i-1)) - i]$.

We can similarly calculate the minimum resistance path between $\vec{y}$ and $\vec{x}$ ($\rho_{\vec{y}\vec{x}}$). Starting from $\vec{y}$, we consider a path in which k consecutive agents switch to x, and none of the agents are influenced by a y-adversary (this occurs with non-zero probability). After this sequence of k switches, every subsequent switch of neighboring agents has zero resistance. Once this growing x segment neighbors a member of T, the node belonging to T will also switch with zero resistance as long as $\alpha < \frac{1}{k}$. The total resistance is therefore $\rho_{\vec{y}\vec{x}} = \sum_{i=1}^{k} [(2k-(i-1)) - (i-1)(1+\alpha)]_+$.

For $\vec{y}$ to be stochastically stable, the condition $\rho_{\vec{x}\vec{y}} < \rho_{\vec{y}\vec{x}}$ must hold. This yields $\alpha < \frac{|T|}{k(|T|+k)}$. As the number of nodes tends to infinity, so does $|T|$, and we are left with the condition $\alpha < \frac{1}{k}$. This yields an efficiency $\frac{1}{1+\alpha}$ regardless of the fractional budget γ. It remains to show any $\pi \in \Pi_{\text{DU}} \setminus \Pi^*$

cannot induce an efficiency less than $\frac{1}{1+\alpha}$. This part of the proof can be found in the online version.

VI. Simulations

In this section, we provide numerical simulations of log-linear learning dynamics for two of the four adversarial models: static and dynamic informed (SI and DI). Simulation results for the other two models are provided in the online version. We verify the tightness of the lower bounds given in Theorems 3.1 and 3.2. We simulate the dynamics on finite ring graphs, and compute the average efficiency the network experiences when the adversary implements an optimal policy. Our results are given in Table I. We observe that the average efficiency approaches the fundamental lower bounds as the size of the graphs are increased.

For SI, we first compute the minimum efficiency SSS given α, budget $\lfloor \gamma \cdot n \rfloor$, and network size n. We then determine an adversarial set S_{xy} that is necessary and sufficient to stabilize the SSS (according to the proof of Theorem 3.1). We then run the LLL dynamics initialized at the SSS and with the static S_{xy}. We perform a similar experiment for DI, where we implement the aggressive policy (Definition 3).

	SI adversary			DI adversary		
n/α	0.3	0.5	0.7	0.3	0.5	0.7
10	0.685	0.667	1	0.638	0.566	0.588
20	0.673	0.650	0.653	0.573	0.566	0.550
30	0.669	0.644	0.659	0.595	0.566	0.537
	Fundamental lower bound					
	0.662	0.640	0.647	0.569	0.542	0.525

TABLE I: Simulation results of log-linear learning for two adversarial models. Each entry represents the averaged efficiency over 30 repetitions, where each repetition consists of 10^6 time steps. We fix the budget as $\gamma = 0.6$ and the learning parameter $\beta = 25$.

VII. Conclusion

This paper investigated the susceptibility of distributed game-theoretic learning algorithms to adversarial influences. We considered a scenario of an adversary intent on maximally degrading a network system's performance guarantees associated with distributed learning algorithms. We asked 1) How susceptible are these algorithms to adversarial interference? In particular, this paper focused on one such algorithm, log-linear learning, that possesses nice properties in non-adversarial settings. 2) How does an adversary's sophistication and system-level knowledge impact the degradation that the adversary can do to the system? We studied both of these questions in the context of graphical coordination games.

In particular, we considered two levels of adversarial sophistication – static and dynamic policies – and two levels of information, informed and agnostic about network structure. In a static policy, the adversary cannot change its influence over time. The dynamic policies we considered are stationary, in which the adversary can respond to the current system state as it evolves. While both types of policies induce asymptotic outcomes characterized by stochastically stable states, it is of interest in future work to consider non-stationary adversarial policies. That is, how can adversaries exploit dynamic policies that may not induce a stable asymptotic outcome?

An important insight gleaned from our analysis is that an adversary with a low resource budget – described by the fraction of agents in the network it can influence – does not benefit as much from system-level information as it would from the ability to employ dynamic strategies. On the other hand, when the adversary's budget is high, the opposite conclusion holds. While the results in this paper are adversarial-centric, these findings provide insight as to what actions a system operator could take to best protect system behavior. For instance, our analysis can inform decisions of whether to obfuscate system-level information from potential adversarial actors, or to disable capabilities of a highly sophisticated attacker.

References

[1] B. Canty, P. N. Brown, M. Alizadeh, and J. R. Marden, "The impact of informed adversarial behavior in graphical coordination games," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 1923–1928.

[2] S. Martínez, J. Cortés, and F. Bullo, "Motion Coordination with Distributed Information," *Control Systems Magazine*, vol. 27, no. 4, pp. 75–88, 2007.

[3] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[4] M. Zhu and S. Martínez, "Distributed coverage games for mobile visual sensors (I): Reaching the set of Nash equilibria," in *Proceedings of the IEEE Conference on Decision and Control*, 2009, pp. 169–174.

[5] V. Ramaswamy and J. R. Marden, "A sensor coverage game with improved efficiency guarantees," in *Proceedings of the American Control Conference*, vol. 2016-July, 2016, pp. 6399–6404.

[6] A. Jadbabaie, J. Lin, and S. A. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules." *Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[7] F. Pasqualetti, F. Dorfler, and F. Bullo, "Attack detection and identification in cyber-physical systems," *IEEE Transactions on Automatic Control*, vol. 58, no. 11, pp. 2715–2729, Nov 2013.

[8] H. Fawzi, P. Tabuada, and S. Diggavi, "Secure estimation and control for cyber-physical systems under adversarial attacks," *IEEE Transactions on Automatic Control*, vol. 59, no. 6, pp. 1454–1467, 2014.

[9] S. Sundaram and B. Gharesifard, "Distributed optimization under adversarial nodes," *IEEE Transactions on Automatic Control*, vol. 64, no. 3, pp. 1063–1076, 2019.

[10] D. H. Wolpert and K. Tumer, "Optimal payoff functions for members of collectives," in *Modeling Complexity in Economic and Social Systems*. World Scientific, 2002, pp. 355–369.

[11] R. D. McKelvey and T. R. Palfrey, "Quantal response equilibria for normal form games," *Games and Economic Behavior*, vol. 10, no. 1, pp. 6–38, 1995.

[12] L. E. Blume, "The Statistical Mechanics of Best-Response Strategy Revision," *Games and Economic Behavior*, vol. 11, no. 2, pp. 111–145, 1995.

[13] J. R. Marden and A. Wierman, "Distributed Welfare Games," *Operations Research*, vol. 61, pp. 155–168, 2013.

[14] Y. Lim and J. Shamma, "Robustness of stochastic stability in game theoretic learning," in *American Control Conference (ACC)*, 2013, pp. 6160–6165.

[15] S. Rahili and W. Ren, "Game theory control solution for sensor coverage problem in unknown environment," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 1173–1178.

[16] C. Sun, "A time variant log-linear learning approach to the set k-cover problem in wireless sensor networks," *IEEE Transactions on Cybernetics*, vol. 48, no. 4, pp. 1316–1325, 2018.

[17] A. Kanakia, B. Touri, and N. Correll, "Modeling multi-robot task allocation with limited information as global game," *Swarm Intelligence*, vol. 10, no. 2, pp. 147–160, 2016.

[18] D. Shah and J. Shin, "Dynamics in Congestion Games," *Proc. ACM SIGMETRICS'10*, vol. 38, p. 107, 2010.

[19] M. Kearns, M. L. Littman, and S. Singh, "Graphical Models for Game Theory," *Proceedings of the 17th conference on Uncertainty in artificial intelligence*, no. Owen, pp. 253–260, 2001.

[20] A. Montanari and A. Saberi, "The spread of innovations in social networks." *Proceedings of the National Academy of Sciences of the United States of America*, vol. 107, no. 47, pp. 20 196–20 201, 2010.

[21] H. P. Young, *Individual strategy and social structure: An evolutionary theory of institutions*. Princeton University Press, 2001.

[22] B. Pradelski and H. P. Young, "Learning Efficient Nash Equilibria in Distributed Systems," *Games and Economic Behavior*, vol. 75, no. 2, pp. 882–897, 2012.

[23] C. Alós-Ferrer and N. Netzer, "The logit-response dynamics," *Games and Economic Behavior*, vol. 68, no. 2, pp. 413–427, 2010.

[24] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, 1982.

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

[26] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 4, pp. 766–781, April 2013.

[27] H. Jaleel, W. Abbas, and J. S. Shamma, "Robustness of stochastic learning dynamics to player heterogeneity in games," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, Dec 2019, pp. 5002–5007.

[28] L. Su and N. Vaidya, "Multi-agent optimization in the presence of byzantine adversaries: Fundamental limits," in *2016 American Control Conference (ACC)*, 2016, pp. 7183–7188.

[29] P. N. Brown, H. P. Borowski, and J. R. Marden, "Security against impersonation attacks in distributed systems," *IEEE Transactions on Control of Network Systems*, vol. 6, no. 1, pp. 440–450, 2018.

[30] D. Monderer and L. S. Shapley, "Potential Games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.

[31] L. E. Blume, "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, vol. 5, no. 3, pp. 387–424, 1993.

[32] J. R. Marden and J. S. Shamma, "Revisiting Log-Linear Learning: Asynchrony, Completeness and a Payoff-based Implementation," *Games and Economic Behavior*, vol. 75, no. 4, pp. 788–808, 2012.

[33] H. P. Young, "The Evolution of Conventions," *Econometrica*, vol. 61, no. 1, pp. 57–84, 1993.

Keith Paarporn received his B.S. in Electrical Engineering from the University of Maryland, College Park in 2013, his M.S. in Electrical and Computer Engineering from the Georgia Institute of Technology in 2016, and his Ph.D. in Electrical and Computer Engineering from the Georgia Institute of Technology in 2018. He is currently a postdoctoral scholar in the Electrical and Computer Engineering Department at the University of California, Santa Barbara. His research interests include game theory, control, and multi-agent systems.

Brian Canty is currently at CACI. He received the B.S. and M.S. degrees in Electrical and Computer Engineering at University of California, Santa Barbara in 2018 and 2019, respectively.

Philip N. Brown Philip N. Brown is an Assistant Professor in the Department of Computer Science at the University of Colorado Colorado Springs. Philip received the Bachelor of Science in Electrical Engineering in 2007 from Georgia Tech, after which he spent several years designing control systems and process technology for the biodiesel industry. He received the Master of Science in Electrical Engineering in 2015 from the University of Colorado at Boulder under the supervision of Jason R. Marden, where he was a recipient of the University of Colorado Chancellor's Fellowship. He received the PhD in Electrical and Computer Engineering from the University of California, Santa Barbara under the supervision of Jason R. Marden. He was finalist for the Best Student Paper Award at the 2016 and 2017 IEEE Conferences on Decision and Control, and received the UCSB Center for Control, Dynamical Systems, and Computation 2018 Best PhD Thesis Award. Philip is interested in the interactions between engineered and social systems.

Mahnoosh Alizadeh received the B.Sc. degree in electrical engineering from the Sharif University of Technology in 2009, and the M.Sc. and Ph.D. degrees in electrical and computer engineering from the University of California at Davis in 2013 and 2014, respectively. She is an Assistant Professor of electrical and computer engineering with the University of California at Santa Barbara. From 2014 to 2016, she was a Post-Doctoral Scholar with Stanford University. Her research interests are focused on designing scalable control and market mechanisms for enabling sustainability and resiliency in societal infrastructures, with a particular focus on demand response and electric transportation systems. She was a recipient of the NSF CAREER Award.

Jason Marden is an Associate Professor in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara. Jason received a BS in Mechanical Engineering in 2001 from UCLA, and a PhD in Mechanical Engineering in 2007, also from UCLA, under the supervision of Jeff S. Shamma, where he was awarded the Outstanding Graduating PhD Student in Mechanical Engineering. After graduating from UCLA, he served as a junior fellow in the Social and Information Sciences Laboratory at the California Institute of Technology until 2010 when he joined the University of Colorado. Jason is a recipient of the NSF Career Award (2014), the ONR Young Investigator Award (2015), the AFOSR Young Investigator Award (2012), the American Automatic Control Council Donald P. Eckman Award (2012), and the SIAG/CST Best SICON Paper Prize (2015). Jason's research interests focus on game theoretic methods for the control of distributed multiagent systems.