Can Competition Outperform Collaboration? the Role of Misbehaving Agents

Can Competition Outperform Collaboration? The Role of Misbehaving Agents

Luca Ballotta, *Member, IEEE*, Giacomo Como, *Member, IEEE*, Jeff S. Shamma, *Fellow, IEEE*, and Luca Schenato, *Fellow, IEEE*

***Abstract*—We investigate a novel approach to resilient distributed optimization with quadratic costs in a multiagent system prone to unexpected events that make some agents misbehave. In contrast to commonly adopted filtering strategies, we draw inspiration from phenomena modeled through the Friedkin–Johnsen dynamics and argue that adding competition to the mix can improve resilience in the presence of misbehaving agents. Our intuition is corroborated by analytical and numerical results showing that 1) there exists a nontrivial tradeoff between full collaboration and full competition and 2) our competition-based approach can outperform state-of-the-art algorithms based on weighted mean subsequence reduced. We also study the impact of communication topology and connectivity on resilience, pointing out insights into robust network design.**

***Index Terms*—Friedkin–Johnsen (FJ) model, misbehaving agents, multiagent systems, resilient consensus.**

I. Introduction

WITH great power comes great responsibility, and networked systems are powerful, indeed. From smart grids managing energy consumption [1], [2] to sensors monitoring vast areas [3] and to autonomous cars for intelligent mobility [4], [5], everyday life relies evermore on the control of connected devices.

While this brings numerous benefits, a major drawback is that malicious agents can locally intrude from any point in the system and cause serious damage at global scale. Recently, Department of Energy secretary stated that enemies of the United States can shut down the U.S. power grid, and it is known that hacking groups around the world have high technological sophistication [6]. Cyberattacks hit Italian health care infrastructures during the COVID-19, disrupting services for weeks [7]. Another concern is accidental failures spreading from single source nodes. Cascading failure damages have notable examples, from city-wide electricity blackouts to denial of service of web applications. Furthermore, as new frontiers through massively connected devices in networked control systems are breached, thanks to powerful communication protocols such as 5G, this problem will only gain in importance.

A. Related Literature

The aforementioned problems have been extensively studied in the literature. A body of work investigates control techniques to overcome the fragility of specific applications. Examples are power outages in smart grids [8], [9], cascading failures in cyber-physical systems [10], [11], [12], [13], denial of service [14], [15], robot gathering [16], and distributed estimation [17], to name a few. From a methodological perspective, control and optimization literature mostly focuses on the robustness of distributed algorithms and control protocols to a fraction of misbehaving agents. This approach can tailor either intentionally malicious agents, such as cyber-attackers, or accidental faults due to, e.g., hardware damage. A fundamental subclass of such approaches is *resilient consensus*, aimed at enforcing consensus of normally behaving (or *regular*) agents in the face of unknown adversaries. The consensus problem has been deeply studied in the past decades [18] and underlies a plethora of application domains. In particular, *average consensus* is a cornerstone in distributed estimation [17] and optimization [19], [20], [21], management of power grids [22], distributed federated learning [23], [24], among others. Unfortunately, the standard consensus protocol is fragile and misbehaving agents can arbitrarily deviate from the system trajectory. To tame this issue, the most common approaches rely on the filtering strategy referred to as "mean subsequence reduced" (MSR), whereby agents discard suspicious messages (largest and smallest values) from updates [25]. The pioneering paper [26] introduced a weighted version (W-MSR) and defined the r-*robustness* of graphs, a suitable index that enables theoretical guarantees for resilient consensus based on W-MSR. Among the many variants and adaptations of W-MSR, Dibaji and Ishii [27] study resilient control for double integrators, Wang et al. [28] tackle mobile adversaries, Usevitch and Panagou [29] focus on the leader–follower framework, Shang [30] targets nonlinear systems with state

Manuscript received 26 June 2023; accepted 28 October 2023. Date of publication 3 November 2023; date of current version 29 March 2024. This work was supported in part by the Italian Ministry of Education, University and Research (MIUR) under the PRIN Project 2017NS9FEY entitled "Realtime Control of 5G Wireless Networks," under the PRIN Project 2017 "Advanced Network Control of Future Smart Grids," and through the initiative "Departments of Excellence" (Law 232/2016), and in part by the Compagnia di San Paolo. Recommended by Associate Editor Z. Li. *(Corresponding author: Luca Ballotta.)*

Luca Ballotta and Luca Schenato are with the Department of Information Engineering, University of Padova, 35131 Padova, Italy (e-mail: ballotta@dei.unipd.it; schenato@dei.unipd.it).

Giacomo Como is with the Department of Mathematical Sciences, Politecnico di Torino, 10129 Torino, Italy (e-mail: giacomo.como@polito.i).

Jeff S. Shamma is with the Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana–Champaign, Urbana, IL 61801 USA (e-mail: jshamma@illinois.edu).

Digital Object Identifier 10.1109/TAC.2023.3329850

constraints, Wen et al. [31] extend the notion of r-robustness to time-varying graphs, and the authors in [32], [33], and [34] consider generic cost functions to achieve resilience in general distributed optimization.

Other approaches in the literature do not filter information from neighbors but explore the enhanced capabilities of regular agents. For example, Dibaji et al. [35] use a buffer to store all values received from other agents and replace the thresholding mechanism with a voting strategy followed by dynamical updates, Abbas et al. [36] study algorithmic robustness enabled by trusted agents, Zhai et al. [37] propose dynamically switching update rule for continuous-time double integrators, and the authors in [38], [39], and [40] use stochastic or heuristic trust scores to filter out potentially malicious transmissions, providing probabilistic bounds on detection, convergence, or deviation from average consensus. While such approaches may overcome limitations of MSR-based strategies, they usually require either stronger assumptions on the network (e.g., trusted agents) or burdening local computation or storage resources.

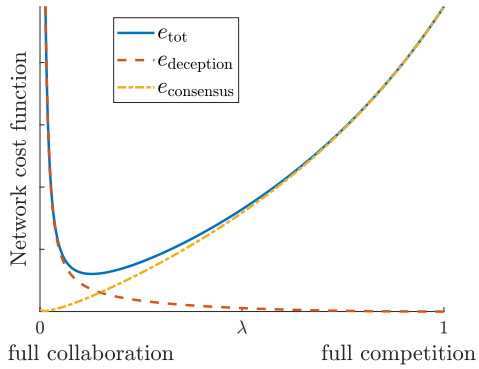

Fig. 1. Competition versus collaboration in distributed quadratic optimization. The global cost e_{tot} is the sum of two contributions that reflect two contrasting attitudes of regular agents: $e_{\text{deception}}$ is caused by (erroneously) trusting misbehaving agents, which makes them drift away from the nominal average while $e_{\text{consensus}}$ is due to the competition among regular agents, which mitigates misbehaviors but also prevents regular agents from reaching a consensus. The tunable parameter $\lambda \in [0, 1]$ allows regular agents to smoothly transition from *full collaboration* ($\lambda = 0$), where they fully trust all agents in the network, to *full competition* ($\lambda = 1$), where they trust no other agent, producing a rich range of behaviors at local and global scales.

B. Novel Contribution

Despite the success of MSR-based strategies, a critical point is the dependence of theoretical guarantees on the r-robustness of the underlying graph, which allows regular agents to reach resilient consensus if such an index is large enough. In fact, it is difficult to characterize the steady-state behavior of agents if some minimal robustness is not met. Even though algorithms might practically work, comprehensive theoretical guarantees are still lacking, and also, some applications require more conservative but safer approaches. In particular, while in some cases, agents may just agree on a common value, other tasks require *average consensus* to succeed. Thus, we depart from classical filtering strategies and seek a framework for resilience that can offer theoretical guarantees in a broader sense.

Toward this goal, we set the stage with two key moves. First, rather than finding conditions that enforce a consensus among regular agents, which only indicates if a system *is resilient*, we aim to measure the *level of resilience*, which we evaluate through the cost of a distributed optimization problem with quadratic costs. Second, we aim to modify the original problem to make it robust to misbehaving agents rather than adapting a consensus protocol. Stepping forward, we propose an update rule based on the celebrated Friedkin–Johnsen (FJ) dynamics [41] to enhance the resilience of the addressed distributed optimization problem. The key feature of the FJ dynamics is a tunable parameter $\lambda \in [0, 1]$ that allows us to smoothly transition from the regime of *full collaboration* ($\lambda = 0$), where each regular agent equally trusts all agents, to the regime of *full competition* ($\lambda = 1$), where each regular agent regards all others as adversaries. We refer to the regime with $\lambda \in (0, 1]$ as *competition-based* because regular agents are forced to (partially) mistrust the others. This approach allows us to study resilience variations that arise from different choices of agents that can trust their neighbors or not, a choice that turns out to be crucial if adversaries are present. In fact, we observe a fundamental performance tradeoff that we name *competition–collaboration tradeoff:* In general, *the optimal resilient strategy is hybrid, namely each regular agent should partially compete with its neighbors*, as depicted in Fig. 1. The global cost (solid blue) is the sum of two conflicting contributions that represent deception due to collaboration with misbehaving agents (dashed red) and inefficiency caused by competition against regular agents (dashed-dotted yellow). To achieve analytical intuition about such a competition–collaboration tradeoff, we leverage the *social power*, a tool drawn from opinion dynamics that sheds light on the twofold effect of the parameter λ used to instantiate the FJ dynamics.

After analytically characterizing the proposed competition-based protocol, we fix the update rule and shift attention to the network in order to assess how it impacts the resilience of regular agents. In particular, we numerically show how network connectivity can mitigate misbehavior and how the performance varies as the network gets sparser or less balanced. In fact, we heuristically observe that not only high connectivity but also degree balance across agents is useful to tame unknown adversaries, which intuitively can exploit highly connected areas to quickly spread damage at a global level.

Besides new results, this article extends the preliminary conference version [42] in two ways. First, we consider a more general prior distribution of the observations of agents. Second, we compare our proposed strategy with both standard W-MSR [26] and recently proposed SABA [35].

C. Organization of the Article

We motivate average consensus for distributed optimization in Section II, model a class of adversaries in Section II-A, and introduce the performance metric used to quantify resilience in Section II-B. In Section III, we propose our competition-based

protocol: We introduce the FJ dynamics in Section III-A, compute the cost function in Section III-B, and formally characterize the cost function and its minimizer in Sections III-C and III-D. In Section IV, we report numerical tests that support our analytical intuition. Then, in Section IV-A, we offer analytical insight on the competition–collaboration tradeoff using the notion of social power. In Section V, we numerically explore the impact of the communication network on resilience. To evaluate our approach, we perform simulations in Section VI and show that it can outperform MSR-based methods. Finally, Section VII concludes this article.

II. Setup and Problem Formulation

We consider a multiagent system composed of N agents labeled as the set $\mathcal{V} = \{1, \dots, N\}$. Each agent $i \in \mathcal{V}$ carries local information encoded by an *observation* $\theta_i \in \mathbb{R}$ and a variable *state* $x_i \in \mathbb{R}$. For notation convenience, we stack all states and observations in the vectors $x \in \mathbb{R}^N$ and $\theta \in \mathbb{R}^N$, respectively.

Within the network, some agents behave according to a control task at hand while others cannot be controlled and may deviate from the task. We call the former agents *regular* and the latter agents *misbehaving*. Because the misbehaving agents cannot be involved in cooperative tasks through their uncontrolled nature, we consider a distributed optimization problem involving only the regular agents. We assume that each regular agent wishes to adjust its state so as to minimize a quadratic mismatch among all observations

$$f_{\text{local}}(x_i) \doteq \sum_{j \in \mathcal{R}} (x_i - \theta_j)^2, \quad i \in \mathcal{R} \tag{1}$$

where $\mathcal{R} \subseteq \mathcal{V}$ gathers all regular agents. By straightforward calculations, (1) can be rewritten as

$$f_{\text{local}}(x_i) = R\left(x_i - \bar{\theta}_{\mathcal{R}}\right)^2 - R\bar{\theta}_{\mathcal{R}}^2 + \sum_{j \in \mathcal{R}} \theta_j^2 \tag{2}$$

where $R \doteq |\mathcal{R}|$ and $\bar{\theta}_{\mathcal{R}}$ is the average of observations $\{\theta_i\}_{i \in \mathcal{R}}$.

The distributed optimization task is then given by

$$\arg\min_x \frac{1}{R} \sum_{i \in \mathcal{R}} f_{\text{local}}(x_i) = \arg\min_x \sum_{i \in \mathcal{R}} \left(x_i - \bar{\theta}_{\mathcal{R}}\right)^2 \tag{3}$$

which is solved if and only if all regular agents reach *average consensus* among them, i.e., $x_i = \bar{\theta}_{\mathcal{R}}$ for all $i \in \mathcal{R}$.

In the nominal scenario where all agents are regular ($\mathcal{V} = \mathcal{R}$), the cost (3) can be minimized via the *consensus dynamics* (or *consensus protocol*) $x(k+1) = W^o x(k)$ where $x(0) = \theta$ and W^o is a doubly stochastic irreducible matrix that leads agents to average consensus. Interpreting W^o as a (weighted) communication matrix, the consensus dynamics allows agent j to communicate its state to agent i if and only if $W_{ij} > 0$.

However, the standard consensus protocol easily fails in the presence of misbehaving agents [26]. We next introduce a misbehavior model that disrupts the nominal protocol.

A. Misbehaving Agents

Misbehaving agents follow state trajectories with no relation to optimization task (3) and broadcast potentially misleading information to neighbors. We denote the subset of misbehaving agents by $\mathcal{M}$ with $M \doteq |\mathcal{M}|$, $\mathcal{V} = \mathcal{M} \cup \mathcal{R}$, and $\mathcal{M} \cap \mathcal{R} = \emptyset$. Also, without loss of generality, we label the agents as $\mathcal{R} = \{1, \dots, R\}$ and $\mathcal{M} = \{R+1, \dots, N\}$. The vectors $x_{\mathcal{R}} \in \mathbb{R}^R$ and $x_{\mathcal{M}} \in \mathbb{R}^M$ stack the states of regular and misbehaving agents, respectively, with $x^\top = [x_{\mathcal{R}}^\top \; x_{\mathcal{M}}^\top]$.

To address a general scenario and remove dependence on the specific values of observations, we assume that these are drawn from a prior distribution.

Assumption 1 (Distribution of observations): Observations $\{\theta_i\}_{i \in \mathcal{V}}$ are distributed as random variables with mean $\mathbb{E}[\theta] = 0$ and covariance matrix $\Sigma \doteq \mathbb{E}[\theta\theta^\top] \succ 0$. We denote $\Sigma_{ii} = \sigma_i^2$ and $\Sigma_{ij} = \sigma_{ij}\, \forall i \neq j$.

While the standard consensus and Assumption 1 are suited to an ideal scenario, misbehaving agents may disrupt the task (3). In the following, we assume that misbehaving agents constantly transmit noisy versions of their observations:

$$x_m(k) = \theta_m + v_m + n_m(k) \quad \forall m \in \mathcal{M}. \tag{4}$$

We refer to the constant input v_m as *(deception) bias* and to the varying input $n_m(k)$ as *(deception) noise*. In other words, the deception bias v_m makes the observation θ_m of the misbehaving agent m an outlier w.r.t. the expected range of values of observations as per Assumption 1. Conversely, the deception noise $n_m(k)$ hides the true state of the misbehaving agent from its neighbors, akin to purposely injected measurement noise.

Assumption 2 (Misbehavior model): We stack biases in the vector $v \in \mathbb{R}^M$ and noises in the vector $n(k) \in \mathbb{R}^M$. Furthermore, we set their statics as $\mathbb{E}[v] = 0$, $\mathbb{E}[vv^\top] = V \succeq 0$, $\mathbb{E}[v\theta^\top] = 0$, $\mathbb{E}[n(k)] = 0$, $\mathbb{E}[n(k)n^\top(h)] = \delta_{kh}Q \succeq 0$, and $\mathbb{E}[n(k)\theta^\top] = 0\, \forall k, h \geq 0$, where $\delta_{kh} = 1$ if $k = h$ and $\delta_{kh} = 0$ otherwise.

Remark 1 (Misbehavior versus intelligent attacks): Assumption 2 is consistent with a portion of the literature on resilient consensus, where algorithms are tested against constant or drifting misbehaving agents that steer their neighbors far off the nominal consensus [28], [35], [38], [40]. In our case, misbehaving agents are stubborn on average but behave in a less trivial (noisy) way. On the other hand, smart (malicious) adversaries may need to be contrasted by sophisticated strategies [43], [44]. This case is outside the scope of this article, where we explore competition as a tool to enhance resilience, and we defer a comprehensive study with intelligent attacks to future work.

B. Performance Metric

In light of problem (3) and assuming that the states of regular agents are updated by a control protocol overtime, we use the following performance metric to measure the resilience of the system, which we refer to as *(average) consensus error*:

$$e_{\mathcal{R}} \doteq \lim_{k \to +\infty} \mathbb{E}\left[\sum_{i \in \mathcal{R}} \left(x_i(k) - \bar{\theta}_{\mathcal{R}}\right)^2\right]. \tag{CE}$$

The error (CE) coincides with the objective cost of the optimization problem (3) (up to additive constants that depend only on the observations) averaged over the stochastic elements

within the system dynamics, such as observations θ of all agents and deception biases v and noises $n(k)$ of misbehaving agents.

While the standard consensus protocol achieves $e_{\mathcal{R}} = 0$ in the nominal scenario where all agents reach average consensus, the presence of unknown misbehaving agents makes $e_{\mathcal{R}}$ grow, degrading the collaborative task (3). In the following section, we propose an update protocol that makes regular agents more resilient by decreasing the consensus error $e_{\mathcal{R}}$ and, hence, improving the performance associated with the task (3).

III. Resilient Average Consensus

A. The Friedkin-Johnsen Dynamics

Because the classical consensus is fragile to misbehaving agents, we look for alternative strategies to minimize (CE).

To this aim, we step back to the optimization problem (3) and search for a way to make it more robust to unexpected behaviors. In particular, we modify the local problems associated with each regular agents $i \in \mathcal{R}$ by integrating the nominal weight matrix W^o and adding a regularization term that penalizes deviations from the local observation

$$\tilde{f}_{\text{local}}(x_i) = \lambda (x_i - \theta_i)^2 + (1-\lambda) \sum_{j \in \mathcal{N}_i} W^o_{ij} (x_i - x_j)^2 . \quad (5)$$

Assumption 3 (Nominal weights): The matrix W^o is irreducible, row stochastic, and $W^o_{ii} = 0, i \in \mathcal{V}$ (no self-loops).

The parameter $\lambda \in [0,1]$ in (5) makes the ith agent anchor to its observation θ_i, so that large deviations of its state x_i from θ_i are discouraged. We then let each agent greedily minimize the modified cost (5) at step $k+1$, which yields the celebrated FJ dynamics [41]

$$x_i(k+1) = \lambda \theta_i + (1-\lambda) \sum_{j \in \mathcal{N}_i} W^o_{ij} x_j(k) \qquad \text{(FJ)}$$

with $x_i(0) = \theta_i$. We interpret the previous rule as a modified consensus protocol where the agents do not fully align with neighbors but also *compete* by tracking their own observations. In particular, we call the parameter λ as *competition*, referring to the case $\lambda = 0$ (equivalent to the consensus protocol) as *full collaboration* and to the case $\lambda = 1$ as *full competition*.

While the dynamics (FJ) is suboptimal if all agents are collaborative because it prevents them from reaching a consensus if $\lambda > 0$, we use it to make regular agents resilient to unknown misbehaving agents. Intuitively, anchoring a regular agent $i \in \mathcal{R}$ to its observation θ_i prevents the agent from being arbitrarily dragged away by misleading values coming from misbehaving agents. In particular, the latter agents obey (4) with no relation to the protocol (FJ) or nominal weights W^o.

In the following, we study how the protocol (FJ) improves system resilience. In fact, tuning λ within the interval $[0,1]$ originates a nontrivial *competition–collaboration tradeoff:* What is the optimal competition λ that makes regular agents most resilient with respect to task (3)? Exploring this tradeoff under misbehaving agents is the main matter of investigation of this article. To this aim, we regard the consensus error as a function of the competition: This allows us to perform analysis and achieve insight about the minimization of $e_{\mathcal{R}}(\lambda)$.

Remark 2 (Connections with game theory and opinion dynamics): The FJ dynamics can be given the following game-theoretic interpretation. The cost (5) is interpreted in games as *cognitive dissonance*, whereby a rational decision-maker gets an incentive both in aligning with the neighbors and in following a local rule. Also, the function (5) with $\lambda = 0$ reduces to the utility used in [45] where Marden et al. analyze the consensus protocol from a game-theoretic perspective. In opinion dynamics, the FJ dynamics is typically used to model *prejudice*, whereby the opinion of an agent is biased toward a personal belief despite interactions with others.

Remark 3 (Competition for resilience): While most works in the literature regard λ as a model parameter, *we purposely design* λ in (FJ). The intuition behind this choice, seemingly counterintuitive for collaborative tasks, is that introducing some competition among agents can mitigate behaviors that are unpredictable at the design stage: Rather than addressing the *binary property* "consensus is (not) achieved" like typical works on resilient consensus, we take a broader viewpoint and interpret resilience as a *real quantity* measured through the cost (CE).

Remark 4 (Heterogeneous competition): While we focus on a single parameter λ for the sake of analysis, the general FJ model with a different parameter λ_i for each agent $i \in \mathcal{V}$ makes the analysis challenging but does not affect the fundamental system behavior. Designing a parameter λ_i for each regular agent $i \in \mathcal{R}$ to improve performance even further is an important topic, whose investigation is left to future work.

B. Computation of the Consensus Error

We now compute the error (CE) with the steady state induced by the dynamics (FJ). To this aim, it is convenient to write the network dynamics associated with all regular agents.

First, we highlight the interactions of regular and misbehaving agents by partitioning the nominal weight matrix as

$$W^o = \left[\begin{array}{c|c} W_{\mathcal{R}} & W_{\mathcal{M}} \\ \hline * & * \end{array}\right] \quad W_{\mathcal{R}} \in \mathbb{R}^{R \times R}, W_{\mathcal{M}} \in \mathbb{R}^{R \times M}. \quad (6)$$

Then, the dynamics of regular agents can be written as follows:

$$\begin{aligned} x_{\mathcal{R}}(k+1) &= A x_{\mathcal{R}}(k) + B x_{\mathcal{M}}(k) + \lambda \theta_{\mathcal{R}} \\ A &\doteq (1-\lambda) W_{\mathcal{R}}, \quad B \doteq (1-\lambda) W_{\mathcal{M}}. \end{aligned} \quad (7)$$

If $(1-\lambda)W_{\mathcal{R}}$ is Schur stable, which happens if the graph is connected, at the steady state, the dynamics (7) induce the following distribution w.r.t. the deception noises $n(k)$:

$$\bar{x}_{\mathcal{R}} \doteq \lim_{k \to +\infty} \mathbb{E}_n \left[x_{\mathcal{R}}(k) \right], \quad P \doteq \lim_{k \to +\infty} \text{Var}_n \left(x_{\mathcal{R}}(k) \right). \quad (8)$$

Defining $S_{\mathcal{R}} \doteq [I_R \mid 0]$, the quantities above amount to

$$\bar{x}_{\mathcal{R}} = S_{\mathcal{R}} L \left(\theta + S_{\mathcal{M}}^\top v \right), \quad L \doteq \left(I - (1-\lambda) W \right)^{-1} \lambda \quad (9)$$

$$P = (1-\lambda)^2 W_{\mathcal{R}} P W_{\mathcal{R}}^\top + (1-\lambda)^2 W_{\mathcal{M}} Q W_{\mathcal{M}}^\top \quad (10)$$

where the matrix W encodes the actual interactions (weights) followed within the network and is defined as follows:

$$W = \left[\begin{array}{c|c} W_{\mathcal{R}} & W_{\mathcal{M}} \\ \hline 0 & I_M \end{array}\right]. \tag{11}$$

In particular, W means that the average state of a misbehaving agent is affected by no other agent, according to (4). The matrix L is row-stochastic with algebraic multiplicity of the eigenvalue 1 equal to $M+1$ and does not induce a consensus.

Let $\mathbb{E}_{x|y}[z] \doteq \mathbb{E}_x[z|y]$ and let $\mathbb{1} \in \mathbb{R}^R$ denote the vector of all ones in $\mathbb{R}^R$. From (CE) and (7)–(10), it follows:

$$\begin{aligned} e_{\mathcal{R}} &= \lim_{k\to+\infty} \mathbb{E}_{\theta,v,n}\left[\left\|x_{\mathcal{R}}(k) - \mathbb{1}\bar{\theta}_{\mathcal{R}}\right\|^2\right] \\ &= \lim_{k\to+\infty} \mathbb{E}_{\theta,v}\left[\mathbb{E}_n\left[\left\|x_{\mathcal{R}}(k) - \mathbb{1}\bar{\theta}_{\mathcal{R}}\right\|^2 |\theta, v\right]\right] \\ &= \mathbb{E}_{\theta,v}\left[\left\|\bar{x}_{\mathcal{R}} - \mathbb{1}\bar{\theta}_{\mathcal{R}}\right\|^2\right] + \lim_{k\to+\infty} \operatorname{var}_n\left(x_{\mathcal{R}}(k) - \mathbb{1}\bar{\theta}_{\mathcal{R}}\right). \end{aligned} \tag{12}$$

Then, standard calculations allow us to rewrite the consensus error as follows:

$$e_{\mathcal{R}}(\lambda) = e_{\mathcal{R},v}(\lambda) + e_{\mathcal{R},n}(\lambda) + \kappa \tag{13}$$

where κ does not depend on λ and

$$e_{\mathcal{R},v}(\lambda) \doteq \operatorname{Tr}\left(\widetilde{\Sigma} E^\top E\right), \quad e_{\mathcal{R},n}(\lambda) \doteq \operatorname{Tr}(P) \tag{14}$$

where we define $\widetilde{\Sigma} = \Sigma + S_{\mathcal{M}} S_{\mathcal{M}}^\top V$, $E = S_{\mathcal{R}} L - C_R S_{\mathcal{R}}$, and $C_R \doteq \frac{1}{R}\mathbb{1}\mathbb{1}^\top$. The expression (13) highlights that the two features of the misbehavior modeled in Assumption 2 generate two different contributions to the consensus error. The error term $e_{\mathcal{R},v}$ is caused by the biased observations of misbehaving agents $\theta + v$ that are constantly injected into the dynamics (7). Instead, the error term $e_{\mathcal{R},n}$ is produced by the deception noises $n(k)$ that make the steady-state drift away.

Lemma 1 (Drift versus competition): The error term $e_{\mathcal{R},n}(\lambda)$ is strictly decreasing with λ and $e_{\mathcal{R},n}(1) = 0$.

Proof: See Appendix B. □

In other words, Lemma 1 implies that setting $\lambda > 0$ makes regular agents more resilient to the deception noise as opposed to the standard consensus protocol. This observation relates to the work in [46] where the authors observe that even small perturbations of a row-stochastic matrix W can result in a large norm of the matrix difference and change of the Perron–Frobenius eigenvector.

C. Competition–Collaboration Tradeoff

To study how our proposed approach performs in the presence of misbehaving agents, we first compare the two extreme cases of full collaboration and full competition to see when the former approach should be ruled out by default.

Proposition 1 (Full competition versus full collaboration): In the presence of misbehaving agents, the dynamics (FJ) with $\lambda = 1$ yields a smaller error than with $\lambda = 0$ if and only if

$$\begin{aligned} \frac{M^2}{R} e_{\mathcal{R},n}(0) + \operatorname{Tr}(V) \geq \frac{M^2}{R}\operatorname{Tr}(\Sigma_{11}) - \frac{2M^2}{R^2}\mathcal{B}(\Sigma_{11}) \\ + \frac{2M}{R}\mathcal{B}(\Sigma_{12}) - \mathcal{B}(\Sigma_{22}) \end{aligned} \tag{15}$$

where $\mathcal{B}(A) = \sum_{i,j} A_{ij}$.

Proof: The statements follow from manipulations of the consensus errors induced by the two considered instantiations of (FJ). The full derivation is reported in Appendix C. □

In other words, Proposition 1 implies that the fully competitive approach outperforms the consensus protocol as soon as the misbehavior disturbances are sufficiently intense compared to the prior correlations between regular and misbehaving agents.

After acknowledging that the proposed competition-based approach can be more resilient than the standard consensus protocol in the presence of misbehaving agents, we now turn to study the optimal resilient strategy. In other words, we are interested in choosing λ so as to reduce the consensus error. In particular, we address the optimal competition λ^*

$$\lambda^* \in \arg\min e_{\mathcal{R}}(\lambda). \tag{16}$$

Such an optimal parameter exists by the Weierstrass theorem because $e_{\mathcal{R}}(\lambda)$ is continuous in (0,1] and has a continuous extension at $\lambda = 0$ through the extended continuity of L [47].

The next result describes when the optimal competition is nontrivial, meaning that the regular agents should compete against their neighbors in order to minimize the error (CE).

Theorem 1 (Competition–collaboration tradeoff.): Let $\Gamma \doteq \lim_{\lambda\to 0^+} \frac{dL}{d\lambda}$ with block partition

$$\Gamma = \left[\begin{array}{c|c} \Gamma_1 & \Gamma_2 \\ \hline 0 & 0 \end{array}\right], \quad \Gamma_1 \in \mathbb{R}^{R\times R},\ \Gamma_2 \in \mathbb{R}^{R\times M} \tag{17}$$

and $C_{RM} \doteq \frac{\mathbb{1}_R \mathbb{1}_M^\top}{M}$. If either of the following conditions holds:

C1. Σ is diagonal;

C2. W^o is symmetric and

$$\begin{aligned} -\frac{\mathrm{d}e_{\mathcal{R},n}(0)}{\mathrm{d}\lambda} - \operatorname{Tr}\left(V\Gamma_2^\top C_{RM}\right) > \operatorname{Tr}\left(-\Sigma_{11}\Gamma_1^\top C_R \right. \\ \left. -\Sigma_{12}\Gamma_2^\top C_R + \Sigma_{12}^\top \Gamma_1^\top C_{RM} + \Sigma_{22}\Gamma_2^\top C_{RM}\right) \end{aligned} \tag{18}$$

then $\lambda^* \in (0,1)$.

Sketch of proof: The result is proven in two phases. First, we show that $\lambda^* < 1$: We compute the first derivative of $e_{\mathcal{R}}(\lambda)$ at $\lambda = 1$ and show that it is positive; hence, $e_{\mathcal{R}}(\lambda)$ is strictly increasing in a left neighborhood of $\lambda = 1$. Second, we show that $\lambda^* > 0$: We compute the right derivative of $e_{\mathcal{R}}(\lambda)$ as $\lambda \to 0^+$ and show that it is negative; hence, the error function is strictly decreasing in a right neighborhood of $\lambda = 0$. The detailed calculations are reported in Appendix D. □

Intuitively, any optimal parameter λ^* is strictly between 0 and 1 if the misbehavior is sufficiently disruptive so that the consensus protocol yields poor performance, similar to what remarked in Proposition 1 while full competition is never optimal under our standing assumptions.

Remark 5 (Optimal competition with general matrices): Even though we assume regular agents have no self-loops, Theorem 1 holds also if this is relaxed. Furthermore, we numerically show that $\lambda^* \in (0,1)$ if W^o is row stochastic and Σ is not diagonal.

Remark 6 (Optimal competition with zero noise): Theorem 1 implies that λ^* may be positive even if V and Q are zero. This is, indeed, consistent with the misbehavior model: Not only misbehaving agents corrupt the consensus value through

deception bias and deception noise but mostly they behave against the prescribed protocol, so that full collaboration is in general a poor strategy even if v and $n(k)$ are trivial.

D. Performance Versus Misbehavior

We now study how the performance of the dynamics (FJ) varies with deception biases v and deception noises $n(k)$.

We first show an intuitive result: More disruptive misbehavior induces larger consensus errors for every λ.

Proposition 2 (Performance versus misbehavior): The error $e_{\mathcal{R}}$ is strictly increasing with V and with Q w.r.t. the partial order of semidefinite matrices.

Proof: See Appendix F. □

We next study what happens to the optimal competition λ^*. Intuitively, the more the nominal system behavior is disrupted, the more regular agents should benefit from competing rather than collaborating with (potential) misbehaving neighbors. Formally speaking, this requires λ^* to increase with the intensities of deception biases and noises. Such a claim is hard to prove analytically because of the involved structure of the cost function. In particular, studying the second derivative of $e_{\mathcal{R}}(\lambda)$ is complicated by the asymmetric matrix inside the trace of $e_{\mathcal{R},v}(\lambda)$, and similarly, a unique root of the first derivative of $e_{\mathcal{R}}(\lambda)$ cannot be proved, in general.

Nonetheless, the next results contribute toward our intuition by describing how the minimum points vary with the misbehavior. For convenience, we denote the diagonal elements of the covariance matrices by $d_m \doteq V_{mm}$ and $q_m \doteq Q_{mm}$.

Proposition 3 (Optimal competition versus misbehavior): Let $\lambda_{\min}$ be a minimum point of $e_{\mathcal{R}}(\lambda)$, then $\lambda_{\min}$ is strictly increasing with $d_m, m \in \mathcal{M}$, and with Q w.r.t. the partial order of semidefinite matrices.

Proof: See Appendix G. □

An immediate consequence of Proposition 3 is that, if there is a unique minimum point for some values of V and Q, then there is a unique minimum point for any "larger" V and Q, which corresponds to λ^*. In other words, a more disruptive misbehavior forces regular agents to progressively become more competitive, in order not to be deceived by misbehaving agents that can draw them away from the nominal average consensus. The next proposition refines this result by describing the optimal parameter λ^* with "extreme" misbehavior.

Proposition 4 (Optimal competition with extreme misbehavior): Let $\lambda_{\min}$ be a minimum point of $e_{\mathcal{R}}(\lambda)$, then $\lim_{d_m\to\infty} \lambda_{\min}(d_m) = 1$ and $\lim_{q_m\to\infty} \lambda_{\min}(q_m) = 1, m \in \mathcal{M}$.

Proof: See Appendix H. □

According to intuition, the (trivial) optimal strategy for regular agents is to fully compete when the misbehavior is too disruptive. However, numerical tests in the next section show that λ^* is significantly smaller than 1 in several cases.

IV. Numerical Experiments

In this section, we perform numerical experiments on the consensus error $e_{\mathcal{R}}$ to achieve intuition about the behavior of FJ dynamics under different topologies and misbehavior and draw insight about effective choices of the parameter λ.

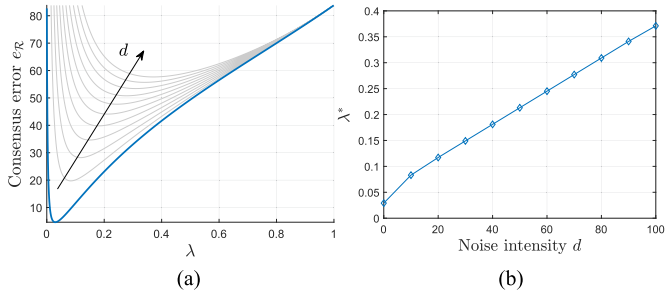

Fig. 2. FJ dynamics with 3-regular graph, exponential decay of observation covariances, and one misbehaving agent. The arrow shows how the error curve varies as the intensity d of the deception bias increases. (a) Average consensus error. (b) Optimal λ as a function of d.

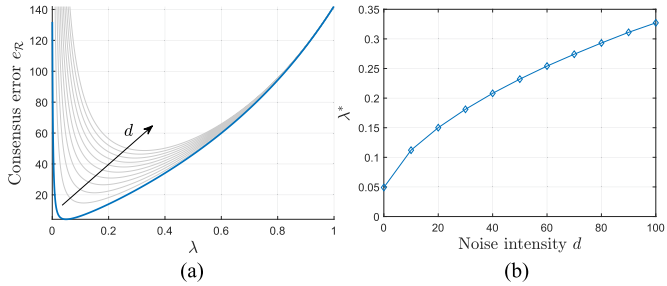

Fig. 3. FJ dynamics CE with 3-regular graph, diagonal prior covariance matrix Σ, and one misbehaving agent. (a) Average CE. (b) Optimal λ as a function of d.

In Fig. 2, we considered a 3-regular communication graph with 100 agents and uniform weights.[1] The prior covariance Σ was chosen such that, for each agent i, the cross-covariances obeyed an exponential decay, $\sigma_{ij} = 10^{-0.2\ell(i,j)}$, $\ell(i,j)$ being the length of the shortest path between i and j, with $\sigma_i^2 \equiv 1$. Furthermore, we randomly selected one misbehaving agent and varied the intensity of its deception bias d within the range $[0, 100]$, with constant intensity of deception noise q.

Fig. 2(a) shows the error curve as d increases. All curves exhibit a unique minimum point λ^*, plotted in Fig. 2(b). Furthermore, both error curve and minimum point increase with d, according to Propositions 2 and 3, showing that the competition level needs to grow with the intensity of deception biases. The same qualitative behavior was observed by varying q.

Fig. 3 shows the same experiment but with a diagonal covariance matrix Σ. We observe the same monotonic behavior of $e_{\mathcal{R}}$ and λ^*. Furthermore, we note that the error curve has a convex shape. In fact, even though it was not possible to prove it formally, all tests performed with diagonal covariance matrices resulted in strictly convex error functions from numerical tests. We next studied what happens when increasing the number of misbehaving agents M. To better visualize changes in the behavior of the system, we fixed the set $\mathcal{R}$ to be a network composed of $R = 100$ regular agents and added misbehaving agents across the network. Fig. 4 shows the error curve when 10 such agents are progressively introduced. In particular, in

[1] In a k-regular graph, each node has exactly k neighbors.

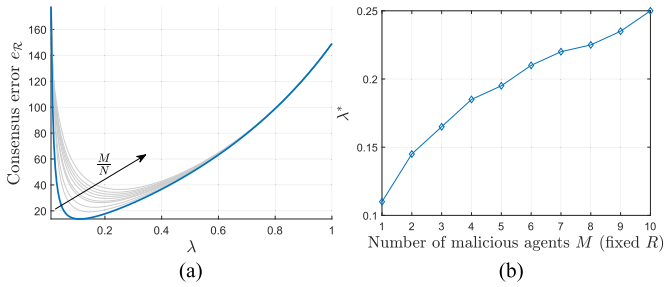

Fig. 4. FJ dynamics with 3-regular graph and diagonal prior covariance matrix Σ. The arrow on the left box shows how the error varies as the number of misbehaving nodes M increases (with $R = 100$). (a) Average consensus error. (b) Optimal λ as a function of M.

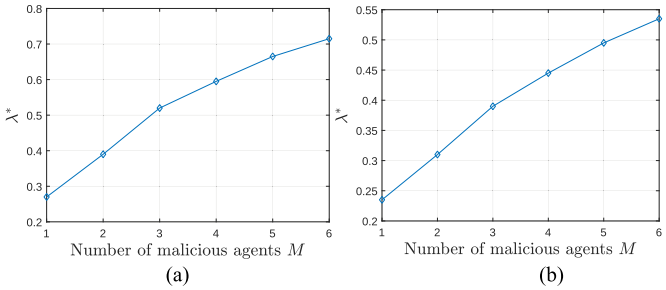

Fig. 5. Optimal λ as a function of M with $d = 10$. Each pair of misbehaving agents affects the same regular agent (e.g., the first two belong to $\mathcal{N}_1$). (a) Exponential observation covariances. (b) Diagonal observation covariance matrix.

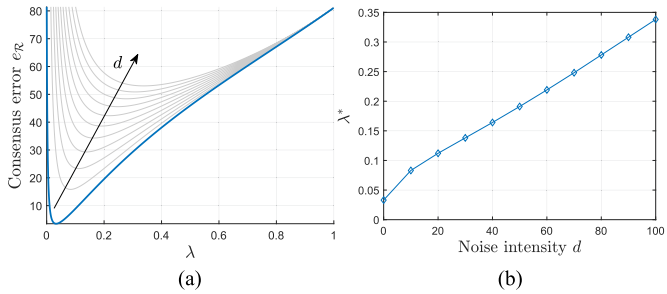

Fig. 6. FJ dynamics with $(3, 4)$-degree communication graph, exponential decay of observation covariances, and one misbehaving agent. (a) Average consensus error. (b) Optimal λ as a function of d.

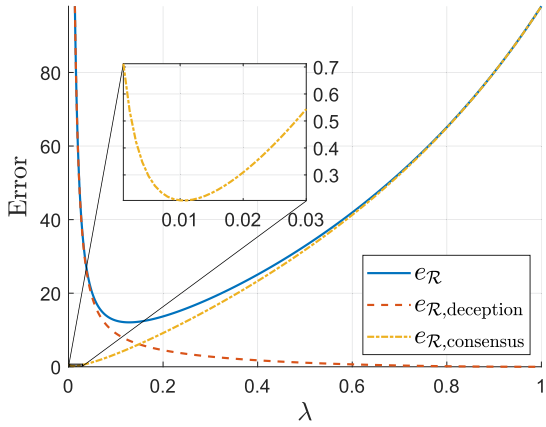

Fig. 7. Consensus error and its two contributions in (19).

this example, all misbehaving agents are selected so as to affect different portions of the network, which allows λ^* to have relatively low values, see Fig. 4(b). Conversely, we note that, in the opposite scenario, some regular agents may be forced to almost freeze their observations (large λ) to not drive the error too large. Fig. 5 shows two cases where the misbehaving agents are connected to the same regular agents. In particular, each couple is added to the neighborhood of one regular agent (e.g., the first two misbehaving agents added to the network are neighbors of agent $1 \in \mathcal{R}$). In this case, λ^* increases faster than Fig. 4(b) because the regular agents affected by multiple misbehaving need to keep their errors small: In other words, they can hardly collaborate because of their misbehaving neighbors. We note that λ^* grows faster when observations of regular agents are correlated [see Fig. 5(a)] because such agents can trust that their states may be similar even before starting dynamical updates, and competing is less risky than collaborating.

Finally, it is interesting to see that the error behavior observed above is consistent also if W^o is only row-stochastic, thus yielding nonzero consensus error even in the nominal scenario. Fig. 6 shows consensus error and λ^* when each node in the graph has degree 3 or 4 and W^o has uniform weights.

Other numerical tests performed with different graphs, observation distributions, and choice of the misbehaving agents show the same quasiconvex behavior of the error function and are omitted in the interest of space. This reinforces and extends the scope of our formal analysis, showing that, indeed, the competition–collaboration tradeoff emerges as a natural resilient mechanism for multiagent systems.

Remark 7 (Value of optimal λ): A remarkable feature of the FJ dynamics that emerges from the tests above is that λ^* is usually small (within the interval $[0.1, 0.2]$ in many cases). This translates into the practical advantage that adding a little competition may be sufficient to get a good level of resilience without forcing too conservative updates by regular agents.

A. Competition–Collaboration Tradeoff: Analytical Insight

As mentioned earlier, the consensus error function $e_{\mathcal{R}}(\lambda)$ is hard to study and an exhaustive analysis seems not possible.

Some intuition can be achieved from a decomposition that we study next. To keep the notation light, we assume a single misbehaving agent (with label m) and a diagonal covariance matrix Σ. Then, we can expand the consensus error as follows:

$$e_{\mathcal{R}} = \underbrace{\sum_{i\in\mathcal{R}} \sigma_i^2 \left\| L_i^{-m} - \frac{\mathbb{1}}{R} \right\|^2}_{\doteq e_{\mathcal{R},\text{consensus}}} + \underbrace{\left(\sigma_m^2 + d\right) \left\| L_m^{-m} \right\|^2 + e_{\mathcal{R},n}}_{\doteq e_{\mathcal{R},\text{deception}}}. \tag{19}$$

In (19), $L_i \in \mathbb{R}^N$ is the ith column of L and $L_i^{-m} \in \mathbb{R}^{N-1}$ is obtained from L_i by removing its mth row (corresponding to the misbehaving agent). The error curves are shown in Fig. 7. Equation (19) allows for an intuitive interpretation of the error, which leverages the notion of *social power* [48], [49].

In opinion dynamics, the social power is used to quantify how much the opinion of an agent affects the opinions of all agents. In particular, when opinions evolve according to the FJ dynamics, the element L_{ij} quantifies the influence of agent j on agent i: As L_{ij} increases, agent i is more affected by the initial opinion of agent j. The total social power of agent j is a symmetric and increasing function of all elements $\{L_{ij}\}_{i\in\mathcal{V}}$.[2]

Borrowing such concepts from opinion dynamics allows us to interpret the two contributions separated in (19). The first, $e_{\mathcal{R},\text{deception}}$, quantifies the impact of the misbehaving agent m on regular agents. The "social power" of m, as quantified through the vector L_m^{-m}, depends on the communication matrix W and on the parameter λ. Each coordinate of L_m^{-m} decreases with λ, meaning that the influence of the misbehaving agent weakens as regular agents anchor more tightly to their observations, and becomes zero when $\lambda = 1$, namely, in the full-competition regime. We formalize this discussion as the following lemma.

Lemma 2: The component $e_{\mathcal{R},\text{deception}}$ is decreasing with λ.

Proof: By computing the derivative of L w.r.t. λ, we see that each element of L_m^{-m} is nonincreasing with λ. Because L is a nonnegative matrix, this and Lemma 1 yield the claim. See Appendix E for the detailed calculations. □

The second contribution $e_{\mathcal{R},\text{consensus}}$ measures "democracy" among regular agents, i.e., it is proportional to the mismatch between how much each regular agent affects the others and the ideal value $1/R$, which means that each agent affects all others equally. This cost is zero if and only if the submatrix of L corresponding to interactions among regular agents is the consensus matrix: This can happen only if they do not interact with the misbehaving agent [47], in which case the vector L_m^{-m} is zero (the misbehavior has no effect). In this special case, $e_{\mathcal{R},\text{consensus}}$ is zero at $\lambda = 0$ and increases monotonically as the network shifts from a democratic system where agents fully collaborate ($\lambda = 0$) to a disconnected system where agents fully compete ($\lambda = 1$). Conversely, with misbehaving agents, $e_{\mathcal{R},\text{consensus}}$ has a nontrivial minimizer (zoomed box in Fig. 7). For small λ, the misbehaving agent overrules all interactions and regular agents hardly affect each other. As λ increases, the interactions among regular agents become more relevant, making $e_{\mathcal{R},\text{consensus}}$ decrease. However, as λ grows further, the competition among regular agents becomes too aggressive and makes them shift away from an ideal democratic system.

Overall, the error (CE) has two concurrent causes that yield two regimes: collaboration with misbehaving agents is most misleading for small λ while for large λ, the error is mainly due to regular agents that compete against each other and reject useful information shared by neighbors. This matches intuition from (FJ) where λ measures conservatism in agent updates.

V. Role of the Communication Network

In the previous sections, we discussed the benefits of using a competition-based approach (FJ dynamics) to tame misbehaving agents. We now shift attention to the communication network, in order to achieve intuition about resilient topologies. In Section V-A, we introduce a second performance metric, which we use to evaluate resilience to attacks. In Section V-B, we observe how performance varies with connectivity.

A. Performance Metrics

Besides consensus error, we also aim to assess energy spent to misbehave. To this aim, we interpret (7) as a controlled system where the misbehaving agents command the input $x_{\mathcal{M}}(\cdot)$. The controllability Gramian in K steps, denoted by $\mathcal{W}_K$, is defined for system (7) as

$$\mathcal{W}_K = \sum_{k=0}^{K-1} A^k B B^\top (A^\top)^k. \tag{20}$$

The controllability Gramian can be used to quantify the control effort: The trace of $\mathcal{W}_K$, called the *controllability index*, is inversely related to the control energy spent in K steps (averaged over the reachable subspace), as shown in the literature [50], [51], [52]. In other words, a small controllability index means that the misbehaviors consume a lot of energy to steer $x_{\mathcal{R}}$ across the reachable space, which may be desired to possibly drain out adversarial resources and hamper an external attack.

If $M = 1$, the controllability index can be written as

$$\text{Tr}\left(\mathcal{W}_K\right) = (1-\lambda)^2 \sum_{k=0}^{K-1} \left\| (1-\lambda)^k W_{\mathcal{R}}^k W_{\mathcal{M}} \right\|^2 \tag{21}$$

resembling the consensus error component $e_{\mathcal{R},\text{deception}}$ in (19)

$$e_{\mathcal{R},\text{deception}} \propto \left\| \sum_{k=0}^{\infty} (1-\lambda)^k \sum_{j=0}^{k-1} W_{\mathcal{R}}^j W_{\mathcal{M}} \right\|^2. \tag{22}$$

Both $\text{Tr}(\mathcal{W}_K)$ and $e_{\mathcal{R},\text{deception}}$ are decreasing with λ (i.e., the more competition, the better) and depend on the vectors $W_{\mathcal{R}}^k W_{\mathcal{M}}$ that describe how attacks spread in k steps. The discount factor $(1-\lambda)^k$ makes the tail of the series in (22) negligible, enhancing the similarity between those two metrics.

Remark 8 (Controllability index): While we use Assumption 2 to compute $e_{\mathcal{R}}$, the controllability Gramian in (20) is independent of the trajectory of the system, and hence, the controllability index evaluates an "average trajectory" of misbehaving agents.

B. Network Connectivity Versus Resilience

We now explore how connectivity of the communication network affects performance and resilience of the dynamics (FJ). While in this section, we attempt to achieve heuristic intuition, an analytical investigation is deferred to future work. To this aim, we fix the parameter $\lambda = 0.1$ and numerically evaluate the theoretical performance as the density of the communication network increases. Specifically, for each evaluated network, we assign uniform weights to the links and compute consensus error $e_{\mathcal{R}}$ and controllability index $\text{Tr}(\mathcal{W}_K)$ (where K is the reachability index) selecting some agents as misbehaving according to either of the following two cases.

1) The worst case misbehaving agent, i.e., $\mathcal{M} = \{m^*\}$ with

$$m^* = \arg\max_{m\in\mathcal{V}} \; e_{\mathcal{R}} \tag{23}$$

$$m^* = \arg\max_{m\in\mathcal{V}} \; \text{Tr}\left(\mathcal{W}_K\right). \tag{24}$$

[2]The authors in [48] and [49] use the arithmetic mean of $\{L_{ij}\}_{i\in\mathcal{V}}$.

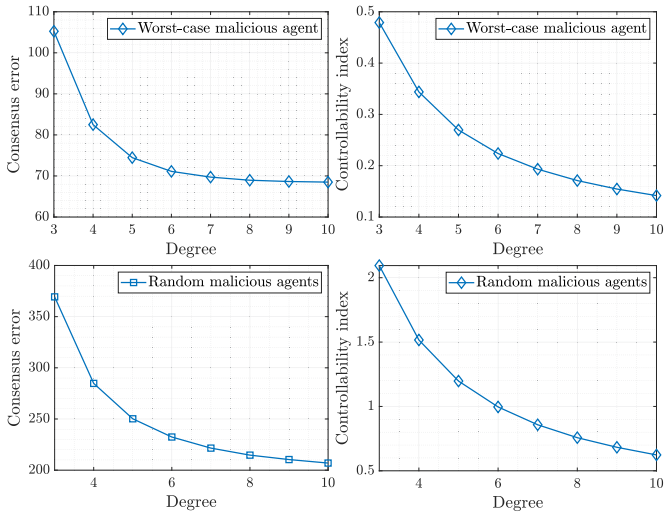

Fig. 8. Average performance metrics for regular graphs.

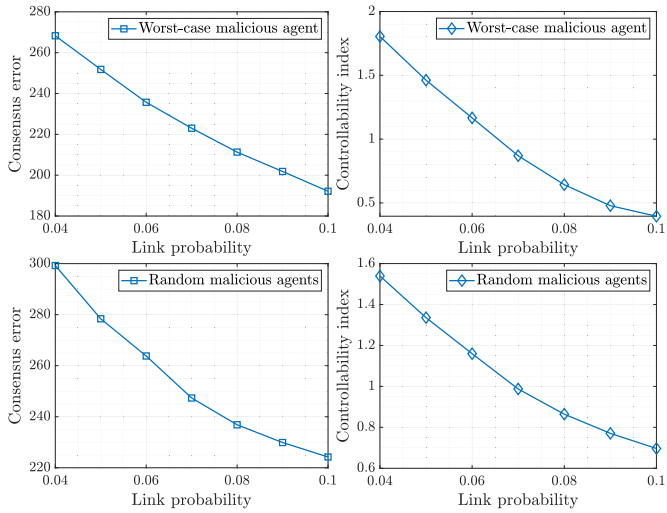

Fig. 9. Average performance metrics for Erdös–Rényi random graphs.

2) Five misbehaving agents randomly drawn from $\mathcal{V}$.

We consider three common classes of graphs: 1) regular graphs with degree Δ, 2) Erdös–Rényi random graphs, where a link between any two nodes exists with probability p, and 3) random geometric graphs, where nodes are randomly placed in $[0, 1]^2$ and any two nodes are linked if their distance is not greater than a radius ρ. While regular graphs induce a doubly stochastic matrix even with simple uniform weights, this is generally not true for the other graphs. Hence, to evaluate the consensus error $e_{\mathcal{R}}$, we considered both the deviation from the nominal average defined in (CE) and the deviation from the consensus value computed from the left Perron eigenvector of the nominal weight matrix W^o. Given that the results were qualitatively equal, we report only the first case in the interest of space.

We consider networks with $N = 100$ agents and compute the performance for each network (i.e., a combination of class of graph and density parameter) by averaging over 1000 random graphs for the worst case misbehaving agent and over 5000 random graphs for the random selection of misbehaving agents. The results are shown in Figs. 8–10, with the consensus error on the left and the controllability index on the right.

The main insight is that, on average, increasing the graph connectivity mitigates attacks with respect to both metrics. Intuitively, this is because high degrees mean many interactions among regular agents that the misbehaving agent cannot control directly. The only remarkable difference is noted in random geometric graphs with the worst case misbehaving agent (top-left box in Fig. 10), for which increasing the radius from 0.35 to 0.5 also increases the consensus error. This might be due to the formation of hubs, that is, densely connected areas that emerge and become denser as the radius increases, which an adversary can exploit to quickly spread damage to a large portion of the network. Notably, this phenomenon is absent both for the same class of graphs with a random selection of misbehaving agents (bottom-left box of Fig. 10) and in the case of Erdös–Rényi random graphs (see Fig. 9), which also typically feature some dense areas—even though not with the small world structure typical of random geometric graphs, see Figs. 14(a) and 15(a).

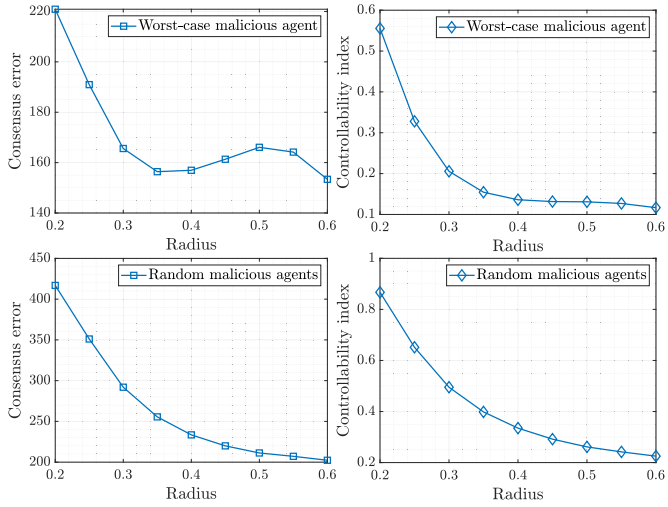

Fig. 10. Average performance metrics for random geometric graphs.

A deeper study of this phenomenon is an interesting direction for future research.

Besides density and number of links, an aspect that also seems to play a role in resiliency is degree balance among nodes. This can be somehow deduced by the plots referred to the same selection strategy of misbehaving agent: For example, with worst case misbehaving agents, regular graphs exhibit the smallest costs, random geometric graphs—where usually nodes have similar number of neighbors—yield worse performance, and Erdös–Rényi random graphs—where both highly connected and almost isolated nodes coexist—have the largest costs.

To more carefully investigate how performance varies with degree balance, we consider *almost-regular* graphs, namely, where nodes have degrees either Δ or $\Delta - 1$ for some Δ. This corresponds to "middle-ways" between Δ- and $(\Delta - 1)$-regular graphs, which could be ideally placed between two consecutive ticks (degrees) Δ and $\Delta - 1$ on the x-axis of Fig. 8.

More specifically, starting from a Δ-regular graph, we iteratively remove one edge at a time so as to minimize performance degradation while selecting the worst case misbehaving agent at

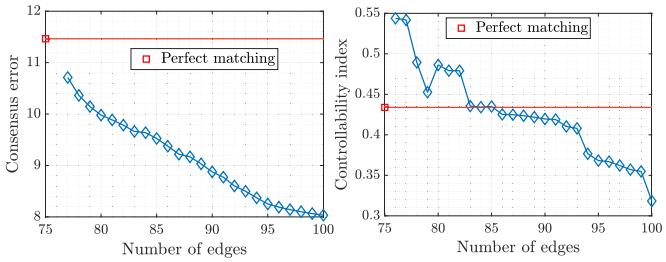

Fig. 11. Consensus error (left) and controllability index (right) for almost-regular graphs starting from a 4-regular graph with $\lambda = 0.2$. Edge removal proceeds from the right (initially, all 100 edges are present) toward the left. At each iteration, one edge is removed so as to minimize performance degradation according to (25)–(26) while enforcing that each node has a degree of either three or four. At the last iteration (leftmost diamonds), most or all nodes have degree three, with possibly a few nodes left with degree four. The red squares show the performance metrics for a 3-regular graph obtained by removing a perfect matching (set of edges) from the initial 4-regular graph.

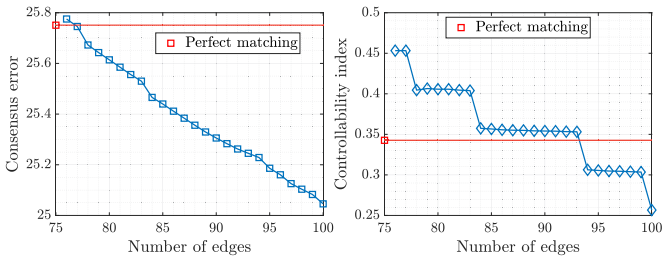

Fig. 12. Consensus error (left) and controllability index (right) for almost-regular graphs starting from a 4-regular graph with $\lambda = 0.7$.

each time. This amounts to removing the edge e that solves

$$\min_{e\in\mathcal{E}} \max_{m\in\mathcal{V}} \quad e_{\mathcal{R}}(\mathcal{E}\setminus\{e\}) \tag{25}$$

$$\min_{e\in\mathcal{E}} \max_{m\in\mathcal{V}} \quad \mathrm{Tr}\left(\mathcal{W}_K(\mathcal{E}\setminus\{e\})\right) \tag{26}$$

where $\mathcal{E}$ is the set of edges (nonzero elements of W) and we set W with uniform weights after each removal. To get almost-regular graphs, we remove at most one edge per node.

Figs. 11 and 12 show the performance obtained starting from a 4-regular graph with 50 nodes (100 edges in total, corresponding to the rightmost point in the plots) and gradually pruning edges according to (25)–(26) (proceeding leftward on the x-axis). Also, the performance with a 3-regular graph obtained by removing perfect matchings from the initial 4-regular graphs is shown for comparison.[3] Remarkably, performance degrades (almost) monotonically for both performance metrics as edges are removed. This may be explained by a combination of lower connectivity and degree unbalance, which allows the adversary to exploit highly connected agents to make more effective damage against low-connected regular agents.

[3]A matching is a set of edges that do not share nodes. A maximum matching is a matching of maximal cardinality, and a perfect matching is a maximum matching such that each node is incident to one edge (*total coverage*). Note that our edge removal strategy need not remove exactly one edge for each node in the graph because we constrain the resulting graphs to be almost regular. For example, the iterative removal corresponding to Fig. 12 stops before reaching a 3-regular graph.

Interestingly, while the consensus error increases smoothly as edges are removed, the controllability index exhibits "jumps." This is evident with large λ, as Fig. 12 shows. Such a behavior suggests the presence of critical subsets of edges and might give indication about critical links to be kept or removed.

Furthermore, in almost all tests (not shown here in the interest of space), the 3-regular graph obtained by removing a perfect matching yielded better performance compared to the last edge removal (leftmost marker on the blue curve). This suggests that increasing connectivity may not be beneficial if it entails less degree balance: In Fig. 12, the 3-regular graph reduces both the consensus error and the controllability index w.r.t. the last graphs obtained by pruning edges (leftmost markers), which have one node with degree 4 and all others with degree 3. In particular, the latter metric is reduced by 22% and is comparable to graphs having most nodes with degree 4. However, as shown in Fig. 11, a regular graph of degree $\Delta - 1$ obtained by removing a perfect matching (not related to performance metrics) from a Δ-regular graph may yield worse performance than almost-regular graphs. This gives further insights: An arbitrary edge selection may perform substantially worse compared to a task-related strategy.

VI. Comparison With Existing Literature

In this section, we test our proposed protocol and compare its performance with other approaches in the literature.

Many techniques have been proposed to mitigate misbehaving agents. However, they usually focus on reaching a generic consensus, possibly while keeping the states of regular agents within a safe region (usually defined by initial conditions), and do not consider the *performance of average consensus*, which here is key to the distributed optimization task, as argued in Section II. Indeed, most resilient consensus strategies aim to make the regular agents agree on, e.g., a common location (such as in robot gathering) in the face of misleading interactions, but need not relate the consensus value to the initial locations.

We compare two strategies: Weighted MSR (W-MSR) [26] and Secure Accepting and Broadcasting Algorithm (SABA) [35]. As noted in Section I-A, many resilient algorithms adapt W-MSR to specific applications and enjoy the same guarantees. W-MSR suffers from two main limitations related to r-robustness, which is the cornerstone of all theoretical analysis. First, while sufficient conditions for resilient consensus are clear, there is little clue about necessary conditions. This translates into an unknown behavior of the system if r-robustness does not hold. While r-robustness has proved a good characterization for update rules based on W-MSR, it raises practical limitations. On the one hand, the communication network may be fixed but not robust enough. On the other hand, checking r-robustness is computationally intractable for large-scale networks [32]. Thus, in some cases, for example, with a sparse structure, a more conservative behavior with provable performance bounds may be preferred. Also, W-MSR requires to estimate the number of misbehaving agents affecting the network. This may be an issue: If the estimate is too low, regular agents may be deceived and average consensus disrupted, whereas, if it is too high, the updates may be too conservative, possibly preventing

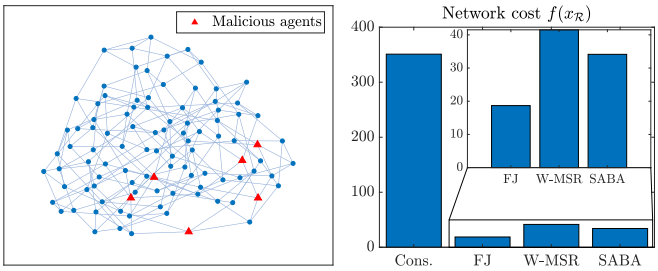

Fig. 13. Comparison among consensus, FJ, W-MSR [26], and SABA [35] with 4-regular graph and six misbehaving agents. (a) Communication network. (b) Cost (3) of regular agents.

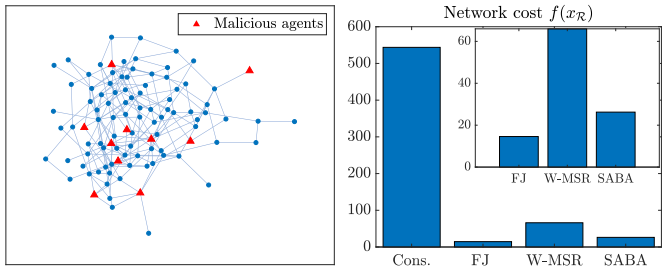

Fig. 14. Comparison among consensus, FJ, W-MSR [26], and SABA [35] with Erdös–Rényi random graph with $p = 4/N$ and 10 misbehaving agents. (a) Communication network. (b) Cost (3) of regular agents.

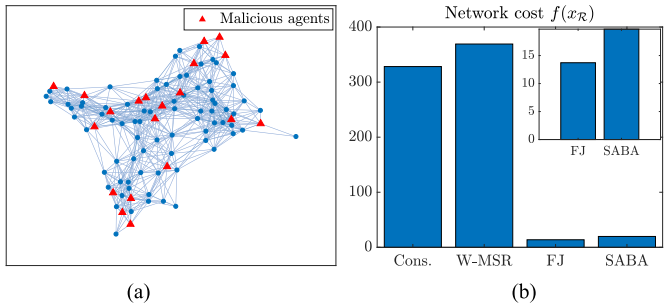

Fig. 15. Comparison among consensus, FJ, W-MSR [26], and SABA [35] with a random geometric graph with $\rho = 0.25$ and 20 misbehaving agents. Communication network. (b) Cost (3) of regular agents.

convergence. Furthermore, misbehaviors could happen in a time-varying fashion and make the r-robustness fail at times, yielding poor performance overall. SABA does not estimate the number of misbehaving agents but stores all received values in a buffer and processes them with a voting strategy. However, this design may impose impractical memory requirements, and the convergence of SABA is still ensured under a minimal r-robustness.

In the next simulations, we consider $N = 100$ agents interacting through sparse communication networks, whose low connectivity hampers W-MSR and SABA, and matrices W^o with homogeneous weights. As the performance metric, we computed the objective cost of the distributed optimization task (3), which equals $e_{\mathcal{R}}$ up to additive constants, cf. Section II-B. The observations are drawn as $\theta \sim \mathcal{N}(0, 0.1I)$ and each misbehaving agent m is assigned a deception bias $v_m \in [2, 6]$. For each scenario, we chose the parameter λ by selecting the minimizer of the theoretical error $e_{\mathcal{R}}(\lambda)$ with $V = 5I_M$.

Fig. 13 illustrates a network where the agents use a regular graph with degree $\Delta = 4$ as a communication network with six misbehaving agents [see Fig. 13(a)]. We implement W-MSR assuming that each regular agent has (at most) one misbehaving neighbor because larger values make updates trivial, i.e., $x_i(k) \equiv x_i(0)$. However, some misbehaving agents communicate with the same regular agents (e.g., the two in the bottom-right portion of the graph), making this scenario challenging for W-MSR and SABA whose r-robustness requirement suffers the sparse communication graph. While both SABA and W-MSR perform poorly [see Fig. 13(b)], our approach mitigates the attacks by setting λ at a suitably large value.

In Fig. 14, we simulate the protocols over an Erdös–Rényi random graph with link probability $p = 4/N$. (hence, each agent has $(N-1)p$ neighbors on average), and ten misbehaving agents (10% of the total number of agents). Note that the matrix W is row stochastic. Also, in this case, the dynamics (FJ) tames the numerous attacks better than the compared approaches.

Finally, we address a random geometric graph with radius $\rho = 0.25$ and 20 misbehaving agents in Fig. 15. Also, in this case, the matrix W is row stochastic. Interestingly, W-MSR is rather challenged by this class of graphs, yielding a large cost. On the other hand, the dynamics (FJ) again manages to keep the error small compared to the other algorithms.

Many other simulations that validated our approach are not reported here in the interest of space and can be found in [53].

Remark 9 (Advantages of FJ dynamics): The previous experiments highlight some advantages of the proposed approach. First, the tunable parameter λ makes the algorithm flexible because it can smoothly adapt to a different attack intensity while still providing decent performance. Furthermore, while the optimal parameterization requires exact knowledge of the adversary, which may not be reasonably assumed, yet our proposed approach proves pretty robust to the choice of a specific λ, as shown in Figs. 2–6 where the error is kept small around λ^*. This also holds with row-stochastic matrices, enabling simple weighing rules to be locally implemented. In contrast, in other approaches, the cost function may be highly sensitive to some design parameters, e.g., the estimated number of misbehaving agents in W-MSR. Furthermore, most works in the literature do not describe the system behavior when resilient consensus is not guaranteed. In fact, they usually either ensure that the states of the agents remain inside the convex hull of the initial conditions [which may be equivalent to setting $\lambda = 1$ in (FJ)], or let agents reach consensus but potentially be steered far away from initial conditions [38]. Finally, computational complexity and memory requirements are minimal, which is typically desired for resource-constrained devices.

VII. Conclusion and Future Work

In this article, we have proposed a competition-based protocol based on the FJ dynamics to mitigate a class of misbehavior that disrupts a quadratic distributed optimization task. We have presented formal results and numerical experiments on performance and optimal parameterization and showed that our approach can

outperform state-of-the-art algorithms. Furthermore, we have discussed the competition–collaboration tradeoff with analytical arguments that are insightful toward a deeper understanding of the fundamental properties of the system in the presence of misbehaviors. Finally, we have addressed network design and explored how resilience relates to graph connectivity, looking both at the optimization performance and at the energy spent to misbehave.

This approach opens several avenues for future research. First, it is desirable to address an effective design of parameters λ_i's in the realistic case where knowledge about the attack is scarce. This may also involve online reweighing of protocol parameters, for example in the realm of recent work where weights are updated via trust information or evidence theory [38], [40].

Second, the more general and challenging scenario of distributed optimization should be addressed. In this case, a common approach is to alternate local descent steps to consensus updates to steer all agents toward a common point [19], [21]. Here, the task-tailored descent steps may critically impact performance even if consensus steps are made resilient.

A third research avenue involves zero-sum games to model interactions among agents [43], [44]. In particular, in asymmetric zero-sum games, one player has more knowledge than the other, which is a suitable model for worst case attacks. In this case, a relevant challenge is determining the optimal strategies for both players, to ultimately derive effective resilient algorithms in the presence of intelligent adversaries.

Finally, it is interesting to deeply investigate the design of the communication network. While graph robustness to node or edge failures has been extensively addressed [54], [55], [56], [57], the novel element given by the dynamics (FJ) calls for a tailored investigation as heuristically motivated in Section V. Also, in the spirit of a graph-theoretic approach, a comparison between classical centrality measures and worst case attacks may be useful to get insight about agents that deserve higher attention.

APPENDIX

A. Useful Lemmas

In this Appendix, we report some standard facts in linear algebra that will be used in the following proofs.

Lemma A.1: Let $\alpha \in \mathbb{R}$ and $A, B \in \mathbb{R}^{n\times n}$ differentiable functions of α, then the derivative of $\mathrm{Tr}(A^\top B)$ is

$$\frac{d\mathrm{Tr}\left(A^\top B\right)}{d\alpha} = \mathrm{Tr}\left(\frac{dA^\top}{d\alpha}B\right) + \mathrm{Tr}\left(A^\top \frac{dB}{d\alpha}\right). \tag{A.1}$$

Lemma A.2: Let $\alpha \in \mathbb{R}$ and $A \in \mathbb{R}^{n\times n}$ invertible and differentiable function of α, then the derivative of A^{-1} is

$$\frac{dA^{-1}}{d\alpha} = -A^{-1}\frac{dA}{d\alpha}A^{-1}. \tag{A.2}$$

Lemma A.3: Let $A \in \mathbb{R}^{n\times n}$ invertible with eigenpair (λ, v), then A^{-1} has eigenpair (λ^{-1}, v).

Corollary A.1: If $A \in \mathbb{R}^{n\times n}$ is diagonalizable, then A and A^{-1} are simultaneously diagonalizable.

Lemma A.4: Let $A \in \mathbb{R}^{n\times n}$ have eigenpair (λ, v), then $(I - \alpha A)$ has eigenpair $((1-\alpha\lambda), v)$.

B. Proof of Lemma 1

We use the implicit function theorem to prove that each diagonal element P_{ii} of P is strictly decreasing with λ. Let

$$g_i(\lambda, P_{ii}) \doteq P_{ii} - (1-\lambda)^2\left[W_{\mathcal{R}}PW_{\mathcal{R}}^\top\right]_{ii} - (1-\lambda)^2\tilde{Q}_{ii} \tag{B.1}$$

where $\tilde{Q} = W_{\mathcal{M}}QW_{\mathcal{M}}^\top$. The implicit function theorem holds for the solutions of $g_i(\lambda, P_{ii}) = 0$ with $\lambda \in (0,1)$: for $i \in \mathcal{R}$

$$\frac{\partial g_i(\lambda, P_{ii})}{\partial P_{ii}} = 1 - (1-\lambda)^2 W_{ii}^2 \neq 0 \quad \forall \lambda \in (0,1). \tag{B.2}$$

Making dependence on λ explicit and for $\lambda \in (0,1)$, we get

$$\begin{aligned}\frac{\mathrm{d}P_{ii}(\lambda)}{\mathrm{d}\lambda} &= -\frac{\partial g_i(\lambda, P_{ii}(\lambda))}{\partial \lambda}\left(\frac{\partial g_i(\lambda, P_{ii}(\lambda))}{\partial P_{ii}(\lambda)}\right)^{-1} \\ &= -\frac{2(1-\lambda)\left(\left[W_{\mathcal{R}}PW_{\mathcal{R}}^\top\right]_{ii} + \tilde{Q}_{ii}\right)}{1-(1-\lambda)^2W_{ii}^2} < 0.\end{aligned} \tag{B.3}$$

Finally, from $e_{\mathcal{R},n}(\lambda) = \sum_{i\in\mathcal{R}} P_{ii}(\lambda)$ and linearity of the derivative, it follows that $e_{\mathcal{R},n}(\lambda)$ is decreasing.

For $\lambda = 1$, we trivially get $P = 0$, and thus, $e_{\mathcal{R},n}(1) = 0$.

C. Proof of Proposition 1

In this and all following proofs, the constant κ in (13) is neglected for the sake of simplicity.

We first compute the consensus error with $\lambda = 1$

$$\begin{aligned}e_{\mathcal{R}}(1) = e_{\mathcal{R},v}(1) &= \mathrm{Tr}\left(\Sigma_{11}(I_R - C_R)\right) \\ &= \frac{R-1}{R}\sum_{i\in\mathcal{R}}\sigma_i^2 - \frac{1}{R}\sum_{i\in\mathcal{R}}\sum_{\substack{j\in\mathcal{R}\\ j\neq i}}\sigma_{ij}.\end{aligned} \tag{C.1}$$

With $\lambda = 0$, the average steady-state consensus value is determined by the biased observations of malicious agents, i.e., $\bar{x}_{\mathcal{R}} = \bar{\bar{\theta}}_{\mathcal{M}} \doteq \frac{1}{M}\sum_{m\in\mathcal{M}}(\theta_m + v_m)$. We have

$$\begin{aligned}e_{\mathcal{R},v}(0) &= \mathbb{E}\left[\left\|\mathbb{1}_R\bar{\bar{\theta}}_{\mathcal{M}} - \mathbb{1}_R\bar{\theta}_{\mathcal{R}}\right\|^2\right] \\ &= \frac{R}{M^2}\sum_{m\in\mathcal{M}} d_m + \frac{R}{M^2}\sum_{m\in\mathcal{M}}\left(\sigma_m^2 + \sum_{\substack{n\in\mathcal{M}\\ n\neq m}}\sigma_{mn}\right) \\ &\quad + \frac{1}{R}\sum_{i\in\mathcal{R}}\left(\sigma_i^2 + \sum_{\substack{j\in\mathcal{R}\\ j\neq i}}\sigma_{ij}\right) - \frac{2}{M}\sum_{i\in\mathcal{R}}\sum_{m\in\mathcal{M}}\sigma_{im}\end{aligned}$$

$$e_{\mathcal{R},n}(0) = \mathrm{Tr}\left(P(0)\right). \tag{C.2}$$

By comparing the expressions in (C.1) and (C.2), it follows that $e_{\mathcal{R}}(0) > e_{\mathcal{R}}(1)$ is equivalent to

$$\frac{M^2}{R}\mathrm{Tr}\left(P(0)\right) + \sum_{m\in\mathcal{M}} d_m > -\sum_{m\in\mathcal{M}}\left(\sigma_m^2 + \sum_{\substack{n\in\mathcal{M}\\ n\neq m}}\sigma_{mn}\right)$$

$$+\frac{M^2}{R}\sum_{i\in\mathcal{R}}\sigma_i^2-\frac{2M^2}{R^2}\sum_{i\in\mathcal{R}}\left(\sigma_i^2+\sum_{\substack{j\in\mathcal{R}\\ j\neq i}}\sigma_{ij}\right)+\frac{2M}{R}\sum_{i\in\mathcal{R}}\sum_{m\in\mathcal{M}}\sigma_{im} \tag{C.3}$$

which leads to condition (15).

D. Proof of Theorem 1

From (13), we get

$$\begin{aligned}\frac{\mathrm{d}e_{\mathcal{R}}(\lambda)}{\mathrm{d}\lambda}&=\frac{\mathrm{d}e_{\mathcal{R},v}(\lambda)}{\mathrm{d}\lambda}+\frac{\mathrm{d}e_{\mathcal{R},n}(\lambda)}{\mathrm{d}\lambda}\\&=\kappa\frac{1}{\lambda}\mathrm{Tr}\left(\widetilde{\Sigma}L^\top\left(I-W^\top L^\top\right)S_{\mathcal{R}}^\top E\right)+\frac{\mathrm{d}e_{\mathcal{R},n}(\lambda)}{\mathrm{d}\lambda}\end{aligned} \tag{D.1}$$

where Lemmas A.1 and A.2 were used and $\kappa>0$.

1) Part One: $\lambda^*<1$**:** From (B.3), $\frac{\mathrm{d}e_{\mathcal{R},n}(\lambda)}{\mathrm{d}\lambda}\big|_{\lambda=1}=0$ and

$$\frac{\mathrm{d}e_{\mathcal{R}}(\lambda)}{\mathrm{d}\lambda}\Big|_{\lambda=1}=\kappa\mathrm{Tr}\left(\widetilde{\Sigma}\left(I-W^\top\right)S_{\mathcal{R}}^\top\left(S_{\mathcal{R}}-C_RS_{\mathcal{R}}\right)\right). \tag{D.2}$$

The argument of the trace in (D.2) has expression

$$\left[\begin{array}{c|c}A&0\\\hline\star&0\end{array}\right],\qquad A\in\mathbb{R}^{R\times R}. \tag{D.3}$$

Condition C1: If Σ is diagonal, the ith diagonal element of A is

$$A_{ii}=\sigma_i^2\left(1-\frac{1}{R}+\frac{1}{R}\sum_{j\in\mathcal{R}}W_{ji}\right)>0. \tag{D.4}$$

Condition C2: If W^o symmetric, the ith diagonal element of A is

$$\begin{aligned}A_{ii}=\sigma_i^2&+\frac{1}{R}\sum_{m\in\mathcal{M}}\sigma_{im}\left(1-\sum_{m'\in\mathcal{M}\setminus\{m\}}W^o_{m'm}\right)-\frac{1}{R}\sigma_i^2\sum_{m\in\mathcal{M}}W^o_{mi}\\&-\sum_{\substack{j\in\mathcal{R}\\ j\neq i}}\sigma_{ij}W_{ij}-\frac{1}{R}\sum_{\substack{j\in\mathcal{R}\\ j\neq i}}\sigma_{ij}\sum_{m\in\mathcal{M}}W^o_{mj}-\sum_{m\in\mathcal{M}}\sigma_{im}W_{im}.\end{aligned} \tag{D.5}$$

It holds

$$\begin{aligned}\frac{1}{R}\sigma_i^2\sum_{m\in\mathcal{M}}W^o_{mi}+\sum_{\substack{j\in\mathcal{R}\\ j\neq i}}\sigma_{ij}W_{ij}+\frac{1}{R}\sum_{\substack{j\in\mathcal{R}\\ j\neq i}}\sigma_{ij}\sum_{m\in\mathcal{M}}W^o_{mj}\\+\sum_{m\in\mathcal{M}}\sigma_{im}W_{im}\leq\max\left\{\frac{1}{R}\sigma_i^2+\sigma_{im^*},\sigma_{ij^*}\right\}\end{aligned} \tag{D.6}$$

where $j^*\doteq\arg\max_{j\in\mathcal{R}\setminus\{i\}}\sigma_{ij}$ and $m^*\doteq\arg\max_{m\in\mathcal{M}}\sigma_{im}$. Inequality (D.6) can be split into the following two cases.

Case $\frac{1}{R}\sigma_i^2+\sigma_{im^*}^2\geq\sigma_{ij^*}^2$*:*

$$\begin{aligned}A_{ii}&\geq\sigma_i^2+\frac{1}{R}\sigma_{im^*}-\frac{1}{R}\sigma_i^2-\sigma_{im^*}\\&=\left(\sigma_i^2-\sigma_{im^*}\right)\left(1-\frac{1}{R}\right)>0.\end{aligned} \tag{D.7}$$

Case $\frac{1}{R}\sigma_i^2+\sigma_{im^*}^2<\sigma_{ij^*}^2$*:*

$$A_{ii}\geq\sigma_i^2-\sigma_{ij^*}+\frac{1}{R}\sum_{m\in\mathcal{M}}\sigma_{im}\left(1-\sum_{m'\in\mathcal{M}\setminus\{m\}}W^o_{m'm}\right)>0. \tag{D.8}$$

The final inequalities in (D.7)–(D.8) follow from $\Sigma\succ 0$ and the Gershgorin circle theorem that implies $\sigma_i^2>\sigma_{ij}\ \forall i,j\in\mathcal{V}$.

It follows that the derivative (D.2) is positive and the CE is increasing in a left neighborhood of 1. By continuity of (D.1), the minimum points satisfy $\lambda^*<1$.

2) Part Two: $\lambda^*>0$**:** From Lemma 1, the error term $e_{\mathcal{R},n}(\lambda)$ has a negative right derivative at $\lambda=0$. By continuity of the derivative of $e_{\mathcal{R},v}(\lambda)$, we can compute the following limit:

$$\begin{aligned}\lim_{\lambda\to0^+}\frac{\mathrm{d}e_{\mathcal{R},v}(\lambda)}{\mathrm{d}\lambda}&=\mathrm{Tr}\left(\widetilde{\Sigma}\lim_{\lambda\to0^+}\frac{\mathrm{d}L}{\mathrm{d}\lambda}^\top S_{\mathcal{R}}^\top\lim_{\lambda\to0^+}E\right)\\&=\mathrm{Tr}\left(\widetilde{\Sigma}\Gamma^\top S_{\mathcal{R}}^\top\left(S_{\mathcal{R}}\overline{W}-C_RS_{\mathcal{R}}\right)\right)\\&=\mathrm{Tr}\left(\widetilde{\Sigma}\Gamma^\top S_{\mathcal{R}}^\top\left[-C_R\,|\,C_{RM}\right]\right)\\&=\mathrm{Tr}\left(-\Sigma_{11}\Gamma_1^\top C_R-\Sigma_{12}\Gamma_2^\top C_R\right.\\&\quad\left.+\Sigma_{12}^\top\Gamma_1^\top C_{RM}+(\Sigma_{22}+V)\Gamma_2^\top C_{RM}\right)\end{aligned} \tag{D.9}$$

where the steady-state consensus matrix $\overline{W}\doteq\lim_{\lambda\to0^+}L$ has block partition (cf. Assumption 2 for the value of $\overline{W}$)

$$\overline{W}=\left[\begin{array}{c|c}0&C_{RM}\\\hline0&I_M\end{array}\right]. \tag{D.10}$$

Matrix Γ can be computed from the spectral decomposition of W. In particular, its elements are finite, Γ_1 is nonnegative, and Γ_2 is nonpositive (details in Appendix E). Putting together (D.9) and Lemma 1, the right derivative of $e_{\mathcal{R}}(\lambda)$ at $\lambda=0$ is negative if and only if the following inequality holds:

$$\begin{aligned}-\frac{\mathrm{d}e_{\mathcal{R},n}(0)}{\mathrm{d}\lambda}-\mathrm{Tr}\left(V\Gamma_2^\top C_{RM}\right)>\mathrm{Tr}\left(-\Sigma_{11}\Gamma_1^\top C_R\right.\\\left.-\Sigma_{12}\Gamma_2^\top C_R+\Sigma_{12}^\top\Gamma_1^\top C_{RM}+\Sigma_{22}\Gamma_2^\top C_{RM}\right)\end{aligned} \tag{D.11}$$

which coincides with (18). If (D.11) holds, $e_{\mathcal{R}}(\lambda)$ is strictly decreasing in the right neighborhood of $\lambda=0$ and $\lambda^*>0$. If Σ is diagonal, then $\Sigma_{12}=0$ and (D.11) is always satisfied.

E. Computation of Matrix Γ

We now show how to derive Γ from W and discuss the sign of its elements. For the sake of simplicity, we assume that the nominal weight matrix W^o is symmetric, which implies that both W^o and W are diagonalizable. If W is not diagonalizable, a similar derivation (with more tedious but conceptually identical calculations) can be carried out by considering the Jordan canonical form. This is because a straightforward extension of Lemma A.3 shows that W and Γ share the same (chain of) generalized eigenvectors.

Computation of $\boldsymbol{\Gamma}$*:* The derivative of L is (Lemma A.2)

$$\frac{dL}{d\lambda}=\tilde{L}-\lambda\tilde{L}\frac{d\tilde{L}^{-1}}{d\lambda}\tilde{L}=\tilde{L}-\lambda\tilde{L}W\tilde{L} \tag{E.1}$$

where $\tilde{L} \doteq (I-(1-\lambda)W)^{-1}$. Let λ_W and v_W an eigenvalue of W and its associated eigenvector, respectively, from Lemmas A.3–A.4, it follows that $\tilde{L}$ has eigenvalue $(1-(1-\lambda)\lambda_W)^{-1}$ with associated eigenvector v_W. Hence, straightforward computations yield

$$\frac{dL}{d\lambda}v_W = \frac{1-(1-(1-\lambda)\lambda_W)^{-1}\lambda\lambda_W}{(1-(1-\lambda)\lambda_W)}v_W. \tag{E.2}$$

In particular, the dominant eigenvector $v_W = \mathbb{1}$ (associated with $\lambda_W = 1$) is in the kernel of $dL/d\lambda$ for any λ. As for the other eigenvectors, by letting λ go to zero in (E.2), one gets

$$\Gamma v_W = (1-\lambda_W)^{-1} v_W. \tag{E.3}$$

Finally, the eigendecomposition of Γ is obtained from eigenvectors v_W and eigenvalues $(1-\lambda_W)^{-1}$, plus the kernel.

Sign of $\boldsymbol{\Gamma_1}$ and $\boldsymbol{\Gamma_2}$: As regards Γ_1, note that the upper-left block in $\overline{W}$ is identically zero and that L is a stochastic matrix for any value of λ: Hence, as λ becomes larger than zero, (some) elements in L_1 become positive, and thus, their derivative at $\lambda = 0^+$ is also positive.

As for Γ_2, define the following block partitions:

$$L = \left[\begin{array}{c|c} L_1 & L_2 \\ \hline 0 & I_M \end{array}\right] \tag{E.4}$$

with $W_1, L_1 \in \mathbb{R}^{R\times R}$ and $W_2, L_2 \in \mathbb{R}^{R\times M}$. Then, it holds

$$\frac{dL}{d\lambda} = \frac{1}{\lambda}L(I-WL) = \left[\begin{array}{c|c} \star & -L_1W_1L_2 - L_1 \\ \hline 0 & 0 \end{array}\right] \tag{E.5}$$

which implies, for any $\lambda \in (0,1)$

$$\frac{dL_{im}}{d\lambda} \leq 0,\ i \in \mathcal{R}, m \in \mathcal{M}. \tag{E.6}$$

In particular, the limit of the derivative of element L_{im} at $\lambda = 0^+$ is nonpositive in virtue of the theorem of sign permanence.

F. Proof of Proposition 2

Dependence on $\boldsymbol{V}$: Note that $e_{\mathcal{R},n}$ is independent of V.From (13), we highlight the contribution of v to the error $e_{\mathcal{R}}$ as follows:

$$e_{\mathcal{R},v} = \mathrm{Tr}\left(L_2 V L_2^\top\right) + \kappa \tag{F.1}$$

where κ does not depend on V and we use the block partition

$$L = \left[\begin{array}{c|c} L_1 & L_2 \\ \hline 0 & I \end{array}\right]. \tag{F.2}$$

The matrix L_2 is positive, see the work in [47] and discussion in Section IV-A). Then, if $V_1 \succ V_2$, it follows that $L_2V_1L_2^\top \succ L_2V_2L_2^\top$, and hence, the trace in (F.1) is strictly increasing with V.

Dependence on $\boldsymbol{Q}$: Note that $e_{\mathcal{R},v}$ is independent of Q. Let P_1 and P_2 denote the solutions of (10) with $Q = Q_1$ and $Q = Q_2$, respectively. If $Q_1 \succ Q_2$, then $\tilde{Q}_1 \succ \tilde{Q}_2$, and it is known that $P_1 \succ P_2$, from which the claim follows.

G. Proof of Proposition 3

Dependence on $\boldsymbol{V}$: In the following, we make the dependence of the error $e_{\mathcal{R}}$ on d_m explicit. Let us compute the partial derivative of the error first w.r.t. λ and then w.r.t. d_m

$$\frac{\partial^2 e_{\mathcal{R}}(\lambda, d_m)}{\partial d_m \partial \lambda} = \frac{1}{\lambda}\mathrm{Tr}\left(L\frac{\mathrm{d}\widetilde{\Sigma}(d_m)}{\mathrm{d}d_m}L^\top\left(I - W^\top L^\top\right)S_{\mathcal{R}}^\top S_{\mathcal{R}}\right). \tag{G.1}$$

It holds

$$L\frac{\mathrm{d}\widetilde{\Sigma}(d_m)}{\mathrm{d}d_m} = \left[\begin{array}{c|c} 0 & L_2S_m \\ \hline 0 & S_m \end{array}\right] \tag{G.2}$$

$$M \doteq I - W^\top L^\top = \left[\begin{array}{c|c} I_R - W_1^\top L_1^\top & 0 \\ \hline -W_2^\top L_1^\top - L_2^\top & 0 \end{array}\right] \tag{G.3}$$

$$L^\top M S_{\mathcal{R}}^\top S_{\mathcal{R}} = \left[\begin{array}{c|c} \star & 0 \\ \hline -L_2^\top W_1^\top L_1^\top - W_2^\top L_1^\top & 0 \end{array}\right] \tag{G.4}$$

and the argument of the trace in (G.1) is

$$\left[\begin{array}{c|c} -L_2S_mL_2^\top W_1^\top L_1^\top - L_2S_mW_2^\top L_1^\top & 0 \\ \hline \star & 0 \end{array}\right] \tag{G.5}$$

whose upper-left block is a negative matrix for all $\lambda \in (0,1)$ and is the zero matrix for $\lambda = 1$. Hence, the derivative of the consensus error w.r.t. λ (D.1) is strictly decreasing with d_m for any $\lambda \in (0,1)$. By continuity of (D.1), the minimum points of $e_{\mathcal{R}}(\lambda)$ are strictly increasing with d_m.

Dependence on $\boldsymbol{Q}$: Note that $e_{\mathcal{R},v}$ is independent of Q. We consider the derivative of $e_{\mathcal{R},n}$ w.r.t. λ

$$\frac{\mathrm{d}e_{\mathcal{R},n}(\lambda)}{\mathrm{d}\lambda} = -\sum_{i\in\mathcal{R}}\frac{2(1-\lambda)\left([W_{\mathcal{R}}PW_{\mathcal{R}}^\top]_{ii} + \tilde{Q}_{ii}\right)}{1-(1-\lambda)^2W_{ii}^2}. \tag{G.6}$$

Let $Q_1 \succ Q_2$, then it holds $P_1 \succ P_2$, which implies $[W_{\mathcal{R}}P_1W_{\mathcal{R}}^\top]_{ii} > [W_{\mathcal{R}}P_2W_{\mathcal{R}}^\top]_{ii}\ \forall i \in \mathcal{R}$. Further, it holds $\tilde{Q}_1 \succ \tilde{Q}_2$ and $[\tilde{Q}_1]_{ii} > [\tilde{Q}_2]_{ii}\ \forall i \in \mathcal{R}$. By combining such two facts, we conclude that (G.6) is strictly decreasing. The statement follows by the same argument of the case above.

H. Proof of Proposition 4

Dependence on $\boldsymbol{V}$: We expand (D.1) to highlight d_m

$$\frac{\mathrm{d}e_{\mathcal{R}}(\lambda)}{\mathrm{d}\lambda} = \frac{N_{mm}}{\lambda}d_m + \kappa \tag{H.1}$$

where N is the nonpositive matrix given by

$$N = -\left(L_2^\top W_1^\top + W_2^\top\right)L_1^\top L_2 \tag{H.2}$$

and κ does not depend on d_m. Note that $N_{mm} \neq 0$ because the opposite implies that the mth malicious agent has no interactions with regular agents. It follows that for any $\lambda < 1$, there exists $d_m \geq 0$ such that (H.1) is negative, which is given by the following inequality:

$$d_m > -\frac{\kappa\lambda}{N_{mm}} - \sum_{\substack{m'\in\mathcal{M}\\ m'\neq m}}\frac{N_{m'm'}}{N_{mm}}d_{m'}. \tag{H.3}$$

The claim follows by combining (H.3) with Proposition 3.

Dependence on $\boldsymbol{Q}$: We use (B.3) to highlight q_m in (D.1)

$$\frac{\mathrm{d}e_{\mathcal{R}}(\lambda)}{\mathrm{d}\lambda} = -\sum_{i\in\mathcal{R}}\frac{2(1-\lambda)\left([W_{\mathcal{R}}PW_{\mathcal{R}}^\top]_{ii}^2 + q_m\tilde{Q}_{im}^2\right)}{1-(1-\lambda)^2W_{ii}^2} + \kappa \tag{H.4}$$

where κ does not depend on q_m. Also, P is increasing with q_m and thus $[W_{\mathcal{R}} P_1 W_{\mathcal{R}}^\top]_{ii}$ also is. Hence, for any $\lambda < 1$, there exists $q_m \geq 0$ such that (H.4) is negative, which is given by the following inequality:

$$\sum_{i \in \mathcal{R}} \frac{2(1-\lambda)\left([W_{\mathcal{R}} P W_{\mathcal{R}}^\top]_{ii}^2 + q_m \tilde{Q}_{im}^2\right)}{1-(1-\lambda)^2 W_{ii}^2} > \kappa. \qquad \text{(H.5)}$$

The claim follows by combining (H.5) with Proposition 3.

Acknowledgment

The views and opinions expressed in this work are those of the authors and may not reflect those of the funding institutions.

References

[1] H. Farhangi, "The path of the smart grid," *IEEE Power Energy Mag.*, vol. 8, no. 1, pp. 18–28, Jan./Feb. 2010.

[2] F. Olivier, P. Aristidou, D. Ernst, and T. Van Cutsem, "Active management of low-voltage networks for mitigating overvoltages due to photovoltaic units," *IEEE Trans. Smart Grid*, vol. 7, no. 2, pp. 926–936, Mar. 2016.

[3] X. Lu, P. Wang, D. Niyato, D. I. Kim, and Z. Han, "Wireless networks with RF energy harvesting: A contemporary survey," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 2, pp. 757–789, Apr.–Jun. 2015.

[4] Z. Niu, X. S. Shen, Q. Zhang, and Y. Tang, "Space-air-ground integrated vehicular network for connected and automated vehicles: Challenges and solutions," *Intell. Converged Netw.*, vol. 1, no. 2, pp. 142–169, 2020.

[5] R. Chen and C. G. Cassandras, "Optimal assignments in mobility-on-demand systems using event-driven receding horizon control," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 3, pp. 1969–1983, Mar. 2020.

[6] S. Raskin, "Energy secretary says enemies are capable of shutting down us power grid," 2021. [Online]. Available: https://nypost.com/2021/06/06/energy-secretary-enemies-are-capable-of-shutting-down-power-grid/

[7] L. Borghese and S. Braithwaite, "Hackers block Italian COVID-19 vaccination booking system in 'most serious cyberattack ever'," 2021. [Online]. Available: https://www.cnn.com/2021/08/02/business/italy-hackers-covid-vaccine-intl/index.html

[8] C. Huang, R. Zhang, and S. Cui, "Optimal power allocation for wireless sensor networks with outage constraint," *IEEE Trans. Commun. Lett.*, vol. 3, no. 2, pp. 209–212, Apr. 2014.

[9] Y. Ye, L. Shi, X. Chu, H. Zhang, and G. Lu, "On the outage performance of SWIPT-based three-step two-way DF relay networks," *IEEE Trans. Veh. Technol.*, vol. 68, no. 3, pp. 3016–3021, Mar. 2019.

[10] S. Gupta, R. Kambli, S. Wagh, and F. Kazi, "Support-vector-machine-based proactive cascade prediction in smart grid using probabilistic framework," *IEEE Trans. Ind. Electron.*, vol. 62, no. 4, pp. 2478–2486, Apr. 2015.

[11] J. Qi, J. Wang, and K. Sun, "Efficient estimation of component interactions for cascading failure analysis by EM algorithm," *IEEE Trans. Power Syst.*, vol. 33, no. 3, pp. 3153–3161, May 2018.

[12] M. Rahnamay-Naeini and M. M. Hayat, "Cascading failures in interdependent infrastructures: An interdependent Markov-chain approach," *IEEE Trans. Smart Grid*, vol. 7, no. 4, pp. 1997–2006, Jul. 2016.

[13] Z. Huang, C. Wang, M. Stojmenovic, and A. Nayak, "Characterization of cascading failures in interdependent cyber-physical systems," *IEEE Trans. Comput.*, vol. 64, no. 8, pp. 2158–2168, Aug. 2015.

[14] Q. Yan, F. R. Yu, Q. Gong, and J. Li, "Software-defined networking (SDN) and distributed denial of service (DDOS) attacks in cloud computing environments: A survey, some research issues, and challenges," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 1, pp. 602–622, Jan.–Mar. 2016.

[15] N. Chaabouni, M. Mosbah, A. Zemmari, C. Sauvignac, and P. Faruki, "Network intrusion detection for IOT security based on learning techniques," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 3, pp. 2671–2701, Jul.–Sep. 2019.

[16] N. Agmon and D. Peleg, "Fault-tolerant gathering algorithms for autonomous mobile robots," *SIAM J. Comput.*, vol. 36, no. 1, pp. 56–82, 2006.

[17] I. D. Schizas, G. Mateos, and G. B. Giannakis, "Distributed LMS for consensus-based in-network adaptive processing," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2365–2382, Jun. 2009.

[18] L. Xiao, S. Boyd, and S.-J. Kim, "Distributed average consensus with least-mean-square deviation," *J. Parallel Distrib. Comput.*, vol. 67, no. 1, pp. 33–46, 2007.

[19] R. Xin, S. Pu, A. Nedić, and U. A. Khan, "A general framework for decentralized optimization with first-order methods," *Proc. IEEE*, vol. 108, no. 11, pp. 1869–1889, Nov. 2020.

[20] S. Shahrampour and A. Jadbabaie, "Distributed online optimization in dynamic environments using mirror descent," *IEEE Trans. Autom. Control*, vol. 63, no. 3, pp. 714–725, Mar. 2018.

[21] D. Varagnolo, F. Zanella, A. Cenedese, G. Pillonetto, and L. Schenato, "Newton-Raphson consensus for distributed convex optimization," *IEEE Trans. Autom. Control*, vol. 61, no. 4, pp. 994–1009, Apr. 2016.

[22] S. S. Kia, B. Van Scoy, J. Cortes, R. A. Freeman, K. M. Lynch, and S. Martinez, "Tutorial on dynamic average consensus: The problem, its applications, and the algorithms," *IEEE Control Syst. Mag.*, vol. 39, no. 3, pp. 40–72, Jun. 2019.

[23] S. Samarakoon, M. Bennis, W. Saad, and M. Debbah, "Distributed federated learning for ultra-reliable low-latency vehicular communications," *IEEE Trans. Commun.*, vol. 68, no. 2, pp. 1146–1159, Feb. 2020.

[24] K. Bonawitz et al., "Towards federated learning at scale: System design," in *Proc. Mach. Learn. Syst.*, 2019, pp. 374–388.

[25] R. Kieckhafer and M. Azadmanesh, "Reaching approximate agreement with mixed-mode faults," *IEEE Trans. Parallel Distrib. Syst.*, vol. 5, no. 1, pp. 53–63, Jan. 1994.

[26] 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.

[27] S. M. Dibaji and H. Ishii, "Consensus of second-order multi-agent systems in the presence of locally bounded faults," *Syst. Control. Lett.*, vol. 79, pp. 23–29, 2015.

[28] Y. Wang, H. Ishii, F. Bonnet, and X. Défago, "Resilient consensus against mobile malicious agents," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 3409–3414, 2020.

[29] J. Usevitch and D. Panagou, "Resilient leader-follower consensus to arbitrary reference values in time-varying graphs," *IEEE Trans. Autom. Control*, vol. 65, no. 4, pp. 1755–1762, Apr. 2020.

[30] Y. Shang, "Resilient consensus in multi-agent systems with state constraints," *Automatica*, vol. 122, 2020, Art. no. 109288.

[31] G. Wen, Y. Lv, W. X. Zheng, J. Zhou, and J. Fu, "Joint robustness of time-varying networks and its applications to resilient consensus," *IEEE Trans. Autom. Control*, vol. 68, no. 11, pp. 6466–6480, Nov. 2023.

[32] S. Sundaram and B. Gharesifard, "Consensus-based distributed optimization with malicious nodes," in *Proc. 53rd Annu. Allerton Conf. Commun. Control Comput.*, 2015, pp. 244–249.

[33] L. Su and N. H. Vaidya, "Byzantine-resilient multiagent optimization," *IEEE Trans. Autom. Control*, vol. 66, no. 5, pp. 2227–2233, May 2021.

[34] 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.

[35] S. M. Dibaji, M. Safi, and H. Ishii, "Resilient distributed averaging," in *Proc. Amer. Control Conf.*, 2019, pp. 96–101.

[36] W. Abbas, A. Laszka, and X. Koutsoukos, "Improving network connectivity and robustness using trusted nodes with application to resilient consensus," *IEEE Control Netw. Syst.*, vol. 5, no. 4, pp. 2036–2048, Dec. 2018.

[37] Y. Zhai, Z.-W. Liu, M.-F. Ge, G. Wen, X. Yu, and Y. Qin, "Trusted-region subsequence reduction for designing resilient consensus algorithms," *IEEE Trans. Netw. Sci. Eng.*, vol. 8, no. 1, pp. 259–268, Jan.–Mar. 2021.

[38] J. S. Baras and X. Liu, "Trust is the cure to distributed consensus with adversaries," in *Proc. 27th Mediterranean Conf. Control Autom.*, 2019, pp. 195–202.

[39] D. G. Mikulski, F. L. Lewis, E. Y. Gu, and G. R. Hudas, "Trust method for multi-agent consensus," *Proc. SPIE*, vol. 8387, pp. 146–159, 2012.

[40] M. Yemini, A. Nedić, A. J. Goldsmith, and S. Gil, "Characterizing trust and resilience in distributed consensus for cyberphysical systems," *IEEE Trans. Robot.*, vol. 38, no. 1, pp. 71–91, Feb. 2022.

[41] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *J. Math. Sociol.*, vol. 15, no. 3/4, pp. 193–206, 1990.

[42] L. Ballotta, G. Como, J. S. Shamma, and L. Schenato, "Competition-based resilience in distributed quadratic optimization," in *Proc. IEEE Conf. Decis. Control*, 2022, pp. 6454–6459.

[43] L. Li, C. Langbort, and J. Shamma, "An LP approach for solving two-player zero-sum repeated Bayesian games," *IEEE Trans. Autom. Control*, vol. 64, no. 9, pp. 3716–3731, Sep. 2019.

[44] L. Li and J. S. Shamma, "Efficient strategy computation in zero-sum asymmetric information repeated games," *IEEE Trans. Autom. Control*, vol. 65, no. 7, pp. 2785–2800, Jul. 2020.
[45] J. R. Marden, G. Arslan, and J. S. Shamma, "Cooperative control and potential games," *IEEE Trans. Syst., Man, Cybern., Part B*, vol. 39, no. 6, pp. 1393–1407, Dec. 2009.
[46] G. Como and F. Fagnani, "Robustness of large-scale stochastic matrices to localized perturbations," *IEEE Trans. Netw. Sci. Eng.*, vol. 2, no. 2, pp. 53–64, Apr.–Jun. 2015.
[47] S. E. Parsegov, A. V. Proskurnikov, R. Tempo, and N. E. Friedkin, "Novel multidimensional models of opinion dynamics in social networks," *IEEE Trans. Autom. Control*, vol. 62, no. 5, pp. 2270–2285, May 2017.
[48] D. Cartwright et al. *Studies in Social Power* (Ser. Publications of the Institute for Social Research: Research Center for Group Dynamics Series). Ann Arbor, MI, USA: Research Center Group Dyn., Inst. Soc. Res., Univ. Michigan, 1959.
[49] Y. Tian, P. Jia, A. MirTabatabaei, L. Wang, N. E. Friedkin, and F. Bullo, "Social power evolution in influence networks with stubborn individuals," *IEEE Trans. Autom. Control*, vol. 67, no. 2, pp. 574–588, Feb. 2022.
[50] M.V. Srighakollapu, R. K. Kalaimani, and R. Pasumarthy, "Optimizing network topology for average controllability," *Syst. Control Lett.*, vol. 158, 2021, Art. no. 105061.
[51] S. Gu et al., "Controllability of structural brain networks," *Nature Commun.*, vol. 6, no. 1, pp. 1–10, 2015.
[52] E. Nozari, F. Pasqualetti, and J. Cortés, "Heterogeneity of central nodes explains the benefits of time-varying control scheduling in complex dynamical networks," *J. Complex Netw.*, vol. 7, no. 5, pp. 659–701, 2019.
[53] L. Ballotta, G. Como, J. S. Shamma, and L. Schenato, "Can competition outperform collaboration? The role of misbehaving agents," 2022, *arXiv:2207.01346*.
[54] W. Ren and R. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Trans. Autom. Control*, vol. 50, no. 5, pp. 655–661, May 2005.
[55] P. Y. Sohouenou, P. Christidis, A. Christodoulou, L. A. Neves, and D. L. Presti, "Using a random road graph model to understand road networks robustness to link failures," *Int. J. Crit. Infrastructure Protection*, vol. 29, 2020, Art. no. 100353.
[56] M. E. Valcher and G. Parlangeli, "On the effects of communication failures in a multi-agent consensus network," in *Proc. Int. Conf. Syst. Theory Control Comput.*, 2019, pp. 709–720.
[57] A. Bojchevski and S. Günnemann, "Certifiable robustness to graph perturbations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, vol. 32. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2019/hash/e2f374c3418c50bc30d67d5f7454a5b4-Abstract.html

Luca Ballotta (Member, IEEE) received the master's degree in automation engineering and the Ph.D. degree in information engineering from the University of Padova, Padova, Italy, in 2019 and 2023, respectively.

He is currently a Research Fellow with the Department of Information Engineering, University of Padova. He was a visiting student with the Massachusetts Institute of Technology in 2020 and 2022. His research interests include multiagent systems and networked control systems under resource constraints, resilient distributed optimization, and learning-based safe control.

Dr. Ballotta is a recipient of the Young Author Prize at the 2020 IFAC World Congress.

Giacomo Como (Member, IEEE) received the B.Sc., M.S., and Ph.D. degrees in applied mathematics from Politecnico di Torino, Turin, Italy, in 2002, 2004, and 2008, respectively.

He is currently a Professor with the Department of Mathematical Sciences, Politecnico di Torino. He is also a Senior Lecturer with Automatic Control Department, Lund University, Lund, Sweden. He was a Visiting Assistant in research with Yale University, New Haven, CT, USA, during 2006–2007, and a Postdoctoral Associate with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA, from 2008 to 2011. His research interests include dynamics, information, and control in network systems with applications to cyberphysical systems, infrastructure networks, and social and economic networks.

Dr. Como is currently a Senior Editor for IEEE TRANSACTIONS ON CONTROL OF NETWORK SYSTEMS, an Associate Editor for *Automatica,* and the Chair of IEEE-CSS Technical Committee on Networks and Communications. He was an Associate Editor for IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING from 2015 to 2021 and IEEE TRANSACTIONS ON CONTROL OF NETWORK SYSTEMS from 2016 to 2022. He was the IPC Chair of the IFAC Workshop NecSys'15 and a Semiplenary Speaker at the International Symposium MTNS'16. He was the recipient of the 2015 George S. Axelby Outstanding Paper Award.

Jeff S. Shamma (Fellow, IEEE) received the Ph.D. degree in systems science and engineering from MIT, Cambridge, MA, USA, in 1988.

He is currently the Department Head and the Dobrovolny Chair with the Department of Industrial and Enterprise Systems Engineering, University of Illinois Urbana-Champaign, Champaign, IL, USA. He is the former Director of the Center of Excellence for NEOM Research and a Professor of Electrical Engineering with the King Abdullah University of Science and Technology (KAUST). Before joining KAUST, he was the Julian T. Hightower Chair in systems and control with the School of Electrical and Computer Engineering, Georgia Tech, and has held faculty positions with the University of Minnesota, The University of Texas at Austin, and the University of California, Los Angeles.

Dr. Shamma is the recipient of an NSF Young Investigator Award, the American Automatic Control Council Donald P. Eckman Award, and the Mohammed Dahleh Award. He is a Fellow of the International Federation of Automatic Control. He is the Editor-in-Chief for IEEE TRANSACTIONS ON CONTROL OF NETWORK SYSTEMS.

Luca Schenato (Fellow, IEEE) received the Dr. Eng. degree in electrical engineering from the University of Padova, Padova, Italy, in 1999, and the Ph.D. degree in electrical engineering and computer sciences from UC Berkeley, Berkeley, CA, USA, in 2003.

He was a Postdoctor in 2004 and a Visiting Professor during 2013–2014 with U.C. Berkeley. He is currently a Full Professor with Information Engineering Department, University of Padova. His interests include networked control systems, multiagent systems, wireless sensor networks, smart grids, and cooperative robotics.

Dr. Schenato is a recipient of the 2004 Researchers Mobility Fellowship by the Italian Ministry of Education, University and Research (MIUR), the 2006 Eli Jury Award in U.C. Berkeley, and the EUCA European Control Award in 2014. He was an Associate Editor for IEEE TRANSACTIONS ON AUTOMATIC CONTROLfrom 2010 to 2014 and is currently a Senior Editor for the IEEE TRANSACTIONS ON CONTROL OF NETWORK SYSTEMS and an Associate Editor for *Automatica*.