A Survey of Graph-Theoretic Approaches for Analyzing the Resilience of Networked Control Systems ⋆

Mohammad Pirani [a], Aritra Mitra [b], Shreyas Sundaram [c]

[a] *Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada.*

[b] *Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA*

[c] *Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA*

Abstract

As the scale of networked control systems increases and interactions between different subsystems become more sophisticated, questions of the resilience of such networks increase in importance. The need to redefine classical system and control-theoretic notions using the language of graphs has recently started to gain attention as a fertile and important area of research. This paper presents an overview of graph-theoretic methods for analyzing the resilience of networked control systems. We discuss various distributed algorithms operating on networked systems and investigate their resilience against adversarial actions by looking at the structural properties of their underlying networks. We present graph-theoretic methods to quantify the attack impact, and reinterpret some system-theoretic notions of robustness from a graph-theoretic standpoint to mitigate the impact of the attacks. Moreover, we discuss miscellaneous problems in the security of networked control systems which use graph-theory as a tool in their analyses. We conclude by introducing some avenues for further research in this field.

Key words: Networked Control Systems, Graph Theory, Resilient Distributed Algorithms

1 Introduction

Networked control systems (NCS) are spatially distributed control systems wherein the control loops are closed through a wired or wireless communication network. A schematic of an NCS is shown in Fig. 1. The network part of the control loop can be in the form of a sensor network, a controller network, or an actuator network. NCSs are often also highly dynamic, with subsystems, actuators, and sensors entering and leaving the network over time.

As opposed to monolithic systems where a single decision-maker (human or machine) possesses all available knowledge and information related to the system, NCSs typically involve multiple decision-makers, each with access to information that is not available to other decision-makers. Coordination among the decision-
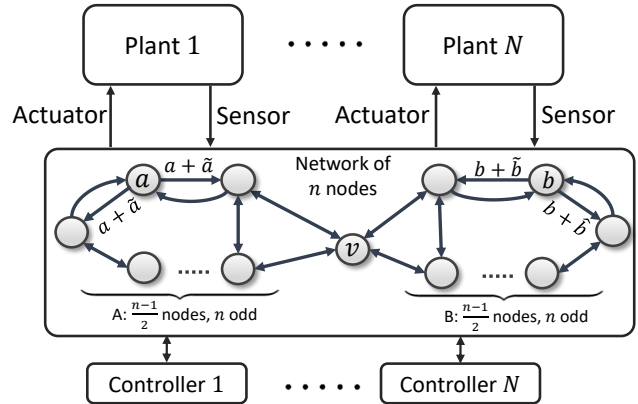

Fig. 1. A schematic figure of a networked control system.

makers in such systems can be achieved through the use of distributed algorithms. These algorithms are executed concurrently by each of the decision-makers, incorporating both their local information and any information received from other decision-makers in the network (via communication channels). However, as the size and the complexity of interconnections in NCSs increase, these distributed algorithms become more prone

to failures, degradation, and attacks. In this survey, we focus our attention on attacks performed by adversarial agents in NCSs. The subtle difference between a fault and an attack is that in the latter, the attacker uses their knowledge about the system model to target the vulnerable parts of an NCS to either maximize their impact, minimize their visibility, or minimize their effort to attack. As the attacker intelligently optimizes its actions, distributed algorithms have to be carefully designed to withstand adversarial actions, rather than more generic classes of faults considered by classical fault-tolerant control methods. Among various approaches to the resilience and security of distributed algorithms, the goal of this survey paper is to focus on the intersection of systems and control theory, graph theory, and communication and computation techniques to provide tools to solve security problems in NCSs. Notably, these are problems that are intractable using any set of methods in isolation.

1.1 Security of Networked Systems

As mentioned above, the main difference between adversarial actions and faults stems from the ability of the attacker to carefully target vulnerable parts of the system, particularly by learning about the system (and the deployed algorithms) before the attack. Cyber-attacks are thus classified into different categories based on their knowledge level and their ability to disrupt resources. In addition to system-theoretic properties, the structure of the underlying large-scale network plays a key role in determining specific attack behaviors. The following simple example shows the role of the graph structure in distinguishing attacks from faults.

Example 1 *Consider the NCS shown in Fig. 1. If the network experiences a single fault distributed uniformly at random, over the total number of n nodes in the network, then any given node becomes faulty with probability $p = \frac{1}{n}$. For large n, this value is small. On the other hand, an attacker that wishes to disconnect the graph can do so by targeting node v. This prevents nodes in set A from receiving true information from nodes in set B.*

The dichotomy between random failures and targeted node removal has been studied in the complex networks literature, particularly in the case of scale-free networks, which arise in a large number of applications [1]. These scale-free networks have a few nodes with large degrees (so called "hub nodes"); when such nodes are targeted for removal [2, 3], they cause the network to split into multiple parts.

Other than simply removing nodes from the network as discussed above, the attacker can perform more complex actions. One such action is to manipulate the dynamics of a subset of nodes in the network by injecting carefully crafted attack signals (or incorrect data) into their update rules. These attacks create a discrepancy between the information that the targeted nodes send to their neighbors and the true information they are supposed to send. Here, based on the communication medium, another network-theoretic feature of attacks arises which distinguishes them from faults: if the communication is point-to-point (as opposed to wireless broadcast), the attacker has the ability to send inconsistent information to different neighbors whereas benign faults result in incorrect but consistent information being sent through the network. An example is shown in Fig. 1 where the node whose true value is a sends a wrong but consistent value of $a + \tilde{a}$ to its neighbors, whereas the node with value b shares wrong and inconsistent information to neighbors, i.e., it sends $b + \tilde{b}$ to one node and $b + \hat{b}$ to the other.

1.2 Applications

Networked control systems have found numerous applications in today's engineering systems, some of which we mention briefly below.

Automotive and Intelligent Transportation Systems

The concept of connected vehicles, denoted by V2X, effectively transforms transportation systems into a network of processors. From this perspective, V2X refers to (i) each vehicle's wireless communications with its surroundings, including other vehicles, road infrastructure, and the cloud, and (ii) wired communication within each vehicle between several electronic control units (ECU) in a controller area network (CAN). At the higher level, the nodes represent vehicles, road infrastructure, the cloud, and any other component that is able to send information. The wireless communication between those nodes is modeled by edges, e.g., the dedicated short range communication system (DSRC) for V2V communication. For the wired intra-vehicle network, the nodes are ECUs and the edges are buses transmitting data. Wireless communications between vehicles and their surroundings are prone to intrusions. Several works have reported different types of attacks on inter-vehicular networks, along with defense mechanisms [4, 5]. Attacks on the wired intra-vehicle network or on safety critical ECUs (e.g., engine control unit, active steering, or brake system) can have life-threatening consequences. Moreover, stealthy attacks on the CAN Bus system, including an attack that embeds malicious code in a car's telematics unit and completely erases any evidence of its presence after a crash, have been reported [6]. Several other attacks using both wired and wireless communications have been studied in the literature. An example is an attack that can enter the vehicle via a Bluetooth connection through the radio ECU and be disseminated to other safety-critical ECUs [7].

Smart Buildings and IoT

Smart buildings are the integration of a vast number of sensors, smart devices, and appliances to control heating, ventilation and air conditioning, lighting, and home security systems through a building automation system (BAS). In a building automation networked system, the home appliances are the nodes and the wireless communications between them and between each appliance and the center are the edges. When home devices are connected to the internet, they form a key part of the *Internet of Things.* The objectives of building automation are to improve occupant comfort, ensure efficient operation of building systems, and reduce energy consumption and operating costs. However, the high level of connectivity, automation, and remote accessibility of devices also makes it critical to protect smart buildings against failures and attacks [8].

Power Systems

The traditional practice in power grids is to institute safeguards against physical faults using protective devices [9]. However, the emergence of new technologies including smart meters, smart appliances, and renewable energy resources, together with available communication technologies introduces further vulnerabilities to potential cyber-attacks [10]. Cyber attacks in power systems can happen at three different levels: (i) generation and transmission level, i.e., automatic generation control (AGC) loops [11], (ii) distribution level, e.g., islanded micro-grids [12], and (iii) market level, e.g., false data injection in electricity markets [13]. Further discussion on the cyber-security of power systems can be found in [14, 15].

Blockchain

A blockchain is a growing list of records, called blocks, that are linked together using cryptography. Each block contains a cryptographic hash of the previous block in a tree structure, called a Merkle tree. As each block contains information about the block previous to it, they form a chain, with each additional block reinforcing the ones before it. Blockchains are considered secure by design and exemplify a distributed computing system with high attack tolerance [16]. One of the most recognized applications of blockchains is in cryptocurrency, e.g., bitcoins.

There are several other applications for which the security of large-scale networked systems plays a crucial role. Examples include swarm robotics (with applications ranging from search and rescue missions to mining and agricultural systems) [17,18], water and waste-water networks [19], nuclear power plants [20], and social networks [21,22]. Further discussion on applications can be found in [23].

1.3 Early Works on the Resilience of Networked Control Systems

We provide a brief literature review on the security of networked control systems; starting from centralized (non-graph-theoretic) approaches and then followed up by graph-theoretic methods.

1.3.1 Centralized Resilient Control Techniques

Centralized fault-tolerant techniques have a long history in the systems and control community [24–27]. These early works focused on detecting and mitigating faults in the system and were not equipped to overcome adversarial actions. Distributed methods for fault-detection that were subsequently developed will be discussed in the next subsection. These later efforts focused on the goal of providing defense mechanisms for control systems against attacks in three layers of *attack prevention, attack detection*, and *attack endurance.* We provide a brief overview of these defense layers below. Detailed discussions can be found in [23, 28].

The first layer of defense is to prevent the attack from happening. Cryptography, network coding, model randomization, differential privacy, and moving target defense are among well-known attack prevention mechanisms used for control systems [29–34]. In many cases, however, it is not always possible to prevent all attacks since regular users may not be distinguishable from the intruders. In those cases, the second layer comes into play which aims to detect and isolate the attack. Observer-based techniques have been proposed to detect the attacks which compare the state estimates under the healthy and the attacked cases [35]. When the control system does not satisfy the required observability conditions, coding-theory, e.g., parity check methods, can be used to detect the attacks [36]. In some cases, an adversary delivers fake sensor measurements to a system operator to conceal its effect on the plant. Certain types of such attacks, referred to as "replay attacks", have been addressed by introducing physical watermarking (by adding a Gaussian signal to the control input) to bait the attacker to reveal itself [37]. The sub-optimality of the resulting control action is the cost paid to detect the attacks in those cases. In addition to the above model-based techniques, anomaly detection methods have been proposed based on machine learning techniques. For example, Neural Networks (NNs) and Bayesian learning have been studied for anomaly detection in the context of security [38–40]. When attack detection is not possible, the system must be at least resilient enough to withstand the attacks or mitigate the impact of the attack. Probabilistic methods for attack endurance in NCSs for both estimation and control were studied in [41–44]. Redundancy-based approaches are used to bypass the attacks by using the healthy redundant parts [45–48]. Such redundancy in

large-scale systems can be in the form of adding parts, e.g., extra sensors, or the connection between the parts, i.e., network connectivity. On the other hand, several control-theoretic methods have also been proposed to mitigate the attack impact, including event-triggered control for tackling denial of service attacks [49, 50]. Robust control techniques have also been shown to be useful tools to mitigate the attack impact [51].

1.3.2 Resilient Distributed Techniques

The theory of distributed algorithms has a long history in computer science with a variety of applications in telecommunications, scientific computing, distributed information processing, and real-time process control [52, 53]. In the control systems community, distributed control algorithms have been studied for several decades [53–57], with an explosion of interest in recent decades due to their applications in distributed coordination of multi-agent systems, formation control of mobile robots, state estimation of power-grids, smart cities, and intelligent transportation systems, and distributed energy systems [58–61].

The earliest works on the security of distributed algorithms can be found in the computer science literature [62, 63], typically with the focus on simple network topologies (such as complete graphs). One of the main approaches to address the resilience and security of distributed systems is to leverage the physical redundancy that the network connectivity provides. Hence, resilient distributed estimation and control algorithms usually use (different types of) network connectivity measures to quantify the resilience against certain adversarial actions [64–66]. To do this, system-theoretic notions, such as controllability (or observability) and detectability are reinterpreted in terms of graph-theoretic quantities with the help of tools such as algebraic graph theory or structured systems theory [35].

In parallel to control-theoretic approaches, several other approaches to the design of resilient distributed algorithms have been developed. Recently, federated learning techniques have found numerous applications in computer networks. The general principle of federated learning is to train local models on local data samples and exchange parameters, e.g., the weights of a deep neural network, between these local models at some frequency to generate a global model [67,68]. With federated learning, only machine learning parameters are exchanged. These parameters can be encrypted before sharing between learning rounds to guarantee privacy. For the cases where these parameters may still leak information about the underlying data samples, e.g., by making multiple specific queries on specific datasets, secure aggregation techniques have been developed [69–72].

When a NCS becomes larger in scale, the notions of attack prevention, detection, and resilience discussed above depend more on the interconnections between components (sensors, actuators, or controllers) in the network. With this in mind, the focus of this survey paper is to present the theoretical works in the literature on graph-theoretic interpretations of the security in NCSs.

Related Survey Papers. There are some recently published survey papers on related topics, including [73], which provides an overview of security and privacy in a variety of cyber-physical systems (e.g., smart-grids, manufacturing systems, healthcare units, industrial control systems, etc.); [74], which focuses on resilient consensus problems; [75], which focuses on distributed statistical inference and machine learning under attacks; and [76], which discusses the applications of resilient distributed algorithms to multi-robot systems. Compared to [73] where the exposition is essentially of a qualitative nature, our survey provides a mathematical treatment of security in networked systems, covering the necessary technical background in linear algebra, graph theory, dynamical systems, and structured systems theory. Our paper differs from [74], [75], and [76] in that it has a much broader scope: we provide a detailed discussion of graph-theoretic measures for the resilience of a variety of distributed algorithms, special cases of which include consensus and distributed statistical inference. Moreover, we provide a comprehensive view of the role of various connectivity measures on the resilience levels for distributed algorithms. We also discuss ways to maintain a desired level of resilience when the network loses connectivity.

2 Mathematical Preliminaries

2.1 Graph Theory

A weighted graph is a pair $(\mathcal{G}, w)$ where $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ is a directed graph in which $\mathcal{V}$ is the set of vertices (or nodes) [1] and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges and $w : \mathcal{E} \to \mathbb{R}$ is a weight function. In particular, $(j, i) \in \mathcal{E}$ if and only if there exists an edge from j to i with some weight $w_{ji} \neq 0$. Graph $\mathcal{G}$ is undirected if $(j, i) \in \mathcal{E}$ implies $(i, j) \in \mathcal{E}$ and $w_{ij} = w_{ji}$. The in-neighbors of vertex $i \in \mathcal{V}$ are denoted $\mathcal{N}_i^{\text{in}} \triangleq \{j \in \mathcal{V} \mid (j, i) \in \mathcal{E}, j \neq i\}$. Similarly, the out-neighbors of i are $\mathcal{N}_i^{\text{out}} \triangleq \{j \in \mathcal{V} \mid (i, j) \in \mathcal{E}, j \neq i\}$. The in-degree (or simply degree) of node i is $d_i = \sum_j w_{ji}$. The minimum degree of a graph $\mathcal{G}$ is denoted by $d_{\min}(\mathcal{G}) = \min_{i \in \mathcal{V}} d_i$. If $(i, i) \in \mathcal{E}$, then i is said to have a self loop (but it is not counted in the degree of i). A subgraph of $\mathcal{G}$ is a graph $\bar{\mathcal{G}} = \{\bar{\mathcal{V}}, \bar{\mathcal{E}}\}$ with $\bar{\mathcal{V}} \subseteq \mathcal{V}$ and $\bar{\mathcal{E}} \subseteq \mathcal{E}$. A subgraph is induced if it is obtained from $\mathcal{G}$ by deleting a set of vertices (and all edges coming into and out of those vertices), but leaving all other edges

---

[1] Throughout this paper, we use words node, vertex, and agent interchangeably.
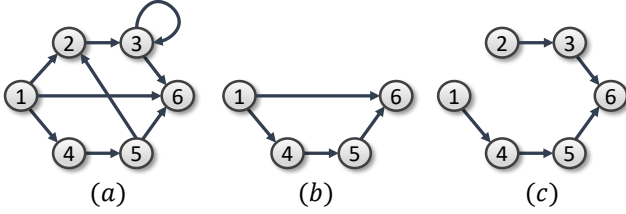
Fig. 2. (a) An example of a digraph $\mathcal{G}$, (b) an induced subgraph of $\mathcal{G}$, (c) a spanning subgraph of $\mathcal{G}$.

intact. The subgraph $\mathcal{H}$ is called spanning if it contains all vertices of $\mathcal{G}$, i.e., $\bar{\mathcal{V}} = \mathcal{V}$. An example of a digraph $\mathcal{G}$ together with an induced and a spanning subgraph of $\mathcal{G}$ is shown in Fig. 2.

2.1.1 Paths and Cycles in Graphs

For subsets $\mathcal{F}, \mathcal{B} \subset \mathcal{V}$, a path from $\mathcal{F}$ to $\mathcal{B}$ is a sequence of vertices $v_1, v_2, ..., v_t$ where $v_1 \in \mathcal{F}$, $v_t \in \mathcal{B}$, and $(v_j, v_{j+1}) \in \mathcal{E}$ for $1 \leq j \leq t-1$. A cycle is a path where $v_t = v_1$. A simple path contains no repeated vertices. A *directed acyclic graph* is a digraph with no cycles. For a subset $\mathcal{X} \subset \mathcal{V}$, an $\mathcal{X}$-rooted path (respectively $\mathcal{X}$-topped path) is a path which starts from a node $v \in \mathcal{X}$ (respectively ends at some $v \in \mathcal{X}$) . Two paths are disjoint if they have no common vertices and two paths are internally disjoint if they have no common vertices except for possibly the starting and ending vertices. A set of paths $P_1, P_2, ..., P_r$ are (internally) vertex disjoint if the paths are pairwise (internally) vertex disjoint. For example, the two paths $P_1 : 1, 2, 3, 6$ and $P_2 : 1, 4, 5, 6$ in Fig. 2 (a) are internally vertex disjoint paths between nodes 1 and 6. Given two subsets $\mathcal{X}_1, \mathcal{X}_2 \subset \mathcal{V}$, a set of r vertex disjoint paths, each with start vertex in $\mathcal{X}_1$ and end vertex in $\mathcal{X}_2$, is called an r-linking from $\mathcal{X}_1$ to $\mathcal{X}_2$.[2] The length of a path is the summation of the edge weights in the path. The distance between a pair of nodes i and j is the length of the shortest path between i and j. The *effective resistance*, $\mathfrak{R}_{ij}$, between two vertices i and j in a graph is the equivalent resistance between these two vertices when we treat the resistance of each edge e as $\frac{1}{w_e}$, where w_e is the edge weight.

2.1.2 Graph Redundancy Measures

A graph $\mathcal{G}$ is called *strongly connected* if there is a path between each pair of vertices $i, j \in \mathcal{V}$. A graph is said to be disconnected if there exists at least one pair of vertices $i, j \in \mathcal{V}$ such that there is no path from i to j. Other than the above binary measures of connectivity, there are several other graph connectivity measures, some of which are mentioned bellow.

---

[2] There are various algorithms to find linkings, such as the Ford-Fulkerson algorithm, which has run-time polynomial in the number of vertices [77].
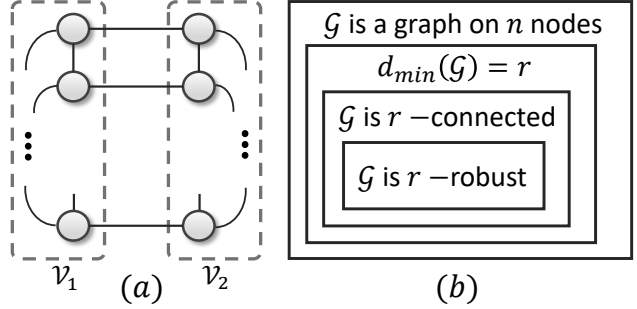

Fig. 3. (a) Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \mathcal{V}_1 \cup \mathcal{V}_2$ and $|\mathcal{V}_1| = |\mathcal{V}_2| = \frac{n}{2}$. $\mathcal{V}_1$ and $\mathcal{V}_2$ are complete graphs. There is a one to one connection between nodes in $\mathcal{V}_1$ and nodes in $\mathcal{V}_2$. (b) Relationships between different graph connectivity measures.

- **Vertex and Edge Connectivity:** A *vertex-cut* in a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ is a subset $\mathcal{S} \subset \mathcal{V}$ of vertices such that removing the vertices in $\mathcal{S}$ (and any resulting associated edges) from the graph causes the remaining graph to be disconnected. A (j, i)-*cut* in a graph is a subset $\mathcal{S}_{ij} \subset \mathcal{V}$ such that if the nodes $\mathcal{S}_{ij}$ are removed, the resulting graph contains no path from vertex j to vertex i. Let κ_{ij} denote the size of the smallest (j, i)-cut between any two vertices j and i. The graph $\mathcal{G}$ is said to have *vertex connectivity* $\kappa(\mathcal{G})$ (or to be κ-*vertex connected*) if $\kappa_{ij} \geq \kappa$ for all $i, j \in \mathcal{V}$. Similarly, the *edge connectivity* $e(\mathcal{G})$ of a graph $\mathcal{G}$ is the minimum number of edges whose deletion disconnects the graph. The vertex connectivity, edge connectivity, and minimum degree satisfy

$$\kappa(\mathcal{G}) \leq e(\mathcal{G}) \leq d_{\min}(\mathcal{G}). \qquad (1)$$

- **Graph Robustness [65, 78]:** For some $r \in \mathbb{N}$, a subset S of nodes in the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is said to be r-*reachable* if there exists a node $i \in S$ such that $|\mathcal{N}_i^{\text{in}} \setminus S| \geq r$. Graph $\mathcal{G}$ is said to be r-*robust* if for every pair of nonempty, disjoint subsets $\mathcal{X}_1, \mathcal{X}_2 \subseteq \mathcal{V}$, either $\mathcal{X}_1$ or $\mathcal{X}_2$ is r-reachable. If $\mathcal{G}$ is r-robust, then it is at least r-connected.

  For some $r, s \in \mathbb{N}$, a graph is said to be (r, s)-robust if for all pairs of disjoint nonempty subsets $\mathcal{X}_1, \mathcal{X}_2 \subset \mathcal{V}$, at least one of the following conditions hold:
  (i) All nodes in $\mathcal{X}_1$ have at least r neighbors outside $\mathcal{X}_1$.
  (ii) All nodes in $\mathcal{X}_2$ have at least r neighbors outside $\mathcal{X}_2$.
  (iii) There are at least s nodes in $\mathcal{X}_1 \cup \mathcal{X}_2$ that each have at least r neighbors outside their respective sets.
  Based on the above definitions, $(r, 1)$-robustness is equivalent to r-robustness.

The gap between the robustness and node connectivity (and minimum degree) parameters can be arbitrarily large, as illustrated by the graph $\mathcal{G}$ in Fig. 3 (a). While the minimum degree and node connectivity of the graph $\mathcal{G}$ is $n/2$, it is only 1-robust (consider subsets $\mathcal{V}_1$ and $\mathcal{V}_2$). In Fig. 3 (b), a Venn diagram is used to show the relationship between various graph connectivity measures.

2.2 Matrix Terminology

For a matrix $\boldsymbol{M}_{m \times n}$ with $n \leq m$, the singular values are ordered as $\sigma_1(\boldsymbol{M}) \leq \sigma_2(\boldsymbol{M}) \leq \ldots \leq \sigma_n(\boldsymbol{M})$. When M is a square matrix, the real parts of the eigenvalues are ordered as $\Re(\lambda_1(\boldsymbol{M})) \leq \Re(\lambda_2(\boldsymbol{M})) \leq \ldots \leq \Re(\lambda_n(\boldsymbol{M}))$. Matrix $\boldsymbol{M}$ is called nonnegative if its elements are nonnegative, and it is a *Metzler* matrix if its off-diagonal elements are nonnegative. We use $\mathbf{e}_i$ to indicate the i-th vector of the canonical basis.

2.3 Spectral Graph Theory and Linear Systems

The adjacency matrix of a graph of n nodes is denoted by $\boldsymbol{A}_{n \times n}$, where $\boldsymbol{A}_{ij} = w_{ij}$ if $(j, i) \in \mathcal{E}$ with the edge weight w_{ij} and $\boldsymbol{A}_{ij} = 0$ otherwise. The Laplacian matrix of the graph is $\boldsymbol{L} \triangleq \boldsymbol{D} - \boldsymbol{A}$, where $\boldsymbol{D} = \mathrm{diag}(d_1, d_2, \ldots, d_n)$. The real parts of the Laplacian eigenvalues are nonnegative and are denoted by $0 = \Re(\lambda_1(\boldsymbol{L})) \leq \Re(\lambda_2(\boldsymbol{L})) \leq \ldots \leq \Re(\lambda_n(\boldsymbol{L}))$.[3] The second smallest eigenvalue of the Laplacian matrix, $\lambda_2(\boldsymbol{L})$, is called the *algebraic connectivity* of the graph and is greater than zero if and only if $\mathcal{G}$ is a connected graph. Moreover, we always have [81]

$$\lambda_2(\boldsymbol{L}) \leq \kappa(\mathcal{G}). \tag{2}$$

Given a connected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, an orientation of the graph $\mathcal{G}$ is defined by assigning a direction (arbitrarily) to each edge in $\mathcal{E}$. For graph $\mathcal{G}$ with m edges, labeled as $e_1, e_2, \ldots, e_m$, its node-edge incidence matrix $\mathcal{B}(\mathcal{G}) \in \mathbb{R}^{n \times m}$ is defined as

$$[\mathcal{B}(\mathcal{G})]_{kl} = \begin{cases} 1 & \text{if node } k \text{ is the head of edge } l, \\ -1 & \text{if node } k \text{ is the tail of edge } l, \\ 0 & \text{otherwise.} \end{cases}$$

The graph Laplacian satisfies $L = \mathcal{B}(\mathcal{G})\mathcal{B}(\mathcal{G})'$.

A discrete-time linear time-invariant system is represented in the state-space form as follows:

$$\begin{aligned} \boldsymbol{x}[t+1] &= \boldsymbol{W}\boldsymbol{x}[t] + \boldsymbol{B}\boldsymbol{u}[t], \\ \boldsymbol{y}[t] &= \boldsymbol{C}\boldsymbol{x}[t] + \boldsymbol{D}\boldsymbol{u}[t], \end{aligned} \tag{3}$$

where $\boldsymbol{x} \in \mathbb{R}^n$ is the state vector, $\boldsymbol{u} \in \mathbb{R}^m$ is the vector of m inputs, $\boldsymbol{y} \in \mathbb{R}^q$ is the vector of q outputs, and $\boldsymbol{W} \in \mathbb{R}^{n \times n}$, $\boldsymbol{B} \in \mathbb{R}^{n \times m}$, $\boldsymbol{C} \in \mathbb{R}^{q \times n}$, and $\boldsymbol{D} \in \mathbb{R}^{q \times m}$, are called state, input, output, and feed-forward matrices, respectively. Similarly, the state-space model of a continuous-time linear system is given by

$$\begin{aligned} \dot{\boldsymbol{x}} &= \boldsymbol{W}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u}, \\ \boldsymbol{y} &= \boldsymbol{C}\boldsymbol{x} + \boldsymbol{D}\boldsymbol{u}. \end{aligned} \tag{4}$$

---

[3] We consider Laplacian matrices for graphs with positive edge weights. Graph spectra for negative edge weights have been studied in [79, 80].

A state space form of a linear system is compactly represented as $(\boldsymbol{W}, \boldsymbol{B}, \boldsymbol{C}, \boldsymbol{D})$ or $(\boldsymbol{W}, \boldsymbol{B}, \boldsymbol{C})$ for cases where there is no feed-forward term. A linear system is called (internally) positive if its state and output are nonnegative for every non-negative input and every non-negative initial state. A continuous-time linear system $(\boldsymbol{W}, \boldsymbol{B}, \boldsymbol{C})$ is positive if and only if $\boldsymbol{W}$ is a Metzler matrix and $\boldsymbol{B}$ and $\boldsymbol{C}$ are non-negative element-wise [82]. Moreover, for such a positive system with transfer function $G(s) = \boldsymbol{C}(s\boldsymbol{I}_n - \boldsymbol{W})^{-1}\boldsymbol{B}$, the system $\mathcal{H}_\infty$ norm is obtained from the DC gain of the system, i.e., $\|G\|_\infty = \sigma_n(G(0))$, where σ_n is the maximum singular value of matrix $G(0)$.

2.4 Structured Systems Theory

Consider the linear time-invariant system (3). With this system, associate the matrices $\boldsymbol{W}_\lambda \in \{0, \lambda\}^{n \times n}$, $\boldsymbol{B}_\lambda \in \{0, \lambda\}^{n \times m}$, $\boldsymbol{C}_\lambda \in \{0, \lambda\}^{q \times n}$, and $\boldsymbol{D}_\lambda \in \{0, \lambda\}^{q \times m}$. Specifically, an entry in these matrices is zero if the corresponding entry in the system matrices is equal to zero, and the matrix entry is a free parameter (denoted by λ) otherwise. This type of representation of (3) shows the structure of the linear system regardless of the specific values of the elements in the matrices. Thus, it is called a *structured system* and can be equivalently represented by a directed graph $\mathcal{G} = \{\mathcal{X}, \mathcal{U}, \mathcal{Y}, \mathcal{E}_{\mathcal{X}\mathcal{X}}, \mathcal{E}_{\mathcal{X}\mathcal{Y}}, \mathcal{E}_{\mathcal{U}\mathcal{X}}, \mathcal{E}_{\mathcal{U}\mathcal{Y}}\}$, where

- $\mathcal{X} \triangleq \{x_1, x_2, \ldots, x_n\}$ is the set of states;
- $\mathcal{Y} \triangleq \{y_1, y_2, \ldots, y_q\}$ is the set of measurements;
- $\mathcal{U} \triangleq \{u_1, u_2, \ldots, u_m\}$ is the set of inputs;
- $\mathcal{E}_{\mathcal{X}\mathcal{X}} = \{(x_j, x_i) | \boldsymbol{W}_{ij} \neq 0\}$ is the set of edges corresponding to interconnections between the state vertices;
- $\mathcal{E}_{\mathcal{U}\mathcal{X}} = \{(u_j, x_i) | \boldsymbol{B}_{ij} \neq 0\}$ is the set of edges corresponding to connections between the input vertices and the state vertices;
- $\mathcal{E}_{\mathcal{X}\mathcal{Y}} = \{(x_j, y_i) | \boldsymbol{C}_{ij} \neq 0\}$ is the set of edges corresponding to connections between the state vertices and the output vertices;
- $\mathcal{E}_{\mathcal{U}\mathcal{Y}} = \{(u_j, y_i) | \boldsymbol{D}_{ij} \neq 0\}$ is the set of edges corresponding to connections between the input vertices and the output vertices.

A structured system is said to have a certain property, e.g., controllability or invertibility, if that property holds for at least one numerical choice of free parameters λ in the system. The following theorem introduces graphical conditions for structural controllability and observability of linear systems.

Theorem 1 ([83]) *The pair $(\boldsymbol{W}, \boldsymbol{B})$ (resp. $(\boldsymbol{W}, \boldsymbol{C})$) is structurally controllable (resp. observable) if and only if the graph $\mathcal{G} = \{\mathcal{X}, \mathcal{U}, \mathcal{E}_{\mathcal{X}, \mathcal{X}}, \mathcal{E}_{\mathcal{U}, \mathcal{X}}\}$ satisfies both of the following properties:*

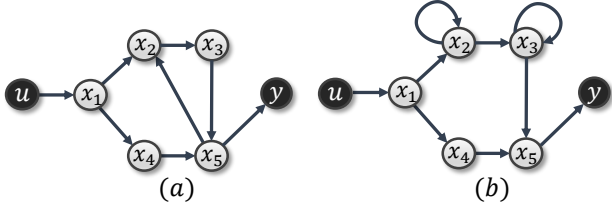*(i) Every state vertex $x_i \in \mathcal{X}$ can be reached by a path*

Fig. 4. (a) A structurally controllable but not observable graph, (b) A structurally controllable and observable graph.

from (resp. has a path to) some input vertex (resp. some output vertex).

(ii) $\mathcal{G}$ contains a subgraph that is a disjoint union of cycles and $\mathcal{U}$-rooted paths (resp. $\mathcal{Y}$-topped paths), which covers all of the state vertices.

In Sec. 8, we will revisit structured systems by discussing structural conditions for the system to be invertible.

Example 2 *The graph shown in Fig. 4 (a) is structurally controllable as it satisfies both conditions in Theorem 1. However, it is not structurally observable: condition (ii) does not hold since the cycle and $\mathcal{Y}$-topped paths are not disjoint. Graph (b) is structurally controllable and observable. The set of disjoint $\mathcal{U}$-rooted paths (respectively $\mathcal{Y}$-topped paths) and cycles is $\mathcal{P} = \big\{\{u, x_1, x_4, x_5\}, \{x_2\}, \{x_3\}\big\}$ (respectively $\mathcal{C} = \big\{\{x_1, x_4, x_5, y\}, \{x_2\}, \{x_3\}\big\}$).*

3   Notions of Resilience

In this section, we discuss the notions of resilience in NCSs which are sought in various distributed algorithms. To classify the resilience against each type of adversary, the first step is to distinguish a regular node from an adversarial one in NCSs.

In distributed algorithms, each agent is given an updating rule which is a function of its own states (and local information) and the states and information obtained from its neighbors. In many distributed algorithms, the agents only need to know their own updating rule and their neighbor set. However, in some problems, the agents not only need to know their neighbor sets, they must also be aware of the entire network topology (and maybe the updating policy of other agents). Based on this, the computational burden and the required storage limits for distributed algorithms may be different. Regular agents (nodes) in a NCS are those who obey the prescribed updating rule. The objective of a regular node can be either to calculate an exact desired value (e.g., a function of the state of other nodes) or an approximation of that value. The required precision of calculation depends on many factors including the cost of computing the exact value. Deviations from normal behaviour can be considered as either a fault or an adversarial action (attack). Faults are those which happen uninten-

tionally and (often) randomly with a given distribution. Communication dropouts, disturbances and noise, and failures due to unmodeled physics are all categorized as faults. Attacks, as discussed in Section 1.1, are failures which intentionally happen to the system and can be viewed as decisions of an intelligent intruder. From the network's perspective, the attacker (or an adversarial agent) is one who intentionally disregards the prescribed updating rule: the attacker updates its state and sends it to its neighbors in an arbitrary (and potentially worst case) manner. Since this type of deliberate adversarial behaviour is the focus of the current paper, we further classify them in the following definition.

Definition 1 (Malicious vs. Byzantine) *An adversarial agent is called malicious if it updates its state in an arbitrary manner. Thus, it sends incorrect but consistent values to all of its out-neighbors at each time-step. An adversarial agent is Byzantine if it can update its state arbitrarily and is capable of sending inconsistent values to different neighbors at each time-step.*

Both malicious and Byzantine agents are allowed to know the entire network topology, the local information of all agents, and the algorithms executed by all agents. Furthermore, both malicious and Byzantine agents are allowed to collude amongst themselves to select their actions. An example of malicious and Byzantine agents was discussed in Section 1.1. Based on the above definition, every malicious node is Byzantine but not vice versa. Malicious attacks may be appropriate for wireless broadcast models of communication or when the state of an agent is directly sensed by its neighbors (e.g., via cameras), whereas Byzantine attacks follow the wired point-to-point model of communication.

In return for providing so much power to the adversarial agents, it is typical to assume a bound on the number of such agents. The following definitions quantify the maximum number of tolerable attacks in a given network.

Definition 2 (f-total and f-local sets) *For $f \in \mathbb{N}$, a set $\mathcal{C} \subset \mathcal{V}$ is said to be f-total if it contains at most f nodes in the network, i.e., $|\mathcal{C}| \leq f$. A set $\mathcal{C} \subset \mathcal{V}$ is f-local if it contains at most f nodes in the neighborhood of each node outside that set, i.e., $|\mathcal{N}_i^{\text{in}} \cap \mathcal{C}| \leq f$ for all $i \in \mathcal{V} \setminus \mathcal{C}$.*

Definition 3 (f-local adversarial model) *For $f \in \mathbb{N}$, a set $\mathcal{F}$ of adversarial nodes is f-locally bounded if $\mathcal{F}$ is an f-local set.*

The set of adversarial nodes $\{a, b\}$ in Fig. 1 is a 2-total and 1-local set. Thus, it is 1-locally bounded. Note that every f-total set is also an f-local set but not vice versa. The f-total adversarial model is predominant in the literature on resilient distributed algorithms [62, 84, 85]. However, in order to allow the number of adversarial agents to potentially scale with the network, several of

the algorithms discussed in this survey allow the adversarial set to be f-local.

Based on the above discussions, the notion of resilience can be stated as follows.

> **Resilient Distributed Algorithm:** Under a given adversarial model (e.g., f-locally bounded or f-total, malicious or Byzantine), a distributed algorithm operating on network $\mathcal{G}$ is called *resilient* if each regular node in $\mathcal{G}$ can compute its desired value (within some specified tolerance) despite the actions of the adversarial agents in $\mathcal{G}$.

Thus, various specific notions of resilience can be considered, based on the above definition and depending on the type and number of adversaries, and the desired value computed by the regular agents.

4 Connectivity: The Earliest Measure of Resilience

In this section, we first discuss the role of network connectivity in reliable information dissemination over networks. Then, with the help of structured systems theory, we tie together the traditional graph property of connectivity with system-theoretic notions to find conditions for reliable calculation of node values in a network.

4.1 Connectivity as a Measure of Resilience

In the following example, we see how the existence of redundant paths between a pair of nodes can facilitate reliable transmission of information between those nodes.

Example 3 *For the graph shown in Fig. 5 (a), suppose that v_1 tries to obtain the true value of v_4, i.e., $\psi_4[0] = 2$. This value can be transmitted to v_1 through v_2 and v_3.[4] Suppose that v_2 is malicious and pretends that $\psi_4[0]$ has a value other than its true value. In this case, as v_1 receives inconsistent information from v_2 and v_3, it can not conclude which value is the true one. In this case, a redundant path can serve as a tie breaker and help v_1 to obtain the true value, Fig. 5 (b).*

One can generalize the observation given in Example 3 by saying that if there are f adversarial nodes in the network, there should be $2f + 1$ disjoint paths between any given pair of nodes in order to make sure that information can be transmitted reliably between those nodes. The number of disjoint paths between node pairs

---

[4] This can be done using a flooding algorithm, i.e., each node reads and stores their neighbors' values and broadcasts them to their out-neighbors in the next time step.
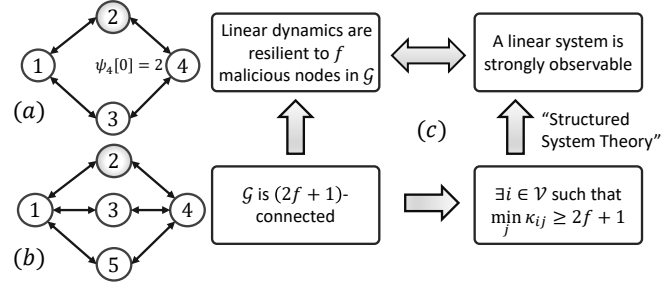

Fig. 5. (a) A graph with (a) two disjoint paths between nodes v_1 and v_4, and (b) three disjoint paths between nodes v_1 and v_5. (c) Schematics of interconnections between system and graph properties.

is related to the vertex connectivity via Menger's theorem [86].

Theorem 2 (Menger's Theorem) *Graph $\mathcal{G}$ has vertex connectivity r if and only if there are r internally vertex disjoint paths between each pair of nodes in $\mathcal{G}$.*

This fundamental observation that $2f+1$ connectivity is required to overcome f adversarial agents is classical in the computer science literature [62]. In the following subsection, we describe how the same result (namely that a connectivity of $2f+1$ is required to reliably exchange information in networks despite malicious agents) arises in the context of linear iterative dynamics for information dissemination in networks. In the process, this will introduce the use of zero-dynamics (and strong observability) together with structured systems theory as a means to analyze the resilience of linear dynamics on graphs.

4.2 A System-Theoretic Perspective on Resilient Exchange of Information

In this subsection, we discuss reliable calculation of node values in a network in the presence of adversaries. In this setting, each agent $i \in \mathcal{V}$ tries to gather the values (measurements, positions, votes, or other data) of all other agents, despite the actions of malicious agents in the network. These values can be later used to calculate any arbitrary function of the agents' values. Here, we consider a broadcast model of communication where each node transmits the same value to all its neighbors. Hence, the adversarial agents are malicious, but not Byzantine. Our goal is to show that the topology of the network (specifically, its connectivity) completely characterizes the resilience of linear iterative strategies to malicious behavior. To this end, we first formally introduce the model of a distributed system under attack.

Distributed System Model Under Attack: Consider a network of n agents (or processors) whose communication is represented by a time-varying graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}[t]\}$. Suppose that each agent $i \in \mathcal{V}$ begins with

some (possibly private) initial value [5] $\psi_i[0] \in \mathbb{R}$ and updates its value over time according to a prescribed rule, i.e.,

$$\psi_i[t+1] = g_i(\psi_j^i[t], t), \quad j \in \mathcal{N}_i[t], t \in \mathbb{Z}_{\geq 0}, \quad (5)$$

where $\psi_j^i[t]$ is the state of node j sent to node i at time step t and $\psi_i^i[t] = \psi_i[t]$. The update rule, g_i, which is designed *a priori*, can be an arbitrary function and may be different for each node. For example, for the standard linear consensus protocol [58], this function is simply some linear combination of the values of node i's neighbors:

$$\psi_i[t+1] = \sum_{j \in \mathcal{N}_i[t]} w_{ij}[t]\psi_j^i[t], \quad (6)$$

where $w_{ij}[t]$ is the weight assigned to node j's value by node i at time step t.

Recall that node i is a *regular* node if it does not deviate from its prescribed update rule $g_i(\cdot)$. The set of regular nodes is denoted by $\mathcal{R}$. A deviation can stem from a failure, e.g., disturbance or noise with a known model, a time delay or signal dropout, or an adversarial action (attack) in the form of arbitrary state updates. Some fundamental differences between faults and attacks were discussed in Section 1.1. Consider a set $\mathcal{F} = \{i_1, i_2, ..., i_f\} \subset \mathcal{V}$ of malicious nodes. One way to represent an adversarial action at time step t is to use an additive attack signal $\zeta_i[t]$ in the updating rule (6). In particular, instead of applying the update equation (6), each node $i \in \mathcal{F}$ updates its state as

$$\psi_i[t+1] = \sum_{j \in \mathcal{N}_i[t]} w_{ij}[t]\psi_j^i[t] + \zeta_i[t]. \quad (7)$$

Here, an agent is malicious in T time steps if $\zeta_i[t] \neq 0$ for at least one time step $0 \leq t \leq T - 1$. Writing (7) in vector form yields

$$\psi[t+1] = \boldsymbol{W}\psi[t] + \underbrace{\begin{bmatrix} \mathbf{e}_{i_1} & \mathbf{e}_{i_2} & ... & \mathbf{e}_{i_f} \end{bmatrix}}_{\boldsymbol{B}_{\mathcal{F}}} \boldsymbol{\zeta}_{\mathcal{F}}[t],$$
$$\boldsymbol{y}[t] = \boldsymbol{C}\psi[t] \quad (8)$$

where $\boldsymbol{\zeta}_{\mathcal{F}}[t] = [\zeta_{i_1}, \zeta_{i_2}, ..., \zeta_{i_f}]$ is the additive error (attack) vector. Matrix $\boldsymbol{W}$ represents the communications between agents and $\boldsymbol{C} = [\boldsymbol{C}_1' \;\; \boldsymbol{C}_2' \;\; ... \;\; \boldsymbol{C}_q']'$, where $\boldsymbol{C}_i$ determines the measurement of node i.

Observability and Connectivity: The objective for each node i is to recover the vector $\psi[0] = [\psi_1[0] \; \psi_2[0], ..., \; \psi_n[0]]'$. To build intuition for the adversarial setting, let us start by looking at the case

---

[5] This value is called the state or the opinion of node i, depending on the context.

without malicious nodes. Moreover, let us consider a time-invariant graph. In this case, based on the linear iterative scheme in Eq. (6), the task of estimating $\psi[0]$ at node i can be equivalently viewed as the problem of recovering the initial condition of the linear discrete-time dynamical system

$$\psi[t+1] = \boldsymbol{W}\psi[t],$$

based on the following observation model:

$$y_i[t] = \boldsymbol{C}_i\psi[t]. \quad (9)$$

Here, $\boldsymbol{C}_i$ is a $(d_i + 1) \times n$ matrix with a single 1 in each row that denotes the states available to node i (these positions correspond to the neighbors of node i, including node i). From basic linear systems theory, it then follows that $\psi[0]$ can be recovered by node i if the pair $(\boldsymbol{W}, \boldsymbol{C}_i)$ is observable. In [87], it is shown that if the underlying network is connected, then the weight matrix $\boldsymbol{W}$ can be designed in a way that ensures observability of $(\boldsymbol{W}, \boldsymbol{C}_i)$.

The above discussion reveals how concepts from systems theory such as observability can be combined with basic graph-theoretic notions such as connectivity to study the process of information diffusion over networks. It is natural to thus wonder whether a marriage of ideas between systems theory and graph theory will continue to be fruitful while analyzing the adversarial setting. The results from [64], summarized below, establish that this is indeed the case.

Suppose that a subset $\mathcal{F}$ of nodes is malicious and deviates from the update rule (6). Thus, the new goal is to recover $\psi[0]$ for the model in (8). As in the non-adversarial case, we start by examining the observation model at node i. We note that the set of all values seen by node i during the first $L + 1$ time-steps of the linear iteration (for any non-negative integer L) is given by

$$y_i[0 : L] = \mathcal{O}_{i,L}\psi[0] + \mathcal{M}_{i,L}^{\mathcal{F}}\boldsymbol{\zeta}_{\mathcal{F}}[0 : L - 1], \quad (10)$$

where $y_i[0 : L] = \begin{bmatrix} y_i'[0] & y_i'[1] & \cdots & y_i'[L] \end{bmatrix}'$ and $\boldsymbol{\zeta}_{\mathcal{F}}[0 : L - 1] = \begin{bmatrix} \zeta_{\mathcal{F}}'[0] & \zeta_{\mathcal{F}}'[1] & \cdots & \zeta_{\mathcal{F}}'[L - 1] \end{bmatrix}'$. Matrices $\mathcal{O}_{i,L}$ and $\mathcal{M}_{i,L}^{\mathcal{F}}$ are the observability and invertibility matrices, respectively (from the perspective of node i), and can be expressed recursively as

$$\mathcal{O}_{i,L} = \begin{bmatrix} \boldsymbol{C}_i \\ \mathcal{O}_{i,L-1}\boldsymbol{W} \end{bmatrix}, \quad \mathcal{M}_{i,L}^{\mathcal{F}} = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathcal{O}_{i,L-1}\boldsymbol{B}_{\mathcal{F}} & \mathcal{M}_{i,L-1}^{\mathcal{F}} \end{bmatrix}, \quad (11)$$

where $\mathcal{O}_{i,0} = \boldsymbol{C}_i$ and $\mathcal{M}_{i,0}^{\mathcal{F}}$ is the empty matrix (with zero columns). The question of interest is the following: *Under what conditions can node i recover $\psi[0]$ based on a sufficiently large sequence of observations, despite the presence of the unknown inputs $\boldsymbol{\zeta}_{\mathcal{F}}$?*

As it turns out, the answer to the above question is intimately tied to the system-theoretic concept of *strong observability*. In particular, the linear system (8) is said to be strongly observable w.r.t. node i if $y_i[t] = 0$ for all t implies $\boldsymbol{\psi}[0] = 0$ (regardless of the values of the unknown inputs $\boldsymbol{\zeta}_{\mathcal{F}}[t]$). Moreover, if such a strong observability condition holds, then this is equivalent to saying that node i will be able to uniquely determine the initial condition $\boldsymbol{\psi}[0]$ based on the knowledge of its output sequence, regardless of the unknown inputs.

Strong Observability and Connectivity: For the non-adversarial setting, connectivity was enough to ensure observability of the pair $(\boldsymbol{W}, \boldsymbol{C}_i), \forall i \in \mathcal{V}$. In a similar vein, we need to now discern how the structure of the underlying network impacts strong observability. To this end, we present a simple argument to demonstrate that if the network is not adequately connected, then system (8) will not be strongly observable w.r.t. certain nodes in the graph. For simplicity, let $\mathcal{G}$ be undirected. Now suppose the connectivity κ of $\mathcal{G}$ is such that $\kappa \leq |\mathcal{F}|$. This implies the existence of a vertex cut $\mathcal{S}_2$ of size at most $|\mathcal{F}|$ that separates the graph into two disjoint parts. Let the vertex sets for these disjoint parts be denoted by $\mathcal{S}_1$ and $\mathcal{S}_3$. After reordering the nodes such that the nodes in $\mathcal{S}_1$ come first, followed by those in $\mathcal{S}_2$ and then $\mathcal{S}_3$, the consensus weight matrix takes the following form:

$$
\boldsymbol{W} = \begin{bmatrix} \boldsymbol{W}_{11} & \boldsymbol{W}_{12} & \boldsymbol{0} \\ \boldsymbol{W}_{21} & \boldsymbol{W}_{22} & \boldsymbol{W}_{23} \\ \boldsymbol{0} & \boldsymbol{W}_{32} & \boldsymbol{W}_{33} \end{bmatrix} .
$$

The structure of the above matrix follows immediately from the fact that the nodes in $\mathcal{S}_1$ can interact with those in $\mathcal{S}_3$ only via the nodes in $\mathcal{S}_2$. Now suppose all the nodes in $\mathcal{S}_2$ are adversarial; this is indeed feasible since $|\mathcal{S}_2| \leq |\mathcal{F}|$. Moreover, let the initial condition $\boldsymbol{\psi}[0]$ be of the form

$$
\boldsymbol{\psi}'[0] = \begin{bmatrix} \boldsymbol{\psi}_{\mathcal{S}_1}[0] & \boldsymbol{\psi}_{\mathcal{S}_2}[0] & \boldsymbol{\psi}_{\mathcal{S}_3}[0] \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{v} \end{bmatrix},
$$

where $\boldsymbol{v}$ is a non-zero vector in $\mathbb{R}^{|\mathcal{S}_3|}$. If the adversarial inputs are of the form $\boldsymbol{\zeta}_{\mathcal{F}}[t] = -\boldsymbol{W}_{23}\boldsymbol{\psi}_{\mathcal{S}_3}[t]$, then it is easy to see that the states of the nodes in $\mathcal{S}_1$ and $\mathcal{S}_2$ remain at zero. Thus, an agent i in $\mathcal{S}_1$ observes a sequence of zeroes. It follows that there is no way for agent i to distinguish the zero initial condition from the non-zero initial condition we considered in this example. Thus, system (8) is not strongly observable w.r.t. node i.

Discussion: The above argument serves to once again highlight the interplay between control- and graph-theory in the context of information diffusion over networks. Moreover, it suggests that in order for every node to uniquely determine the initial condition, the connectivity of the network has to somehow scale with the number of adversaries. Using a more refined argument than the one we presented above, it is possible to show that a connectivity of $2f + 1$ is necessary for the problem under consideration [64,66], where f is the maximum number of malicious nodes in the network.

As shown in [64], a connectivity of $2f + 1$ is also sufficient for the linear iterative strategy to reliably disseminate information between regular nodes in the network despite the actions of up to f malicious adversaries.

Theorem 3 ([64]) *Given a fixed network with n nodes described by a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, let f denote the maximum number of malicious nodes that are to be tolerated in the network, and let κ_{ij} denote the size of the smallest (i, j)-cut between any two vertices j and i. Then, regardless of the actions of the malicious nodes, node i can uniquely determine all of the initial values in the network via a linear iterative strategy if and only if $\min_j \kappa_{ij} \geq 2f + 1$. Furthermore, if this condition is satisfied, i will be able to recover the initial values after the nodes run the linear iterative strategy with almost any choice of weights for at most n time-steps.*

A key ingredient in the proof of Theorem 3 is establishing that if $\mathcal{G}$ is $(2f + 1)$-connected, then the tuple $(\boldsymbol{W}, \boldsymbol{B}_{\mathcal{F}}, \boldsymbol{C}_i)$ is strongly observable (i.e., does not possess any zero dynamics) $\forall i \in \mathcal{V}$, for almost all choices of the weight matrix $\boldsymbol{W}$, and under any f-total adversarial set $\mathcal{F}$. [6] This interdependence between system and graph-theoretic properties is schematically illustrated in Fig. 5(c).

One important point to note is that the approaches developed in [64] and [66] to combat adversaries require each regular agent to possess complete knowledge of the network structure, and to perform a large amount of computation to identify the malicious sets. For large-scale networks, this may be infeasible. One may thus ask whether it is possible to resiliently diffuse information across a network when each regular agent only has *local* knowledge of its own neighborhood, and can run only simple computations. In the subsequent sections, we provide an overview of work showing that this is indeed possible. However, as we shall see, the lack of global information will dictate the need for stronger requirements on the network topology (relative to $(2f+1)$-connectivity).

Remark 4 (Byzantine Attacks) *For point-to-point communications, which are prone to Byzantine attacks, in addition to the network connectivity, which has to be at least $2f + 1$, the total number of nodes must satisfy $n \geq 3f + 1$. This is because of the fact that if i receives*

---

[6] Note that when $\mathcal{F} = \emptyset$ (i.e., in the absence of adversaries), we immediately recover that connectivity of $\mathcal{G}$ implies observability of $(\boldsymbol{W}, \boldsymbol{C}_i), \forall i \in \mathcal{V}$, for almost all choices of the weight matrix $\boldsymbol{W}$.
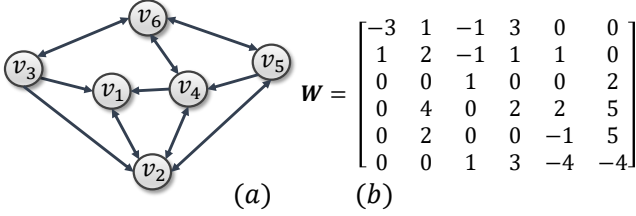
Fig. 6. (a) Graph discussed in Example, (b) The weight matrix corresponding to graph (a).

j's value reliably, it still does not know what j told other nodes in the network. Thus, there must be a sufficient number of non-Byzantine nodes in the network in order for i to ascertain what j told 'most' of the nodes [88].

Example 4 *Consider the graph shown in Fig. 6(a). The objective is for node 1 to calculate the function $\sum_{i=1}^{6} \psi_i^2[0]$ despite the presence of a malicious node in the network, i.e, $f = 1$. In this graph, nodes v_2, v_3, and v_4 are neighbors of v_1, and v_5 and v_6 have three internally vertex-disjoint paths to v_1. Thus, $\kappa_{j1} \geq 3$ for all j and, based on Theorem 3, v_1 is able to calculate the desired function after running the linear iteration (with almost any choice of weights) for at most $n = 6$ time-steps. We take each of the edge and self loop weights to be i.i.d. random variables from the set $\{-5, -4, -3, -2, -1, 1, 2, 3, 4, 5\}$ with equal probabilities. These weights produce the matrix shown in Fig. 6 (b). Since node 1 has access to its own state and the states of its neighbors, we have $C_1 = [\mathbf{I}_4 \quad \mathbf{0}]$. Based on these values, matrices $\mathcal{O}_{1,L}$ and $\mathcal{M}_{1,L}^{\mathcal{F}}$ are obtained and v_1 can calculate the initial state vector $\psi[0]$ using (10). Suppose that the initial values of the nodes are $\psi[0] = [3 \ -1 \ 4 \ -4 \ 7 \ 11]'$ and v_4 is a malicious node. At time steps 1 and 2, v_4 adds an additive error of $\zeta_4[0] = -8$ and $\zeta_4[0] = -12$ to its updating rule. The values of all nodes over the first three time-steps of the linear iteration are given by $\psi[0] = [3 \ -1 \ 4 \ -4 \ 7 \ 11]'$, $\psi[1] = [-26 \ 0 \ 26 \ 49 \ 46 \ -80]'$, and $\psi[2] = [199 \ 43 \ -134 \ -222 \ -446 \ 309]'$. The values seen by v_1 at time-step t are given by $y_1[t] = C_1\psi[t]$; node v_1 can now use $y_1[t]$ to calculate the vector of initial values, despite the efforts of the malicious node. Node v_1 has to find a set $\mathcal{F}_i$ for which $y_1[0:2]$ falls into the column space of $\mathcal{O}_{1,2}$ and $\mathcal{M}_{1,2}^{\mathcal{F}}$. In this example, v_1 can figure out that this holds for $j = 4$. Then, it finds vectors $\bar{\psi}$ and $\zeta_{\mathcal{F}_4}[0:1]$ such that $y_i[0:2] = \mathcal{O}_{1,2}\bar{\psi} + \mathcal{M}_{1,2}^{\mathcal{F}_4}\zeta_{\mathcal{F}_4}[0:1]$ as $\zeta = [3 \ -1 \ 4 \ -4 \ 7 \ 11]'$ and $\zeta_{\mathcal{F}_4}[0:1] = [-8 \ -12]'$. Node v_1 now has access to $\psi[0] = \bar{\psi}$ and can calculate $\sum_{i=1}^{6} \psi_i^2[0] = 212$.*

It is worth noting that for the network in Fig. 6(a), we have $\kappa_{26} = 2$, since the set $\mathcal{F} = \{v_4, v_5\}$ forms a $(2, 6)$-cut (i.e., removing nodes v_4 and v_5 removes all paths from v_2 to v_6). Thus, node v_6 is not guaranteed to be able to calculate any function of node v_2's value when there is a faulty node in the system. In particular, one can verify that in the example above, where node v_4 is malicious

and updates its values with the errors $\zeta_4[0] = -8$ and $\zeta_4[0] = -12$, the values seen by v_6 during the first three time steps of the linear iteration are the same as the values seen by node v_6 when $\psi_1[0] = 4$, $\psi_1[0] = -3$ and node 5 is malicious, with $\psi_5[0] = 4$ and $\psi_5[0] = 6$. In other words node v_6 cannot distinguish the case when node v_4 is faulty from the case where node v_5 is faulty (with different initial values in the network).

5 Resilient Distributed Consensus

Distributed consensus is a well studied application of information diffusion in networks. In distributed consensus, every node in the network has some information to share with the others, and the entire network must come to an agreement on an appropriate function of that information [58, 61, 89, 90]. In the resilient version of distributed consensus, the algorithm has to be modified in such a way that it maintains the consensus value in a desired region despite the actions of adversarial agents who attempt to steer the states outside that region (or disrupt agreement entirely). The desired steady state value can vary according to the application of interest [91, 92].

Remark 5 (A Fundamental Limitation) *In the standard linear consensus dynamics, a single malicious agent (shown in black in Fig. 7) can drive the consensus value towards its own state simply by keeping its value constant,[7] as shown in Fig. 7. More generally, since the initial values of the nodes are assumed to be known only to the nodes themselves, an adversarial node can simply change its own initial value and participate in the rest of the algorithm as normal. This would allow the adversary to affect the final consensus value (through its modified local value), but never be detected. Thus, perfect calculation of any function of initial values is generally impossible under adversarial behaviour. This is a fundamental limitation of all distributed algorithms for any problem where each agent holds data that is required by others to compute their functions (e.g., as in consensus, function calculation, or distributed optimization), as will be discussed later.*

5.1 Classical Approaches

Classical results on distributed consensus in the presence of Byzantine agents date back to the computer science literature [95] showing that the regular agents can always reach a consensus if and only if (1) the number of Byzantine agents is less than $\frac{1}{2}$ of the network connectivity, and (2) less than $\frac{1}{3}$ of the total number of agents. Similar results have been derived to detect and identify malicious nodes in consensus dynamics, as discussed in the previous section [64, 66]. However, these

---

[7] These types of adversarial agents are referred to as *stubborn* agents in the literature [92–94].
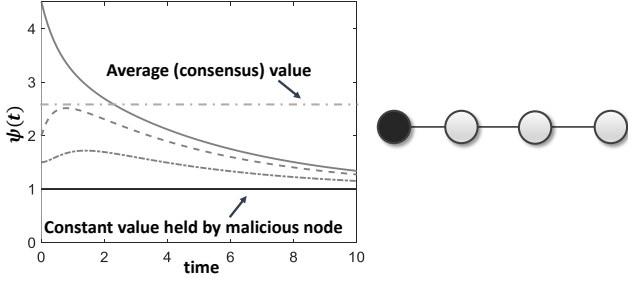
Fig. 7. State trajectories of agents in the presence of a stubborn agent.

works require each regular node to have full knowledge of the network topology, and for each regular node to perform complex computations. The following subsection describes an alternative scalable and "purely local" method for resilient consensus.

5.2   Purely Local Approaches

By imposing stronger conditions on the network topology (beyond being just $(2f + 1)$-connected), one can formulate algorithms that can handle worst case f-local Byzantine attacks with much less computational cost. In this class of algorithms, which were first named *approximate agreement* [95], each node disregards the largest and smallest f values received from its neighbors at each iteration and updates its state to be the average of a carefully chosen subset of the remaining values (such quantities are known as *trimmed means* in the robust statistics literature [96]). These methods were extended to a class of algorithms named *Mean-Subsequence-Reduced (MSR)* algorithms [97]. In [98] a continuous-time variation of the MSR algorithms, named the Adversarial Robust Consensus Protocol (ARC-P) was proposed. In what follows, we discuss an extension of MSR algorithms, called Weighted-Mean-Subsequence Reduced (W-MSR) in [65], which can handle f-local adversarial agents.

The algorithm is as follows.

1) Let $\psi_i[t] \in \mathbb{R}$ be the value maintained at each time step by each regular node i. At each time-step t, each regular node i receives values $\psi_i[t]$, $j \in \mathcal{N}_i$ from all of its neighbors, and ranks them from largest to smallest.
2) If there are f or more values larger than $\psi_i[t]$, normal node i removes the f largest values. If there are fewer than f values larger than $\psi_i[t]$, regular node i removes all of these larger values. This same logic is applied to the smallest values in regular node i's neighborhood. Let $\mathcal{R}_i[t]$ denote the set of nodes whose values were removed by i at time-step t.
3) Each regular node i updates its value as

$$\psi_i[t+1] = w_{ii}[t]\psi_i[t] + \sum_{j \in \mathcal{N}_i[t] \setminus \mathcal{R}_i[t]} w_{ij}[t]\psi_j[t], \quad (12)$$

where $w_{ii}[t]$ and $w_{ij}[t]$ satisfy the following conditions:

- $\sum_{j=1}^n w_{ij} = 1, \quad \forall i \in \mathcal{V}, t \in \mathbb{Z}_{\geq 0}$;
- $w_{ij}[t] = 0$ whenever $j \notin \mathcal{N}_i[t], t \in \mathbb{Z}_{\geq 0}$;
- $w_{ij}[t] \geq \alpha$ for some $\alpha > 0$ whenever $j \in \mathcal{N}_i[t], t \in \mathbb{Z}_{\geq 0}$.

We call the largest number of values that each node could throw away the *parameter* of the algorithm (it is equal to $2f$ in the above algorithm). Before presenting conditions for resilient consensus, in the following example we show that large network connectivity is no longer sufficient to guarantee consensus under W-MSR algorithms.

Example 5 *In the graph shown in Fig. 3 (a), suppose that the initial value of nodes in set $\mathcal{V}_1$ and set $\mathcal{V}_2$ are zero and 1, respectively. For $f = 1$, if each node disregards the largest and smallest values in its neighborhood, then the value of nodes in sets $\mathcal{V}_1$ and $\mathcal{V}_2$ remains the same as their initial value for all $t \geq 0$. As a result, consensus will not be achieved even though there are no malicious nodes. This lack of consensus is despite the fact that the connectivity of the graph is $\frac{n}{2}$, and arises due to the fact that the local state-dependent filtering in W-MSR causes all nodes between sets $\mathcal{V}_1$ and $\mathcal{V}_2$ to be disconnected at each iteration.*

Although the network connectivity is no longer an appropriate metric for analyzing the resilience of W-MSR dynamics, the notion of graph robustness from [65, 78] (see Section 2.1.2) turns out to be the key concept. We first start with the following concept. Denote the maximum and minimum values of the normal nodes at time-step t as $M[t]$ and $m[t]$.

Definition 4 (f-local safe) *Under the f-local adversarial model, the W-MSR algorithm is said to be f-local safe if both of the following conditions are satisfied: (i) all regular nodes reach consensus for any choice of initial values, and (ii) the regular nodes' values (including the final consensus value) are always in the range $\big[m[0], M[0]\big]$.*

The following result provides conditions under which the W-MSR algorithm guarantees (or fails) to be f-local safe.

Theorem 6 ([65]) *Under the f-local Byzantine adversary model, the W-MSR algorithm with parameter $2f$ is f-local safe if the network $\mathcal{G}$ is $(2f + 1)$-robust. Furthermore, for any $f > 0$, there exists a $2f$-robust network which fails to reach consensus based on the W-MSR algorithm with parameter $2f$.*

As discussed in Section 2.1.2, the robustness condition used in Theorem 6 is much stronger than the network connectivity condition which was required in classical distributed consensus algorithms. However, this stronger condition can be considered as the price to be paid for a computationally tractable resilient consensus algorithm
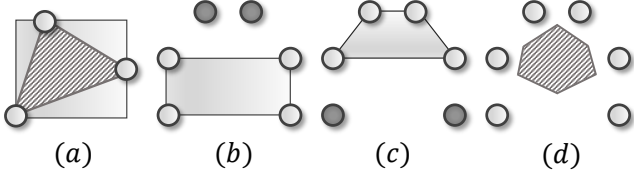
12

Fig. 8. (a) The box over-estimating the convex-hull of nodes in $\mathbb{R}^2$. (b),(c) show non-intersecting convex hull of the regular nodes. (d) the dark region is the centerpoint region of all six nodes.

which is able to tolerate worst-case Byzantine attacks. Furthermore, under this condition, one gains the ability to tolerate f-local adversaries (rather than f-total adversaries) "for free".

Several other variations of the above approach, including extensions to second order consensus with asynchronous time delay and applications to formation control of mobile robots, are discussed in [99–101]. Applications of resilient consensus in multi-robot systems using Wi-Fi communication is studied in [102]. Resilient flocking in multi-robot systems requires the extension of the above techniques to time varying networks, which is studied in [103]. There, it is shown that if the required network robustness condition is not satisfied at all times, the network can still reach resilient consensus if the union of communication graphs over a bounded period of time satisfies $(2f + 1)$-robustness. Moreover, a control policy to attain such resilient behavior in the context of perimeter surveillance with a team of robots was proposed.

5.3 Resilient Vector Consensus

The W-MSR algorithm described in the previous section considered the case where agents maintain and exchange scalar quantities and remove "extreme" values at each iteration. However, the extension to multi-dimensional vectors requires further considerations since there may not be a total ordering among vectors. One option is to simply run W-MSR on each component of the vector separately. If the graph is $(2f + 1)$-robust, this would guarantee that all regular agents reach a consensus in a hypercube formed by the initial vectors of the regular agents (despite the action of any f local set of Byzantine agents). Keeping the agents' states within the convex hull of the initial vectors, however, requires some subtleties, as we will discuss in this section. We know that the convex hull of a set of vectors is a subset of the region (the box) formed by the convex hull of their components separately. This is shown in Fig. 8 (a) in which the triangle is the convex hull of the three points in $\mathbb{R}^2$ and the grey rectangle is the box formed by calculating the convex hull of each component of the three vectors separately. Thus, the component-wise convex hull gives an overestimate of the actual convex hull of the vectors.

First attempts to address Byzantine resilient vector consensus to the convex hull of the initial values of the reg-

ular agents were provided in [104, 105], and were further developed in the context of rendezvous multi robot systems by [106]. To describe these approaches, we first need to extend the notion of f-local safe points.

Definition 5 ([106]) *Given a set of n nodes in $\mathbb{R}^d$ of which at most f are adversarial, a point p that is guaranteed to lie in the interior of the convex hull of $(n - f)$ regular points is called an f-safe point.*

Based on the above definition, the resilient vector consensus algorithm relies on the computation of f-safe points by each node. In particular,

- Let $\boldsymbol{\psi}_i[t] \in \mathbb{R}^d$ be the value maintained at each time step by each regular node i. In the iteration t, a regular node i gathers the state values of its neighbors $\mathcal{N}_i[t]$.
- Each regular node i computes an f_i-safe point, denoted by $s_i[t]$, of points corresponding to its neighbors' states.
- Each regular agent i then updates its state by moving toward the safe point $s_i[t]$, i.e.,

$$\boldsymbol{\psi}_i[t + 1] = \alpha_i[t]s_i[t] + (1 - \alpha_i[t])\,\boldsymbol{\psi}_i[t], \qquad (13)$$

where $\alpha_i[t] \in (0, 1)$ is a dynamically chosen parameter whose value depends on the application.

It was shown in [106] that if all regular agents follow the above routine, they are guaranteed to converge to some point in the convex hull of their initial states. The following proposition shows conditions of the existence of f-safe points.

Proposition 7 ([106]) *Given n points in $\mathbb{R}^d$, where $d \in \{1, 2, ..., 8\}$, and at most f points belong to adversaries, then there exists an f-safe point if $n \geq (f + 1)(d + 1)$. It also holds for $d > 8$ if Reay's conjecture is true [107].*

In [108], it is shown that $n \geq (f + 1)(d + 1)$ is also a necessary condition for the existence of an f-safe point. An example is shown in Fig. 8 (b) and (c). The malicious nodes are shown with darker colors. Here, for $f = 2$ malicious nodes, there is no 2-safe point (an interior point in the intersection of convex hull of four regular nodes).

The main question is how to find these f-safe points from a given set of points in $\mathbb{R}^d$. In [104, 106], this is done via the *Tverberg partitioning algorithm*, which partitions points into subsets such that the convex hull of the partitions have a non-empty intersection, provided that the number of nodes is sufficiently large; see [109] for more details. A similar approach is recently adopted in [110]. However, finding Tverberg partitions is computationally hard in practice (although it is not proved that the problem is NP hard). To achieve fast algorithms, one has to pay the price of reducing the number of parts in

the partition, i.e., the number of malicious nodes. Linear time approximation algorithms to find Tverberg points, i.e., f-safe nodes, have been proposed in [111], provided that $f \leq \lceil \frac{n}{2^d} \rceil - 1$. On the other hand, Tverberg partitioning provides strong conditions for f-safe points, i.e., the outcome of the Tverberg partitioning algorithm are f-safe points, but the reverse is not true. A less conservative approach to find f-safe points via the notion of a centerpoint was developed in [112], as explained below.

Definition 6 *Given a set X of n points in $\mathbb{R}^d$, a centerpoint p is a point, not necessarily from X, such that any closed half-space[8] of $\mathbb{R}^d$ containing p also contains at least $\frac{n}{d+1}$ points from X.*

By the Centerpoint Theorem, every finite set of points in $\mathbb{R}^d$ has a centerpoint [113]. It is shown that for a set of n points in $\mathbb{R}^d$ and $f \leq \lceil \frac{n}{d+1} \rceil - 1$, the region of f-safe points is *equivalent to* the centerpoint region [112]. The centerpoint of six points in shown in Fig. 8 (d). The following result provides conditions for finding f-safe points.

Proposition 8 ([112]) *Given n points in $\mathbb{R}^d$ for which at most f points belong to adversaries, then an f-safe point can be computed (using centerpoint) if*

$$
\begin{aligned}
f &\leq \lceil \frac{n}{d+1} \rceil - 1 \quad d = 2, 3, \\
f &\leq \lceil \frac{n}{d^2} \rceil - 1 \qquad d > 3.
\end{aligned} \quad (14)
$$

6 Resilient Distributed Optimization

While the consensus problem discussed in the previous section considered the scenario where each agent has a static initial value, a more general setting is that of distributed optimization. In this setting, each agent $i \in \mathcal{V}$ has a convex function $g_i : \mathbb{R} \to \mathbb{R}$ (with bounded subgradients) which is only available to agent i.[9] The objective is for the agents to solve the following global optimization problem in a distributed manner:

$$
\min_{\psi} \; g(\psi) = \frac{1}{n} \sum_{i=1}^n g_i(\psi). \quad (15)
$$

A common approach to solve this problem is to use a synchronous iterative consensus-based protocol in which agents use a combination of consensus dynamics and gradient flow to find a minimizer of $g(\psi)$ [114]. More specifically, at every time-step $t \in \mathbb{N}$, each agent i maintains an

---

[8] A closed half-space in $\mathbb{R}^d$ is a set of the form $\{x \in \mathbb{R}^d : a'x \geq b\}$ for some $a \in \mathbb{R}^d \setminus \{0\}$.

[9] We discuss the case of multi-dimensional functions later in this section.

estimate $\psi_i(t) \in \mathbb{R}$ of the solution to (15), and updates it based on the information received from its neighbors, as follows

$$
\psi_i[t+1] = w_{ii}[t]\psi_i[t] + \sum_{j \in \mathcal{N}_i[t]} w_{ij}[t]\psi_j[t] - \alpha_t d_i[t]. \quad (16)
$$

In the above update rule, $d_i[t]$ is a subgradient of g_i evaluated at $w_{ii}[t]\psi_i[t] + \sum_{j \in \mathcal{N}_i[t]} w_{ij}[t]\psi_j[t]$, and α_t is the step size sequence corresponding to the influence of the subgradient on the update rule at each time-step. Dynamics (16) can be represented in the following vector form as

$$
\boldsymbol{\psi}[t+1] = \boldsymbol{W}\boldsymbol{\psi}[t] - \alpha_t \boldsymbol{d}[t], \quad (17)
$$

where $\boldsymbol{W}$ is a doubly-stochastic matrix. The following result shows that the update rule (16) allows the nodes in the network to distributively solve the global optimization problem (15).

Proposition 9 ([114]) *Suppose that $\mathcal{G}$ is strongly connected at each time step and that the subgradients of each of the local functions g_i are bounded. For the update rule (16) with step sizes satisfying $\sum_{t \in \mathbb{N}} \alpha_t = \infty$ and $\sum_{t \in \mathbb{N}} \alpha_t^2 < \infty$, we have*

$$
\lim_{t \to \infty} \|\psi_i[t] - \psi^*\| = 0, \quad (18)
$$

for all $i \in \mathcal{V}$, where ψ^* is the global minimizer of $g(\psi)$.

Our objective here is to summarize the vulnerabilities of such protocols to adversarial agents, and to provide an overview of secure distributed optimization algorithms that have provable safety guarantees despite the presence of such agents. Recent works have focused on some gradient-based metrics to detect and identify malicious agents in a distributed optimization algorithm [115]. However, similar to previous sections, our focus is on graph-theoretic methods. As before, we assume that adversarial nodes can update their states in a completely arbitrary manner.

Remark 10 (Fundamental Limitation Revisited) *Following the fundamental limitation discussed in Remark 5, one can easily argue that it is generally impossible to compute ψ^* when there are adversarial agents in the network, since one can never infer their local functions accurately. As an example, suppose that node n is adversarial and wishes the states to converge to $\bar{\psi} \in \mathbb{R}$. It simply chooses a function $\bar{g}(\psi)$ such that the minimizer of $\sum_{i=1}^{n-1} g_i(\psi) + \bar{g}_n(\psi)$ is $\bar{\psi}$. For a vanishing step size, i.e., $\lim_{t \to \infty} \alpha_t = 0$, all regular nodes will asymptotically converge to $\bar{\psi}$ when following the distributed optimization dynamics (16). Since the functions g_i are arbitrary and known only to the nodes themselves, such deceptions cannot be detected.*

The above fact is formally stated as follows.

Theorem 11 *[116, 117] Suppose Γ is a distributed algorithm that guarantees that all nodes calculate the global optimizer of* (15). *Then a single adversary can cause all nodes to converge to any arbitrary value when they run algorithm Γ, and furthermore, will remain undetected.*

Theorem 11 indicates that it is impossible to develop an algorithm that always finds optimal solutions in the absence of adversaries and that is also resilient to carefully crafted attacks. Thus, the price that should be paid for resilient distributed optimization is a *loss in optimality*. In what follows, we describe resilient consensus-based distributed optimization protocols.

Suppose the adversarial nodes form an f-local set. At each time step, every regular node gathers and sorts the states of all of its neighbors and, similar to the W-MSR algorithm in Section 5, each agent disregards the highest f and lowest f states from the gathered states, denoted by set $\mathcal{R}_i[t]$, and updates its state as

$$\psi_i[t+1] = w_{ii}[t]\psi_i[t] + \sum_{j \in \mathcal{N}_i[t] \setminus \mathcal{R}_i[t]} w_{ij}[t]\psi_j[t] - \alpha_t d_i[t], \tag{19}$$

where $d_i[t], \alpha_t$ are the same as (16).

In [116], graph-theoretic conditions for agents to reach consensus in the presence of f-total malicious nodes under the linear filtering rule in (19) are discussed. The arguments rely on the fact that in $(2f + 1)$-robust networks, the weight matrix $\boldsymbol{W}$ corresponding to the regular nodes is rooted at each time-step, i.e., there is a node with a directed path to every other node in $\mathcal{G}$. Other sufficient conditions (similarly assuming the existence of rooted nodes in the set of regular nodes $\mathcal{R}$) is presented in [117].

The following result shows that (19) provides a safety guarantee for distributed optimization, i.e., convergence to the convex hull of local minimizers, under additional conditions on the step size α_t.

Theorem 12 ([116]) *Suppose that one of the following conditions holds:*

(i) *The adversarial nodes are f-total malicious and the network is $(f + 1, f + 1)$-robust; or*

(ii) *The adversarial nodes are f-local Byzantine and the network is $(2f + 1)$-robust.*

Furthermore, for each node $i \in \mathcal{R}$, let the local function $g_i(.)$ have minimizer m_i. Define $\bar{M} = \max\{m_i | i \in \mathcal{R}\}$ and $\underline{M} = \min\{m_i | i \in \mathcal{R}\}$. If stepsizes α_t satisfy $\sum \alpha_t = \infty$ and $\lim_{t\to\infty} \alpha_t = 0$, then $\limsup_{t\to\infty} \psi_i[t] \leq \bar{M}$ and $\liminf_{t\to\infty} \psi_i[t] \geq \underline{M}$ for all $i \in \mathcal{R}$, regardless of the actions of the adversarial nodes and the initial values.

Remark 13 (Lack of Convergence) *Despite the fact that the above resilient distributed optimization technique*

guarantees that the optimizer remains in the convex hull of local optimizers, it does not guarantee convergence to a constant value under certain type of adversarial actions and specific classes of step sizes. See examples in [116].

6.1 Factors that Affect the Performance of Resilient Distributed Optimization Algorithms

The following example from [116] shows that under the dynamics (19), the nature of the individual optimization functions together with the network topology determine how far away the convergence point is from the minimizer of the average of the regular nodes' functions.

Example 6 *Consider network $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ which is $2f + 1$ robust and let $\mathcal{T} \subset \mathcal{V}$ be an f-local set. Suppose all nodes are regular. Pick an $a \in \mathbb{R}$ and let nodes in $\mathcal{T}$ have local functions $g_a(\psi) = (\psi - a)^2$ and the nodes in $\mathcal{V} \setminus \mathcal{T}$ have local functions $g_b(\psi) = \psi^2$ (both functions can be modified to have their gradients capped at sufficiently large values, so as to not affect the minimizer of any convex combination of the functions). Let $g(\psi)$ be the average of all of the functions, with minimizer $\psi^* = \frac{|\mathcal{T}|}{n}a$. Then, under the local filtering dynamics with parameter f, all nodes converge to the value $\bar{\psi} = 0$ and thus $\bar{\psi} - \psi^* = \frac{|\mathcal{T}|}{n}a$ and $g(\bar{\psi}) - g(\psi^*) = \frac{|\mathcal{T}|^2}{n^2}a^2$.*

The above result shows that if the network contains a large f-local set (in relation to the total number of nodes) or the local functions have minimizers that are very different (corresponding to a large $|a|$ in the above result), then the value computed by (19) will have a greater divergence from the globally optimal solution. Note that an f-local set in a graph will have size at least equal to f (since any set of size f is f-local).

6.2 Extension to Multi-Dimensional Functions

In this subsection, we consider the case where the value for each agent is a d-dimensional vector. Thus, the local cost function for node i becomes $g_i : \mathbb{R}^d \to \mathbb{R}$, and the objective is to collaboratively solve (15) over $\psi \in \mathbb{R}^d$.

The extension to general multi-dimensional functions is a challenging problem as even the region containing the true minimizer of the functions is not easy to characterize. In particular, unlike the scalar case for which the minimizer of $g(\psi)$ lies within the convex hull of the minimizers of the individual functions, for the multi-dimensional case, the true minimizer vector may lie outside the convex hull of the individual minimizer vectors [118]. However, there have been some recent attempts to address the resilient distributed multi-dimensional optimization problem, as we now briefly summarize. Assuming the malicious agents behave in a prescribed manner, [115] proposed an attack detection and isolation technique before the execution of the

distributed optimization algorithm. In [119], the authors consider a resilient decentralized machine learning problem, and show that by utilizing a block coordinate descent method, the states of the regular agents will converge to the statistical minimizer with high probability. However, the analysis in [119] is restricted to i.i.d. training data across the network.

In contrast to [119], the authors in [120] develop a two-step filtering technique and provide convergence guarantees that do not make any statistical assumptions on the agents' objective functions. For an f-local Byzantine attack model, under the assumption that $\mathcal{G}$ is $((2d+1)f+1)$-robust, the approach in [120] guarantees asymptotic consensus of the states of all regular agents within a bounded region containing the global minimizer. [10] Simply applying the W-MSR algorithm to each coordinate of the parameter vector does not immediately lead to the above result. Instead, the approach in [120] relies on a carefully designed second filtering step.

Exact Fault Tolerance: Instead of settling for convergence to a *proximity* of the global minimizer (as in [116, 120]), one may ask whether it is possible to converge *exactly* to the minimizer of the sum of the objective functions of the regular agents, despite Byzantine attacks. As argued in [116], this is impossible unless additional assumptions are made on the agents' functions. In this context, the authors in [121] show that under a $2f$-redundancy assumption on the agents' objective functions, one can indeed achieve *exact* convergence even in the multi-dimensional case based on a norm filter. In fact, such an assumption turns out to be necessary for guaranteeing exact convergence, as established in [122].

The analysis in [121] is carried out for a complete peer-to-peer network under an F-total Byzantine attack model. Extending the results in [121] to general networks remains an open direction of research. Moreover, investigating whether the graph-theoretic conditions in [120] can be relaxed is also an interesting open problem.

7 Resilient Distributed Estimation and Inference

Another canonical distributed problem involves estimating/tracking an unknown state of interest based on measurements that are collected by a network of sensors. Within this broad setting, there can be several variations: the unknown state may be static or may evolve based on a dynamical model; the measurements may be noise-free or may be corrupted by stochastic noise; and the goal could be to estimate the state asymptotically, or to derive finite-time guarantees.

---

[10] Note that the requirement on the network topology scales with the dimension d of the parameter.

All of the above variations share a common unifying feature: *new information flows into the network at every time-step.* This is a key distinction between the distributed estimation setup, and the consensus and optimization problems that we discussed earlier. Another important difference stems from the fact that agents typically have *heterogeneous measurement/observation models* in an estimation problem. As a consequence, some agents may be more "informative" than others. This disparity in information content across the network is another key feature that is absent in the standard consensus or distributed optimization formulations. As such, the algorithmic techniques and graph-theoretic conditions that we will cover in this section will differ significantly from those in Sections 5 and 6; see Remark 1 in this context.

Before formally discussing resilient estimation algorithms, we outline two important considerations: the nature of the unknown quantity to be estimated, and the nature of the threat model.

Static Parameter Estimation vs. Dynamic State Estimation: As the name suggests, in static parameter estimation, the goal is to estimate a static parameter θ^* based on noisy sensor observations acquired by the agents. In contrast, the task in (dynamic) state estimation is to track a state $\mathbf{x}[t]$ that evolves based on a dynamical system model (such as a linear time-invariant (LTI) model). Even in the absence of adversaries, tracking the state of an *unstable* system based on dispersed measurements is a significantly challenging task. Thus, we will discuss the relatively simpler resilient distributed parameter estimation problem first, and then move on to the dynamic state estimation setting.

Sensor Attacks vs. Byzantine Attacks: The works that we will review can also be broadly classified in terms of the threat model. In particular, there are two predominant attack models that are studied in the resilient distributed estimation literature: sensor attacks and Byzantine attacks. In the former case, measurement streams of certain agents are corrupted by an additive attack signal; these signals may or may not be bounded. However, *all* agents behave normally, i.e., they follow the prescribed protocol at all times. In contrast, a Byzantine agent can act *arbitrarily*. Thus, the Byzantine attack model subsumes the sensor attack model. As we shall see, the nature of the threat model has significant implications for the graph-theoretic properties needed to combat attacks.

7.1 Parameter Estimation

In the distributed parameter estimation problem, each agent $i \in \mathcal{V}$ receives measurements as follows:

$$y_i[t] = \mathbf{H}_i \theta^* + w_i[t]. \tag{20}$$

Here, $\theta^* \in \mathbb{R}^d$ is the true unknown parameter, $y_i[t] \in \mathbb{R}^{r_i}$ is the measurement vector for agent i, $\mathbf{H}_i \in \mathbb{R}^{r_i \times d}$ is the local observation matrix for agent i, and $w_i[t]$ is the measurement noise that is typically assumed to independent and identically (i.i.d.) distributed over time, with zero mean and finite variance. Moreover, the noise sequences across different agents are assumed to be independent.

Objective: In the non-adversarial setting, the goal is to design a *consistent* distributed estimator, i.e., an estimator that ensures that the estimates of all agents converge to θ^* asymptotically almost surely. This is typically achieved by designing "consensus + innovations"-type estimators [123–126] that work under two standard assumptions: (i) the joint observation model is *globally observable*, i.e., $\sum_{i \in \mathcal{V}} \mathbf{H}_i' \mathbf{H}_i$ is invertible, and (ii) the graph $\mathcal{G}$ is *connected*.

In the resilient version of the above problem, a certain subset $\mathcal{A} \subseteq \mathcal{V}$ of the agents is corrupted either due to sensor attacks or due to Byzantine attacks. We now discuss the key algorithmic approaches to tackle such attacks. For each approach, we will focus on highlighting (i) the threat model; (ii) the main technique; (iii) the guarantees provided by the approach; and (iv) the assumptions on the observation model and the underlying graph needed to provide such guarantees.

1) Methods based on Adversary Detection: In [127], the authors consider a Byzantine attack model, and propose the *Flag Raising Distributed Estimation* ($\mathcal{FRDE}$) algorithm where agents simultaneously perform parameter estimation and adversary detection. Specifically, for parameter estimation, the regular agents employ a consensus+innovations update rule, similar to those in [123–126]. The consensus part of the update rule is based on a weighted average of neighbors' parameter estimates, while the innovation part processes the agent's own local measurements. For adversary detection, an agent computes the Euclidean distance between its own estimate and the estimates of its neighbors. If this distance exceeds a time-varying threshold, then an attack flag is raised. The design of this adaptive threshold constitutes the key part of the $\mathcal{FRDE}$ algorithm.

The $\mathcal{FRDE}$ algorithm is analyzed under two main assumptions: (i) the joint observation model of the *regular* agents is globally observable, i.e., $\sum_{i \in \mathcal{R}} \mathbf{H}_i' \mathbf{H}_i$ is invertible, where $\mathcal{R} = \mathcal{V} \setminus \mathcal{A}$; and (ii) the *induced* sub-graph $\mathcal{G}_\mathcal{R}$ of the regular agents is connected. Under these assumptions, it is shown in [127] that either all regular agents detect the presence of adversaries, or their local estimates converge to θ^* asymptotically almost surely. In other words, "strong" attacks get detected while "weak" attacks fail to disrupt the process of estimation.

Discussion: The assumption that the parameter is globally observable w.r.t. the joint measurements of the

regular agents is quite intuitive, and in fact necessary (under the Byzantine attack model). The necessity of the graph condition in [127] is, however, an open question. Based on the $\mathcal{FRDE}$ algorithm, if the presence of adversaries is detected, the system needs to go through an external "repair" phase; multiple such repair phases could potentially be quite expensive. An alternative is to thus design algorithms that *always* allow the regular agents to estimate the true parameter, despite the presence of adversaries. We now discuss such methods.

2) Saturating Adaptive Gain Methods: We will discuss this technique in some detail since it has recently been used in the context of resilient distributed state estimation as well [128]. To isolate the core idea, we will review the simplest version of this method introduced in [129] to tackle sensor attacks. For an agent $i \in \mathcal{A}$ under attack, its measurement model is as follows:

$$y_i[t] = \theta^* + a_i[t], \qquad (21)$$

where $a_i[t]$ is the attack signal injected in the measurements of agent i. For agents whose measurements have not been corrupted, the attack signal is identically zero at all times. Note immediately that the measurement model is *homogeneous* and *noise-free*, and that all uncompromised agents can directly measure θ^*. For this model, the authors in [129] propose the *Saturated Innovation Update* ($\mathcal{SIU}$) algorithm where all agents employ a consensus+innovations estimator with a time-varying gain applied to the local innovation term. Specifically, agent i's estimate $x_i(t)$ of θ^* is iteratively updated as

$$\begin{aligned} x_i[t+1] = x_i[t] - \beta_t \sum_{j \in \mathcal{N}_i} (x_i[t] - x_j[t]) \\ + \alpha_t K_i[t] (y_i[t] - x_i[t]), \end{aligned} \qquad (22)$$

where α_t, β_t, and $K_i[t]$ are strictly positive, scalar-valued design parameters. The time-varying gain $K_i[t]$ is defined as

$$K_i[t] = \begin{cases} 1, & \|y_i[t] - x_i[t]\|_2 \le \gamma_t \\ \frac{\gamma_t}{\|y_i[t] - x_i[t]\|_2} & \text{otherwise,} \end{cases} \qquad (23)$$

where γ_t is an adaptive threshold. For a detailed description of how the parameter sequences $\{\alpha_t\}$, $\{\beta_t\}$, and $\{\gamma_t\}$ are designed, we refer the reader to [129]. In what follows, we briefly explain why the design of γ_t is a delicate matter. Indeed, if γ_t is chosen to be too small, then the innovation gain $K_i[t]$ will limit the impact of adversaries; however, a very small innovation again may also prevent correct identification of θ^*. On the other hand, if γ_t is too large, then it may provide the adversaries with enough flexibility to direct the agents' estimates away from θ^*. Thus, striking the right balance in the design of γ_t is critical. The following theorem from [129] characterizes the performance of the $\mathcal{SIU}$ algorithm.

Theorem 1 ([129]) *Suppose the following conditions hold. (i) The graph $\mathcal{G}$ is connected. (ii) The true parameter θ^* is bounded, i.e., $\|\theta^*\|_2 \leq \eta$, for some finite η that is known a priori to all agents. (iii) Less than half of the agents are under sensor attack, i.e., $|\mathcal{A}|/|\mathcal{V}| < 1/2$. Then, the parameters $\alpha_t, \beta_t,$ and γ_t can be designed such that the update rule in Eq.* (22) *ensures*

$$\lim_{t \to \infty} (t+1)^{\tau_0} \|x_i[t] - \theta^*\|_2 = 0, \forall i \in \mathcal{V}, \quad (24)$$

for all τ_0 such that $0 \leq \tau_0 < \tau_1 - \tau_2$, where τ_1, τ_2 are design parameters satisfying $0 < \tau_2 < \tau_1 < 1$. [11]

The above theorem tells us that $\mathcal{SIU}$ is a consistent estimator, and that the rate of convergence is of the order of $1/t^{\tau_0}$, for any τ_0 satisfying $0 \leq \tau_0 < \tau_1 - \tau_2$. Building on the main idea of using an adaptive threshold to design the innovation gain, the authors in [129] later generalized their results to account for heterogeneous measurement models corrupted by noise; see [130] and [131]. In [132], it was shown that the saturating adaptive gain idea is also effective in the context of resilient distributed field estimation under measurement attacks.

Discussion: We now highlight two subtle implications of the choice of threat model. First, note that the guarantee in Theorem 1 holds for *all* agents, as opposed to just the regular agents. This is the typical guarantee one provides for measurement/sensor attack models. It is instructive to compare such a result with those for the Byzantine setting (see, for instance, Theorem 14) where the goal is to enable only the regular agents to estimate the unknown quantity of interest.

The second key observation pertains to the graph condition in Theorem 1. All that is needed is connectivity of the underlying network - the exact same condition even in the absence of adversaries. Thus, the main takeaway here is that *the graph-theoretic conditions for solving the distributed parameter estimation problem are the same with and without sensor attacks.* The main reason for this can be attributed to the fact that even if an agent's measurements are compromised, it does not try to actively disrupt the flow of information between regular agents; however, a Byzantine agent might. This necessitates much stronger graph-theoretic conditions to tackle Byzantine attacks, as we shall see in Section 7.2.

3) Methods based on Online Optimization: Yet another way to approach the distributed parameter estimation problem is to view it from the lens of online optimization. This is precisely the method adopted in [133], where the authors consider a Byzantine attack model. To explain this method, for each agent $i \in \mathcal{V}$, define its

---

[11] In [129], the adversarial set $\mathcal{A}$ is allowed to change over time.

local *asymptotic* loss function $g_i : \mathbb{R}^d \to \mathbb{R}$ as

$$g_i(x) = \frac{1}{2}\mathbb{E}\left[\|\mathbf{H}_i x - y_i\|_2^2\right], \quad (25)$$

where y_i is as in Eq. (20), and the expectation is taken w.r.t. the measurement noise $w_i[t]$. Since the distribution of the noise sequence is unknown to agent i, it cannot access the above loss function. Nonetheless, agent i can use all the measurements it has acquired up to each time-step t to compute an *empirical approximation* of $g_i(x)$:

$$g_{i,t}(x) = \frac{1}{2t}\sum_{s=1}^{t} \|\mathbf{H}_i x - y_i[s]\|_2^2. \quad (26)$$

The algorithm in [133] essentially combines local gradient descent on the above empirical loss functions, followed by coordinate-wise trimming to aggregate neighboring information; trimmed means are used to account for the presence of adversaries. This algorithm enables each regular agent to estimate the true parameter asymptotically almost surely. Moreover, as a departure from existing results on this problem, the authors provide *finite-time concentration bounds* that hold with high probability. To arrive at the above results, the conditions imposed on the graph topology are the same as those for Byzantine-resilient scalar consensus [134]. In [133], certain additional assumptions are made on the observation model that may not be necessary.

Additional Results: Before moving on to the state estimation setting, we briefly comment on a couple of related works. For a somewhat different observation model than in Eq. (20), the authors in [135] provide guarantees against Byzantine attacks by drawing on the techniques and graph-conditions in [65]. To the best of our knowledge, this is the earliest work on resilient distributed parameter estimation.

Very recently, the authors in [136] proposed a *min-switching* technique to account for the presence of Byzantine agents in the context of least-squares static estimation. The main idea behind the approach in [136] is to first construct an appropriate local Lyapunov function at each regular agent. The filtering technique then comprises of using only those neighboring estimates that lead to maximum decrease of the Lyapunov function. It is shown that this method can help relax the graph-theoretic conditions in both [133] and [137].

7.2 Dynamic State Estimation and Inference

In a typical distributed state estimation problem, the goal is to track the state of a linear time-invariant system of the following form:

$$\boldsymbol{x}[t+1] = \boldsymbol{A}\boldsymbol{x}[t], \quad (27)$$

where $\boldsymbol{x}[t] \in \mathbb{R}^N$ is the state vector and $\boldsymbol{A}$ is the state transition matrix. The system is monitored by a network $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ consisting of n nodes. The measurement model of the i-th node is given by

$$y_i[t] = \boldsymbol{C}_i \boldsymbol{x}[t], \qquad (28)$$

where $y_i[t] \in \mathbb{R}^{r_i}$ and $\boldsymbol{C}_i \in \mathbb{R}^{r_i \times N}$. We use $\boldsymbol{C} = [\boldsymbol{C}_1', \boldsymbol{C}_2', \cdots, \boldsymbol{C}_N']'$ to collect all the individual node observation matrices and $\boldsymbol{y}[t] = [y_1'[t], y_2'[t], \cdots, y_N'[t]]'$ to aggregate all the individual measurement vectors; accordingly, $\boldsymbol{y}[t] = \boldsymbol{C}\boldsymbol{x}[t]$. Each node i maintains an estimate $\hat{\boldsymbol{x}}_i[t]$ of the state $\boldsymbol{x}[t]$, and the goal is to ensure that these estimates converge to $\boldsymbol{x}[t]$ asymptotically.

Even in the absence of adversaries, the distributed state estimation problem is quite challenging, and only recently were necessary and sufficient conditions discovered for this problem [138–144]. The key technical challenge arises from the fact that $(\boldsymbol{A}, \boldsymbol{C}_i)$ may not be detectable w.r.t. the measurements of any individual node i. This is precisely what necessitates communication between nodes in the graph. The difficulty of tracking an unstable dynamical process based on dispersed measurements only gets exacerbated in the presence of adversaries. Following the same style of exposition as in Section 7.1, we now discuss the main techniques for resilient distributed state estimation.

1) Methods based on Observable Decompositions and Local Filtering: We start by reviewing the approach developed in [137] for solving the resilient distributed state estimation problem subject to an f-local Byzantine adversary model. In order to focus on the core ideas behind this approach, we assume that the system matrix $\boldsymbol{A}$ has real and simple eigenvalues; extensions to general spectra can be found in [137]. As a first step, we diagonalize $\boldsymbol{A}$ using a coordinate transformation matrix $\boldsymbol{V} = [\boldsymbol{v}^1, \boldsymbol{v}^2, ..., \boldsymbol{v}^N]$ formed by N linearly independent eigenvectors of $\boldsymbol{A}$. In the new coordinate system where $\boldsymbol{z}[t] = \boldsymbol{V}^{-1}\boldsymbol{x}[t]$ is the state, the dynamics (27) and (28) take the following form:

$$\begin{aligned} \boldsymbol{z}[t+1] &= \boldsymbol{M}\boldsymbol{z}[t], \\ y_i[t] &= \bar{\boldsymbol{C}}_i \boldsymbol{z}[t]. \end{aligned} \qquad (29)$$

Here, $\boldsymbol{M} = \boldsymbol{V}^{-1}\boldsymbol{A}\boldsymbol{V}$ is a diagonal matrix with diagonal entries comprising of the distinct eigenvalues $\lambda_1, \lambda_2, ..., \lambda_N$ of $\boldsymbol{A}$, and $\bar{\boldsymbol{C}}_i = \boldsymbol{C}_i \boldsymbol{V}$. Based on the above decomposition, each node i can immediately identify (locally) the set of eigenvalues $\mathcal{O}_i$ that are detectable w.r.t. its own measurements. The key observation made in [137] is the following. Each node i can estimate the components of the state vector $\boldsymbol{z}[t]$ corresponding to its detectable eigenvalues $\mathcal{O}_i$ without interacting with any neighbor. It needs to employ consensus only for estimating those components that correspond to its undetectable eigenvalues $\mathcal{U}\mathcal{O}_i$. Specifically, every regular node i employs the following scheme.

(i) For each $\lambda_j \in \mathcal{O}_i$, node i uses a standard Luenberger observer to estimate $z^{(j)}[t]$ - the component of $\boldsymbol{z}[t]$ corresponding to λ_j.

(ii) For each $\lambda_j \in \mathcal{U}\mathcal{O}_i$, node i uses a "local filtering" technique to estimate $z^{(j)}[t]$.

We now elaborate on item (ii). Let $\Lambda_U(\boldsymbol{M})$ represent the set of all unstable and marginally stable eigenvalues of $\boldsymbol{M}$. Moreover, for each $\lambda_j \in \Lambda_U(\boldsymbol{M})$, let $\mathcal{S}_j$ be the set of nodes that can detect λ_j. To enable the nodes in $\mathcal{V} \setminus \mathcal{S}_j$ to estimate $z^{(j)}[t]$, the following two requirements (stated loosely here, and more formally later) turn out to be critical.

- **Information redundancy:** The set $\mathcal{S}_j$ needs to be sufficiently large. Otherwise, if the sources of information for the mode λ_j are corrupted by too many adversaries, then the other nodes cannot hope to recover the correct information about that mode.
- **Network-structure redundancy:** There must exist sufficient disjoint paths in $\mathcal{G}$ that link the nodes in $\mathcal{S}_j$ to those in $\mathcal{V} \setminus \mathcal{S}_j$. Otherwise, the adversaries can form a bottleneck and disrupt the flow of information from $\mathcal{S}_j$ to $\mathcal{V} \setminus \mathcal{S}_j$.

In [137], the authors introduce a graph-theoretic construct called the Mode Estimation Directed Acyclic Graph (MEDAG) to capture the above requirements. Essentially, a MEDAG $\mathcal{G}_j$ for mode λ_j is a subgraph of $\mathcal{G}$ that provides a secure *uni-directional* medium of information-flow from $\mathcal{S}_j$ to $\mathcal{V} \setminus \mathcal{S}_j$.[12] The uni-directional aspect is important to ensure stability of the estimation error dynamics; it has nothing to do with adversaries. Once a MEDAG $\mathcal{G}_j$ has been constructed for each $\lambda_j \in \Lambda_U(\boldsymbol{M})$, an agent $i \in \mathcal{V} \setminus \mathcal{S}_j$ uses the estimates of only its neighbors $\mathcal{N}_i^{(j)}$ in the MEDAG $\mathcal{G}_j$ to update its estimate $\hat{z}_i^{(j)}[t]$ of $z^{(j)}[t]$ as follows:

$$\hat{z}_i^{(j)}[t+1] = \lambda_j \sum_{\ell \in \mathcal{M}_i^{(j)}[t]} w_{i\ell}^{(j)}[t] \hat{z}_\ell^{(j)}[t]. \qquad (30)$$

In the above update rule, $\mathcal{M}_i^{(j)}[t] \subset \mathcal{N}_i^{(j)} \subseteq \mathcal{N}_i$ is the set of those neighbors from whom node i accepts estimates of $z^{(j)}[t]$ at time-step t, after removing the f largest and f smallest estimates of $z^{(j)}[t]$ from $\mathcal{N}_i^{(j)}$; the properties of a MEDAG $\mathcal{G}_j$ ensure that $\mathcal{M}_i^{(j)}[t]$ is always non-empty. The weights in (30) are non-negative and chosen to satisfy $\sum_{\ell \in \mathcal{M}_i^{(j)}[t]} w_{i\ell}^{(j)}[t] = 1$. The overall approach we described above is called the Local-Filtering based Resilient Estimation (LFRE) algorithm in [137].

---

[12] For a precise description of the properties of a MEDAG, see [137] where a distributed algorithm is provided to construct such subgraphs.
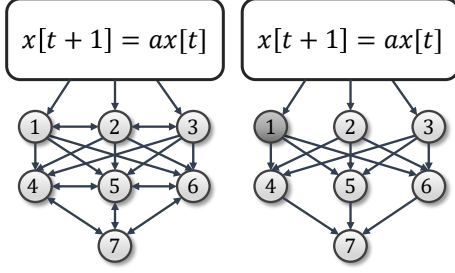
Fig. 9. (**Left**) A scalar unstable plant is monitored by a network of 7 nodes. The network is strongly 3-robust w.r.t. the set of source nodes $\mathcal{S} = \{1, 2, 3\}$. (**Right**) A subgraph of the original graph satisfying the properties of a MEDAG.

To analyze the performance of the LFRE algorithm, we need to first understand when a given network $\mathcal{G}$ contains a MEDAG. The following graph-theoretic property is what we need in this context.

Definition 7 (*strongly* r-*robust graph* w.r.t. $\mathcal{S}_j$) *For* $r \in \mathbb{N}_+$ *and* $\lambda_j \in \Lambda_U(\mathbf{A})$, *a graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ *is strongly* r-*robust w.r.t. the set of source nodes* $\mathcal{S}_j$ *if for any nonempty subset* $\mathcal{C} \subseteq \mathcal{V} \setminus \mathcal{S}_j$, $\mathcal{C}$ *is* r-*reachable.*

For an illustration of the above definition, consider the setup in Fig. 9 (taken from [137]) where a scalar unstable plant is monitored by a network of nodes. Nodes 1, 2, and 3 are the source nodes for this system, i.e., $\mathcal{S} = \{1, 2, 3\}$. The graph on the left in Fig. 9 is an example of a network that is strongly 3-robust w.r.t. the set of source nodes $\mathcal{S}$. Specifically, all subsets of $\{4, 5, 6, 7\}$ are 3-reachable (i.e., each such subset has a node that has at least 3 neighbors outside that subset). The graph on the right is an example of a MEDAG. The next result exemplifies the role played by *strong-robustness* in resilient distributed state estimation.

Theorem 14 ([137]) *Suppose* $\mathcal{G}$ *is strongly* $(2f + 1)$-*robust w.r.t.* $\mathcal{S}_j, \forall \lambda_j \in \Lambda_U(\mathbf{A})$. *Then, the LFRE algorithm guarantees* $\lim_{t \to \infty} \|\hat{\boldsymbol{x}}_i[t] - \boldsymbol{x}[t]\| = 0$ *for every regular agent* i, *despite the presence of any* f-*locally bounded set of Byzantine adversaries.*

Discussion: While the above result provides a sufficient condition for tolerating Byzantine adversaries in the context of distributed state estimation, separate necessary conditions are also identified in [137]. The latter are a blend of requirements on the observation model and the network structure, and generalize the conditions for centralized state estimation subject to attacks [44, 145].

The techniques in [137] were later generalized to account for time-varying networks in [146] and [147]. Moreover, in [148], the authors formally showed how one can incorporate the ideas of trust and diversity to relax some of the stringent redundancy requirements in [137].

One important takeaway is that the ideas of informa-

tion redundancy and network-structure redundancy are quite general, and as such, applicable beyond the specific estimation problem we considered here. Indeed, we will later briefly comment on the fact that these ideas also turn out to be crucial in the context of resilient distributed hypothesis testing/statistical inference.

Remark 1 *Unlike the* r-*robustness property that is coNP-complete to check [149], the strong* r-*robustness property in Definition 7 can be checked in polynomial-time. Aside from algorithmic differences, this highlights that the graph-theoretic condition required for resilient distributed state estimation (strong* r-*robustness) is fundamentally different from that for resilient consensus and optimization (* r-*robustness).*

2) Methods based on Robust Control Theory: In [150], the authors consider distributed state estimation of a continuous-time LTI system where both the state and measurement dynamics are subject to L_2-integrable disturbances. Given the nature of the disturbances, the authors build on the theory of distributed H_∞ filters developed in [151, 152]. The threat model is that of a *biasing attack* where an attacker injects an additive attack signal directly to the state estimator/observer dynamics of certain agents.

To tackle such biasing attacks, the authors in [150] develop certain attack detection filters. Specifically, the attack detection filter at each agent i takes as input two different innovation signals. The first innovation signal is the gap between the actual measurement $y_i[t]$ and the predicted sensor measurement $C_i \hat{x}_i[t]$. Intuitively, if agent i's observer dynamics is compromised, then one should expect $C_i \hat{x}_i[t]$ to be biased, leading to a significant deviation from $y_i[t]$. [13] With a similar motivation, a second innovation signal is computed based on the deviations of agent i's state estimate from those of its neighbors. It is shown that the problem of designing the parameters of the above detector can be recast as the problem of stabilizing a distributed dynamical system via output injection. The latter problem is addressed by drawing on ideas from vector dissipativity theory [151, 152]. Since the overall design procedure is quite intricate, we refer the reader to [150] for details.

As their main result, the authors in [150] show that if certain LMI's are feasible, then their approach ensures attack-detection and guarantees a desired level of H_∞ disturbance attenuation. The requirements on the network structure are implicitly captured by the LMI's.

Discussion: It is instructive to compare the results

---

[13] As in [150], the innovation signal $(y_i[t] - H_i \hat{x}_i[t])$ plays a key role in the saturating adaptive gain method of [129]. However, while $y_i[t]$ is accurate and $C_i \hat{x}_i[t]$ is biased in [150], the situation is exactly the opposite in [129], where the measurements are biased but the agents' estimators are not.

in [137] with those in [150]. In [137], accounting for a worst-case Byzantine attack model necessitates the requirement of sufficient redundancy in the underlying network, as captured by the "strong robustness" condition in Definition 7. Moreover, for the problem to be meaningful in [137], only a subset of the agents can be adversarial. In contrast, since a specific class of biasing attacks is considered in [150], the network requirements are relatively less stringent, and it is plausible for the entire network to be under attack.

Following up on [150], more general biasing attacks are considered in [153] where an attacker is allowed to bias both the state observer and the attack detector at an agent. Moreover, unlike [150] where the design of the filter gains involves solving LMI's that are coupled across agents, the design procedure is carried out locally at every agent in [153].

3) Methods based on Saturating Gains and Attack Detection: Recently, in [128], the authors developed secure distributed filters for tackling measurement attacks. Their main approach relies on a saturating adaptive gain technique, similar to [129]. There are, however, considerable differences with [129] that stem from the fact that [128] considers dynamical state estimation while [129] studies static parameter estimation. One such difference is that [128] employs a two-time-scale estimation technique: between two consecutive time-steps of the dynamics, the agents are allowed to perform multiple consensus steps to bridge the gap between their estimates. Under reasonable assumptions on observability, and connectivity of the graph, the estimation error is shown to be uniformly bounded.

Discussion: Notably, the analysis in [128] applies to *time-varying adversarial sets*. When the adversarial set is fixed, an attack detection algorithm is further developed in [128] that leads to tighter error bounds. It should be noted, however, that the two-time-scale approach is crucial to the stability of the distributed filter in [128].

Remark 2 *In some very recent work, the authors in [154] have studied the resilient distributed state estimation problem from a dynamic average consensus perspective. In [155], the setting where some of the communication links can also be corrupted (in addition to node attacks) has been explored.*

7.3 Hypothesis Testing and Statistical Inference

In this subsection, we will briefly discuss an approach for tackling Byzantine attacks in the context of distributed hypothesis testing [156–165] - a problem similar in flavor to the static parameter estimation setting we considered in Section 7.1. In this problem, each agent in a network receives a sequence of stochastic measurements generated by a common underlying distribution that is parameterized by an *unknown, static* state θ^*. Each agent

is equipped with a local likelihood model, and is aware that θ^* belongs to a finite set $\Theta = \{\theta_1, \ldots, \theta_m\}$ of m candidate hypotheses. The goal is for the agents to collaboratively identify θ^* from Θ. The challenge arises from the fact that θ^* may not be locally identifiable w.r.t. the likelihood model of any specific agent. In other words, no one agent can, in general, eliminate every false hypothesis on its own. Instead, we assume *global identifiability* of the joint observation model, i.e., θ^* can be uniquely identified based on the collective observations of the agents. [14]

In a typical approach to solving the above problem, each agent maintains a belief vector $\boldsymbol{\mu}_{i,t}$ which is a distribution over Θ. Formally, the objective is to design belief-update and propagation rules that ensure $\mu_{i,t}(\theta^*) \to 1, \forall i \in \mathcal{V}$, almost surely. The predominant approach is to employ some form of "belief-averaging" protocol to update the belief vectors [156–163]. In a departure from these algorithms, a min-rule was recently developed in [165]; the asymptotic learning rate of this rule is strictly better than those based on averaging. However, all of the above approaches are vulnerable in the face of attacks: a single malicious agent can essentially cause all good agents to eliminate the true hypothesis.

The Min-Rule: A desirable feature of the min-rule in [165] is that it admits a simple, computationally-efficient extension that is robust to worst-case attacks. We first describe the basic min-rule, and then its adversarial extension. Each agent $i \in \mathcal{V}$ maintains an auxiliary local belief vector $\boldsymbol{\pi}_{i,t}$ that is updated in a Bayesian manner based on just the observations of agent i. For every false hypothesis $\theta \in \Theta \setminus \{\theta^*\}$ that agent i can eliminate on its own, we will have $\pi_{i,t}(\theta) \to 0$ almost surely. Thus, agent i only needs to interact with neighbors for eliminating those false hypotheses that it cannot rule out on its own. Let $\mathcal{S}(\theta^*, \theta)$ be those agents that can distinguish between θ^* and θ, i.e., these agents can eliminate θ individually. The main idea is to transmit low beliefs on θ from agents in $\mathcal{S}(\theta^*, \theta)$ to the rest of the network. This is achieved via the following rule at each agent i:

$$\mu_{i,t+1}(\theta) \propto \min\{\{\mu_{j,t}(\theta)\}_{j \in \mathcal{N}_i \cup \{i\}}, \pi_{i,t+1}(\theta)\}. \quad (31)$$

The above beliefs are normalized to ensure that $\boldsymbol{\mu}_{i,t}$ is a valid distribution at every time-step. For the adversarial setting, the approach is very similar, except that the min-rule is applied to a set of *moderate beliefs*. Specifically, each regular agent i updates its belief on a state θ by first rejecting the highest f and lowest f beliefs on θ

---

[14] Global identifiability for distributed hypothesis testing is the exact analogue of global observability for distributed state estimation.

received from $\mathcal{N}_i$, and then employing:

$$\mu_{i,t+1}(\theta) \propto \min\{\{\mu_{j,t}(\theta)\}_{j \in \mathcal{M}_{i,t}^\theta}, \pi_{i,t+1}(\theta)\}, \qquad (32)$$

where $\mathcal{M}_{i,t}^\theta$ are those agents that do not get rejected in the above filtering step. This is known as the Local-filtering based Resilient Hypothesis Elimination (LFRHE) algorithm. The correctness of the LFRHE algorithm once again rests on the two key ingredients we identified in Section 7.1, namely, information-redundancy and network-structure redundancy. In particular, for every pair θ_p, θ_q, we need $\mathcal{S}(\theta_p, \theta_q)$ to be large enough, and we also need a sufficient number of disjoint paths from $\mathcal{S}(\theta_p, \theta_q)$ to $\mathcal{V} \setminus \mathcal{S}(\theta_p, \theta_q)$. These requirements are succinctly captured in the following theorem.

Theorem 2 ([165]) *Suppose that for every pair of hypotheses $\theta_p, \theta_q \in \Theta$, the graph $\mathcal{G}$ is strongly $(2f+1)$-robust w.r.t. the source set $\mathcal{S}(\theta_p, \theta_q)$. Moreover, suppose each regular agent i has a non-zero prior belief on every hypothesis, i.e., $\pi_{i,0}(\theta) > 0$ and $\mu_{i,0}(\theta) > 0$, $\forall \theta \in \Theta$. Then, the LFRHE algorithm guarantees that $\mu_{i,t}(\theta^*) \to 1$ almost surely for every regular agent i, despite the actions of any f-local set of Byzantine adversaries.*

Discussion: One of the main takeaways from the above result is that just like the resilient distributed estimation problem, the strong-robustness property in Definition 7 ends up playing a crucial role when it comes to tolerating Byzantine attacks for distributed hypothesis-testing as well. We conjecture that this graph-theoretic property will prove to be useful for other distributed learning problems where information is diffused across the network. A recent work that studies resilient distributed best-arm identification for stochastic multi-armed bandits supports this conjecture [166].

Additional Results: The algorithm in [165] was later extended in [167] to account for time-varying networks. In [162], the authors proposed an alternate approach to tackling adversaries by building on the log-linear belief-update rule in [159–161]. Their approach requires the agents to compute Tverberg partitions (see Section 5.3); however, as discussed in that section, there is no known algorithm that can compute an exact Tverberg partition in polynomial time for a general d-dimensional finite point set [111].

8 Attack Detection and Identification Over Networks

In the preceding sections, we covered several techniques for solving a variety of distributed information-processing problems subject to attacks (e.g., consensus, optimization, and estimation). As discussed in those sections, depending on the attack model and nature of the information available to the nodes, detection

and identification of (worst-case) adversarial behavior may be impossible in general. In particular, the "local-filtering" algorithms discussed in those sections did not explicitly rely on detection/identification of adversarial behavior. However, in other settings, detection and identification of adversarial behavior may indeed be possible. The purpose of this section is to briefly summarize algorithms for such settings.

Since the precise nature of the attack detection algorithm is usually dictated by the specific distributed task at hand, we will not be able to cover all such detection mechanisms here. Instead, we will primarily restrict our attention to the attack model in Eq. (8) of Section 4.2 that we studied in the context of distributed function calculation. Using this model, we will discuss graph-theoretic requirements for detecting and identifying attacks in a network. In particular, we will demonstrate how *structured systems theory* plays a key role in this context. Before delving into the technical details, we remind the reader that the f-total attack model under consideration involves a set of malicious nodes $\mathcal{F}$, where $|\mathcal{F}| \leq f$ for a known $f \geq 0$.

8.1 Attack Detection

Centralized and distributed detection techniques can be used to detect attacks, see [35]. Here, we characterize graph-theoretic conditions for detecting attacks. We start by considering that the initial states, $\psi[0]$, are known. An attack ζ is called *undetectable* or *perfect* if $y(\psi[0], \zeta, t) = y(\psi[0], 0, t)$ for all $t \geq 0$, i.e., the measurement is the same as the case of no attack. The notion of a perfect attack has an equivalent algebraic condition, which is based on the following definition.

Definition 8 *The generic normal rank (gnr) of the matrix pencil of dynamics* (8)

$$\boldsymbol{P}(z) = \begin{bmatrix} \boldsymbol{W} - z\mathbf{I}_n & \boldsymbol{B}_{\mathcal{F}} \\ \boldsymbol{C} & \boldsymbol{0} \end{bmatrix},$$

is the maximum rank of the matrix over all choices of free (nonzero) parameters in $(\boldsymbol{W}, \boldsymbol{B}_{\mathcal{F}}, \boldsymbol{C})$ and $z \in \mathbb{C}$.

It is shown in [83] that having a perfect attack and the generic normal rank of $\boldsymbol{P}(z)$ being less than $n + |\mathcal{F}|$ are equivalent . Recalling the input set $\mathcal{U}$ and measurement set $\mathcal{Y}$ for a structured system from Section 2.4, the following result interprets the generic normal rank of $\boldsymbol{P}(z)$ in terms of the disjoint paths in the graph of structured system $\mathcal{G}$.

Lemma 15 ([168]) *The generic normal rank of the matrix pencil $\boldsymbol{P}(z)$ is equal to $n + r$, where r is the size of the largest linking in $\mathcal{G}$ from the input vertices, $\mathcal{U}$, to the output vertices, $\mathcal{Y}$.*

Note that the generic normal rank of the matrix pencil is at least n, since the matrix $\boldsymbol{W} - z\mathbf{I}_n$ will have generic rank n for any choice of parameters in $\boldsymbol{W}$ and any z that is not an eigenvalue of $\boldsymbol{W}$. Lemma 15 implies that to prevent perfect attacks, parameter r has to be equal to the number of attacks, i.e., $r = |\mathcal{F}|$. This implicitly indicates that the number of sensors must be at least $|\mathcal{F}|$. Lemma 15 along with Menger's theorem and Expansion lemma, c.f. [86], yield the following graph-theoretic result on attack detectability.

Theorem 16 *Suppose that dynamics* (8) *with measurement* (9) *is subject to a set of f attacked nodes and that the initial states, $\boldsymbol{\psi}[0]$, are known. To prevent a perfect attack, it is sufficient for graph $\mathcal{G}$ to be $(f+1)$-connected.*

Remark 17 *As the number of sensors is often limited and the underlying network may be sparse, detecting all attacked nodes may not be always possible. An alternative approach is to place the available sensors on key nodes in the network in order to maximize r in Lemma 15, i.e., detecting maximum number of attacks. The sensor placement problem for optimal attack detection is discussed in Section 10.2.*

For the cases where the initial condition of the system is unknown, an undetectable attack $\boldsymbol{\zeta}$ is characterized by the existence of a pair of initial states $\boldsymbol{\psi}_1[0]$ and $\boldsymbol{\psi}_2[0]$ such that $\boldsymbol{y}(\boldsymbol{\psi}_1[0], 0, t) = \boldsymbol{y}(\boldsymbol{\psi}_2[0], \boldsymbol{\zeta}, t)$ for all $t \geq 0$. In such cases, one needs to first recover the initial conditions of the system in order to detect the attack. This demands the system to be strongly observable. Recall from Section 4 and Theorem 3 that system (8) is strongly observable if the graph is $(2f+1)$-connected. This extra level of graph connectivity, compared to Theorem 16, is a price paid for the lack of the knowledge of the initial states.

Theorem 18 *Suppose that dynamics* (8) *with measurement* (9) *is subject to a set of f attacked nodes and unknown initial conditions. To prevent a perfect attack, it is sufficient for graph $\mathcal{G}$ to be $(2f+1)$-connected.*

8.2 Attack Identification Procedure

In order to identify the attacked nodes, first, node i must find the true initial value of all other nodes (we assume that it satisfies the condition in Theorem 3). After obtaining the vector of initial states, $\boldsymbol{\psi}[0]$, and assuming that the interaction matrix, $\boldsymbol{W}$, is known to i, it can apply dynamics (8) to obtain $\boldsymbol{\psi}[1] - \boldsymbol{W}\boldsymbol{\psi}[0] = \boldsymbol{B}_{\mathcal{F}}\boldsymbol{\zeta}[0]$. Every nonzero component in the vector on the left hand side of this equation indicates an additive error injected by the corresponding node. Thus, every node that is malicious during time-step 0 can be identified by this method. The same process can be repeated to find all nodes that were malicious during the first L time steps from the transmitted values $y_i[0:L]$ in (10). Note that using iteration policy (8) to identify the attacks requires

the system to be free of noise or external disturbances, as otherwise, they mislead us in the above procedure. Further details on centralized and distributed attack detection and identification techniques can be found in [35].

8.3 Other Attack Detection and Identification Approaches

Our discussion in this section has thus far focused on the attack model in Eq. (8) for distributed function calculation. Before closing this section, we briefly summarize certain other attack detection mechanisms that are relevant in other contexts. For resilient distributed parameter estimation, attack detection constitutes a key component of the $\mathcal{FRDE}$ algorithm in [127] that we discussed in Section 7.1. For state estimation, agents under attack are identified based on robust control techniques in [150], and a saturating adaptive gain method in [128]; see Section 7.2. At a high-level, the approaches in [127], [150], and [128] share a common principle for detecting whether an agent i is under attack: they involve computing an appropriate innovation signal that captures the extent to which the estimates/measurements of agent i differ from those of its neighbors. For distributed optimization, [115] proposed a heuristic gradient based method to detect the misbehaving agents. The insight is that the attackers' biasing actions can result in a large gradient value. Hence, a regular agent can attempt to detect such attacks by approximating the gradient of each neighbor and track it over time relative to the mean of the gradients of the remaining neighbors. Finally, we note that in the context of multi-robot coordination, the authors in [169, 170] propose methods to tackle the so called "Sybil attack", where an attacker spoofs or impersonates the identities of existing agents to gain a disproportionate advantage in the network. The key idea in these works is to detect such spoofing attacks by exploiting the physics of wireless signals. For specific details, we refer the reader to the respective papers.

9 Graph-Theoretic Interpretations of the Attack Impact

In the previous sections, we discussed how to use network connectivity to withstand the adversarial actions in several distributed algorithms. In some cases, however, it is not possible to completely nullify the attacker's actions. In those cases, an alternative approach is to mitigate its impact (as was done in the local filtering approaches for consensus and optimization). In this section, we quantify the attacker's impact on networked control systems in terms of the topology of the underlying network.

9.1 Controllability of Networks Under Attack

One way to define the attacker's impact is via the largest subset of nodes which can be controlled (or reached) by
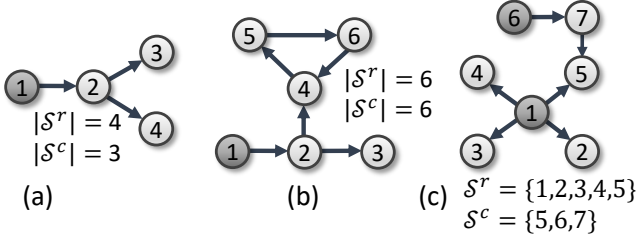
Fig. 10. Graphs and their largest reachable and controllable subsets.

the set of the attacked nodes. There is a vast literature that studies the controllability of networks with a limited number of actuators [83, 171–174]. In this section, our focus is on structural controllability as a discrete measure of controllability, and Gramian-based methods as continuous measures.

9.1.1 Structural and Strong Structural Controllability

As mentioned in Section 2.4, a structured system is said to be controllable if this property holds for at least one numerical choice of free (nonzero) parameters in the system. Theorem 1 provided graph-theoretic conditions for structural controllability. Condition (i) in Theorem 1 is called the reachability condition. The reachability, by itself, can be a measure of the attack impact. In particular, when the attacker's goal is to disseminate a signal throughout the network and infect as many nodes as possible, maximizing reachability is beneficial to the attacker. On the other hand, when the attacker's objective is to steer the states towards its desired direction, controllability is the appropriate measure. We present the following definition.

Definition 9 *The reachable set, $\mathcal{S}^r(\mathcal{F})$ (respectively controllable set, $\mathcal{S}^c(\mathcal{F})$) for a set of attacked nodes $\mathcal{F}$, is a subset of nodes in graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ whose elements can all be reached (controlled) by the set of attacked nodes $\mathcal{F}$ and is maximal in the sense that no additional node can be added to $\mathcal{S}^r(\mathcal{F})$ ($\mathcal{S}^c(\mathcal{F})$) without breaking this property. Elements of the attacked set, $\mathcal{F}$, are called the source nodes of $\mathcal{S}^r(\mathcal{F})$ (or of $\mathcal{S}^c(\mathcal{F})$).*

Based on Theorem 1, it is clear that $|\mathcal{S}^c(\mathcal{F})| \leq |\mathcal{S}^r(\mathcal{F})|$. Fig. 10 (a) shows a graph whose largest reachable set, for $|\mathcal{F}| = 1$, is larger than the largest controllable set. In graph (b) the largest reachable and controllable sets are identical. Note that $\mathcal{S}^c(\mathcal{F})$ is not necessarily a subset of $\mathcal{S}^r(\mathcal{F})$ as shown in graph (c). The source nodes are shown with a darker color. Due to resource constraints, the attacker naturally tries to solve either of the following problems:

(i) Controlling (or reaching) the *largest* possible subset of nodes in the network with a given number of attacked nodes, or

(ii) Controlling (or reaching) the whole network with *minimum* number of attacked nodes.

The latter has been investigated under the context of minimal structural controllability problems [175]. It is shown that the problem of finding the minimum number of control input (attacked) nodes is in general NP hard and in specific cases (such as dedicated inputs) can be solved in polynomial time. When reachability is the objective, the problem is related to estimating the reachable set of nodes which can be solved in polynomial time [176].

A system is called strong structurally controllable if $\text{rank}(\boldsymbol{W}, \boldsymbol{B}) = n$ for *all* (nonzero) choice of free parameters in $\boldsymbol{W}$ and $\boldsymbol{B}$. The dimension of the strong structurally controllable subspace is the minimum rank of the controllability matrix, $\mathcal{C}(\boldsymbol{W}, \boldsymbol{B})$. There are graph-theoretic bounds on this quantity for consensus dynamics as stated below.

Theorem 19 ([177]) *For set $\mathcal{F}$ chosen by the attacker, the dimension of the strong structurally controllable subspace is lower bounded by*

$$\text{rank}\,\mathcal{C}(\boldsymbol{W}, \boldsymbol{B}) \geq \max_{i \in \mathcal{V}, j \in \mathcal{F}} \text{dist}(i, j) + 1, \qquad (33)$$

where $\text{dist}(i, j)$ is the shortest distance between nodes i and j. Tighter lower bounds for the dimension of the strong structurally controllable subspace can be found in [177]. According to Theorem 19, from the attacker's perspective, the optimal decision is to select a node with maximum distance from the rest of the nodes in the graph. Another interpretation from the above result, which had been discussed before in [171], is the reverse effect of network connectivity on the controllability. Specifically, sparse networks may contain pairs of nodes that are far apart, and consequently may have large controllability subspaces, while well-connected networks may have smaller controllability subspaces.

9.1.2 Gramian-Based Controllability

Unlike the (discrete) structural or rank-based controllability measures of dynamical systems [178], the controllability Gramian provides a continuous measure for this property, in the form of the energy required to drive the dynamical system towards specific directions in the state space. The T-step controllability Gramian is defined as

$$\mathcal{W}_{\mathcal{F}, T} \triangleq \sum_{\tau=0}^{T-1} \boldsymbol{W}^\tau \boldsymbol{B}_{\mathcal{F}} \boldsymbol{B}'_{\mathcal{F}} (\boldsymbol{W}')^\tau \qquad (34)$$

where $\boldsymbol{B}_{\mathcal{F}}$ corresponds to the set of attacked nodes and defined in (8). The controllability Gramian $\mathcal{W}_{\mathcal{F}, T}$ is positive definite if and only if the system is controllable

in T steps [179]. However, even if a system is controllable, certain directions of the state space may be hard to reach [180]. The smallest eigenvalue of the Gramian, $\lambda_1(\mathcal{W}_{\mathcal{F},T})$, is inversely related to the amount of energy required to move the system in the direction that is the most difficult to control, i.e., the eigenvector corresponding to $\lambda_1(\mathcal{W}_{\mathcal{F},T})$. Other controllability metrics, such as $\mathrm{trace}(\mathcal{W}_{\mathcal{F},T})$ and $\mathrm{trace}(\mathcal{W}_{\mathcal{F},T}^{-1})$, quantify the energy needed on average to move the system around on the state space. From the attacker's perspective, the system should be easily controllable. Thus, it targets nodes in which one of the above mentioned spectra is optimized. In particular, the attacker attempts to minimize its effort to steer the system by maximizing $\lambda_1(\mathcal{W}_{\mathcal{F},T})$ and $\mathrm{trace}(\mathcal{W}_{\mathcal{F},T})$ or minimizing $\mathrm{trace}(\mathcal{W}_{\mathcal{F},T}^{-1})$.

The selection of control nodes to optimize the spectrum of the Gramian does not generally admit a closed-form solution. The exception is $\mathrm{trace}(\mathcal{W}_{\mathcal{F},T})$ for which [181]

$$\mathrm{trace}(\mathcal{W}_{\mathcal{F},T}) = \sum_{i\in\mathcal{F}}\left(\sum_{\tau=0}^{T-1}\boldsymbol{W}^{2\tau}\right)_{ii}.$$

When $\boldsymbol{W}$ is Schur stable, then $\sum_{\tau=0}^{\infty}\boldsymbol{W}^{2\tau} = (I - \boldsymbol{W}^2)^{-1}$. For continuous time systems, if $\boldsymbol{W}$ is Hurwitz, this closed form solution becomes $\mathrm{trace}(\mathcal{W}_{\mathcal{F}}) = \sum_{i\in\mathcal{F}}\left(\boldsymbol{W}^{-1}\right)_{ii}$. For other metrics, selecting optimal nodes, from the attacker's perspective, is a combinatorial problem and (in general) hard to solve. However, recent studies on submodularity and monotonicity of some of those metrics indicate that greedy algorithms for selecting the control nodes result in a sub-optimal solution with a guaranteed performance bound [182–185].

In the following subsection, we outline a graph-theoretic interpretation of the attacker's strategy to optimize $\mathrm{trace}(\mathcal{W}_{\mathcal{F},T})$ in consensus dynamics.

9.1.3 Case Study: Consensus Dynamics

We consider Gramian-based controllability on two types of consensus dynamics on undirected graphs, namely average consensus and leader-following consensus.

1. Edge Attack in Average Consensus: Assume that the attack happens in the form of a flow which enters one node and exists from another node. In particular, we say that $ik \in \mathcal{F}$ if the pair i, k are chosen by the attacker and their dynamics are

$$\dot{\psi}_i = \sum_{j\in\mathcal{N}_i^{\mathrm{in}}}(\psi_j - \psi_i) + \zeta_{ik},$$

$$\dot{\psi}_k = \sum_{j\in\mathcal{N}_k^{\mathrm{in}}}(\psi_j - \psi_k) - \zeta_{ik}, \qquad (35)$$
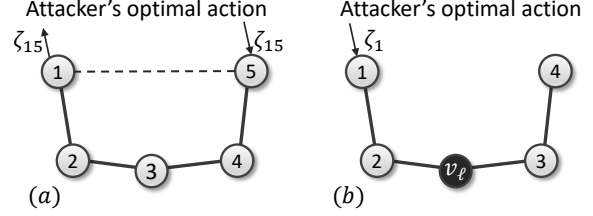

Fig. 11. (a) Attacker's optimal decisions based on (39), (b) attacker's optimal decision based on (41).

where ζ_{ik} is the attack flow. This type of input signal, as schematically shown in Fig. 11 (a), happens in power systems (DC input links) and distribution networks [186, 187]. If we write the dynamics in vector form, it becomes

$$\dot{\boldsymbol{\psi}} = -L\boldsymbol{\psi} + \mathcal{B}\boldsymbol{\zeta}, \qquad (36)$$

where $\mathcal{B}$ is the incidence matrix of the graph induced by the attacked edges. The controllability Gramian is

$$\mathcal{W}_c = \int_0^\infty \mathbf{e}^{-L\tau}\mathcal{B}\mathcal{B}'\mathbf{e}^{-L'\tau}d\tau. \qquad (37)$$

Since L is marginally stable, the infinite integral does not exist. However, the eigenvector of the marginally stable eigenvalue belongs to the subspace corresponding to the consensus value, which is of little interest to the attacker (as otherwise no attack would be needed). We remove this subspace by grounding one node (removing a row and column corresponding to that node) which makes the Laplacian non-singular. The *grounded Laplacian* matrix induced by the grounded node v is denoted by L_{fv} (or simply L_f). Then we have

$$\mathrm{tr}(\mathcal{W}_c) = \mathrm{tr}\left(\mathcal{B}'\int_0^\infty\mathbf{e}^{-2L_f\tau}d\tau\mathcal{B}\right) = \frac{1}{2}\mathrm{tr}\left(\mathcal{B}'L_f^{-1}\mathcal{B}\right). \qquad (38)$$

It was shown in [188] that the above value is independent of the choice of the grounded node, i.e., we have

$$\mathrm{tr}\left(\mathcal{B}'L_f^{-1}\mathcal{B}\right) = \mathrm{tr}\left(\mathcal{B}'L^\dagger\mathcal{B}\right) = \sum_{ij\in\mathcal{F}}\mathfrak{R}_{ij}, \qquad (39)$$

where $L^\dagger$ is the Moore–Penrose inverse of L and $\mathfrak{R}_{ij}$ is the effective resistance between nodes i and j. Equation (39) indicates that the trace of the controllability Gramian is the summation of effective resistances between node pairs chosen by the attacker. Thus, if the attacker seeks to maximize (39), i.e., minimize the attack energy, by choosing m node pairs, it should choose m pairs with *the largest effective resistance* in the graph. Figure 11 (a) is an example of an optimal attack.

2. Node Attack in Leader-Follower Consensus: We consider a leader-follower dynamical system on undi-

rected graphs which is widely studied in formation control problems [171, 189, 190]. Based on this model, there is a leader, which has access to the control input or determines the set-point, and a set of followers, which follow the state of the leader. The dynamics of leader and followers are given by

$$\begin{bmatrix} \dot{\boldsymbol{\psi}}_f(t) \\ \dot{\boldsymbol{\psi}}_\ell(t) \end{bmatrix} = -\underbrace{\begin{bmatrix} L_f & L_{\ell f} \\ L_{f\ell} & L_\ell \end{bmatrix}}_{L} \begin{bmatrix} \boldsymbol{\psi}_f(t) \\ \boldsymbol{\psi}_\ell(t) \end{bmatrix} + \begin{bmatrix} \boldsymbol{B}_{\mathcal{F}} \\ 0 \end{bmatrix} \boldsymbol{\zeta}_{\mathcal{F}}(t),$$

(40)

where $\boldsymbol{\psi}_f(t), \boldsymbol{\psi}_\ell(t)$ and $\boldsymbol{\zeta}_{\mathcal{F}}(t)$ are the state of followers, the state of the leader, and the attack vector, respectively. The leader is not affected by communication attacks. It also keeps its state constant and does not incorporate the states of the followers, i.e., $\dot{\boldsymbol{\psi}}_\ell(t) = 0$. Thus, we have $L_{f\ell} = L_\ell = 0$. Matrix $\boldsymbol{B}_{\mathcal{F}}$, formed from indicator vectors as in (8), determines the nodes that are under attack and $L_{\ell f}$ determines the connection of the leader to the followers. Matrix L_f is the grounded Laplacian matrix discussed before. One of the key properties of this matrix, which is used in this analysis, is that $[L_f^{-1}]_{ii} = \mathfrak{R}_{i\ell}$, where $\mathfrak{R}_{i\ell}$ is the effective resistance between node i and the leader. Similar to the case of edge attack, the objective of the attacker is to maximize the trace of the controllability Gramian. Calculating the Gramian integral yields

$$\text{trace}(\mathcal{W}_{\mathcal{F}}) = \sum_{i \in \mathcal{F}} \left(L_f^{-1} \right)_{ii} = \sum_{i \in \mathcal{F}} \mathfrak{R}_{i\ell}.$$

(41)

Based on (41), if the attacker wants to minimize the average energy by attacking m nodes, it must select m nodes with the largest effective resistance from ℓ. Figure 11 (b) is an example of an optimal attack.

9.2 System Norm Approaches

Another way to quantify the attacker's impact is through the use of system norms from the attack signal to the output of interest, e.g., the state of the agents. Two widely used system norms in robust control are $\mathcal{H}_2$ and $\mathcal{H}_\infty$ norms. Since these methods were initially developed on continuous time systems, we focus on such systems in this section.

Suppose that the evolution of a network of agents is described by (4) where $\boldsymbol{D} = 0$ and $\boldsymbol{B} = \boldsymbol{B}_{\mathcal{F}}$ which corresponds to the set of attacked nodes defined in (8). The transfer function of the error dynamics from attack input $\boldsymbol{\zeta}_{\mathcal{F}}(t)$ to output $\boldsymbol{y}(t)$ in the Laplace domain is $G(s) \triangleq \boldsymbol{C}(sI - \boldsymbol{W})^{-1}\boldsymbol{B}_{\mathcal{F}}$. The system $\mathcal{H}_2$ and $\mathcal{H}_\infty$

norms are defined as

$$\|G\|_2 \triangleq \left(\frac{1}{2\pi} \text{trace} \int_0^\infty G^*(j\omega)G(j\omega)d\omega \right)^{\frac{1}{2}},$$
$$\|G\|_\infty \triangleq \sup_{\omega \in \mathbb{R}} \sigma_{\max}(G(j\omega)),$$

(42)

where $\sigma_{\max}(.)$ is the maximum singular value of a matrix. The system $\mathcal{H}_2$ norm can also be calculated based on the observability Gramian $\mathcal{W}_o$, which is the solution of the following Lyapunov equation

$$\boldsymbol{W}'\mathcal{W}_o + \mathcal{W}_o\boldsymbol{W} = -\boldsymbol{C}'\boldsymbol{C},$$
$$\|G\|_2 = \text{trace}(\boldsymbol{B}_{\mathcal{F}}'\mathcal{W}_o\boldsymbol{B}_{\mathcal{F}}).$$

(43)

Unlike the approach of the structured system theory to resilient NCSs, which was based on nonzero patterns of system matrices and not the magnitude of the elements, the system norm approach depends on the exact value of the matrix elements. Thus, having knowledge about the nonzero patterns of dynamic matrices is not sufficient and one has to specify the type of matrices which describe the interactions between agents. We revisit the leader-follower consensus dynamics discussed in Section 9.2.1 to further explain this fact.

Remark 3 ($\mathcal{H}_2$ *versus* $\mathcal{H}_\infty$ *norm*) *From a security perspective, either the system $\mathcal{H}_2$ or $\mathcal{H}_\infty$ norm can be used to quantify the attack impact. If the frequency content of the attack signal is unknown, using a $\mathcal{H}_2$ norm is a more reasonable choice as it is calculated over all frequencies. However, if the objective is to find the worst-case attack impact over all frequencies, the system $\mathcal{H}_\infty$ norm is an appropriate choice.*

9.2.1 Case Study: Consensus Dynamics Revisited

1. Edge Attack in Average Consensus: Consider the consensus dynamics under an attack flow as in (36). Suppose that the attacker measures the difference between node values under attack, i.e., $\boldsymbol{y} = \mathcal{B}'\boldsymbol{\psi}$. Dynamics (36) together with this measurement form a symmetric system, i.e., L is symmetric and $\boldsymbol{B} = \boldsymbol{C}'$. Thus, the $\mathcal{H}_\infty$ norm is equal to the DC gain of the system [191], i.e., for the transfer function from $\boldsymbol{\zeta}$ to $\boldsymbol{y}$ we have

$$\|G\|_\infty = \sigma_{\max}\left(\mathcal{B}' L^\dagger \mathcal{B}\right)$$

(44)

Unlike the trace of $\mathcal{B}' L^\dagger \mathcal{B}$, in (39), interpreting its largest singular value is hard. We consider the simple case where only one node pair, i and j, is under attack, i.e., $\mathcal{B} = \mathbf{e}_{ij}$. In this case, (44) becomes scalar and we have $\|G\|_\infty = \sigma_{\max}\left(\mathbf{e}_{ij}' L^\dagger \mathbf{e}_{ij}\right) = \mathfrak{R}_{ij}$. Thus, in order to have a large impact, the attacker must choose node pairs with the largest effective resistance in the network.

2. Node Attack in Leader-Follower Consensus: Here, a single attacker i targets a set of nodes

26

and $\boldsymbol{B}_{\mathcal{F}} = C'$. As before, the $\mathcal{H}_{\infty}$ norm is equal to the DC gain of the system [82] and we have $\|G\|_{\infty} = \sigma_{\max}\left(\boldsymbol{B}_{\mathcal{F}}' L_f^{-1} \boldsymbol{B}_{\mathcal{F}}\right)$. For the case of a single node under attack, we have $\boldsymbol{B}_{\mathcal{F}} = \mathbf{e}_i$ and $\|G\|_{\infty} = \mathbf{e}_i' L_f^{-1} \mathbf{e}_i = \Re_{i\ell}$. Hence, to have a large impact on a target node, the attacker must choose a node in the network with the largest effective resistance from the leader.

10 Related Problems in Resilient Networked Control Systems

In this section, we briefly discuss other problems on the resilience of networked control systems which use graph theory as a tool in the analysis.

10.1 Resilience to Actuator/Sensor and Link Removals

So far, we discussed the case where the agents or the communications between the agents are under attack. In that context, we assume that the control inputs are not affected by the attacker. In some situations, however, the attacker may choose to remove certain sensors, actuators, or communication links entirely. In those cases, the main concern is to retain the controllability of the system. This can be written in terms of a *robust structural controllability* problem. In particular, the objective is to maintain a system's controllability despite the removal of a subset of actuators.

Consider the linear time invariant system (3) where a subset of control inputs are removed (potentially due to adversarial actions). In this case, the minimum number of actuators which retain the controllability of the system is determined by the following problem

$$\begin{array}{ll} \underset{\boldsymbol{B} \in \mathbb{R}^{n \times m}}{\arg\min} & \|\boldsymbol{B}\|_0 \\ \text{such that} & (\boldsymbol{W}, \boldsymbol{B}_{\mathcal{U} \backslash \mathcal{F}}) \text{ is struc. cont. } \forall \mathcal{F} \subset \mathcal{U}, \\ & |\mathcal{F}| \le m, \end{array} \quad (45)$$

where $\boldsymbol{B}_{\mathcal{U} \backslash \mathcal{F}}$ corresponds to the structure of the input matrix $\boldsymbol{B}$ whose columns corresponding to set $\mathcal{F}$ are removed (i.e., actuators are failed). The number of actuator faults are upper bounded by m. It is shown that the above problem is NP-hard and polynomial time algorithms to approximate the solution of those problems have been proposed [192, 193].

In other set of problems, the attacker targets a set of links in the network to remove. Several performance measures may be affected by such an action. When robust controllability (or observability) is of interest, the problem can

be written as follows [192]

$$\begin{array}{ll} \underset{\boldsymbol{B} \in \mathbb{R}^{n \times m}}{\arg\min} & \|\boldsymbol{B}\|_0 \\ \text{such that} & (\boldsymbol{W}_{\mathcal{E}_{\mathcal{X},\mathcal{X}} \backslash \mathcal{E}_{\mathcal{F}}}, \boldsymbol{B}) \text{ is struc. cont. } \forall \mathcal{E}_{\mathcal{F}} \subset \mathcal{E}_{\mathcal{X},\mathcal{X}}, \\ & |\mathcal{E}_{\mathcal{F}}| \le \bar{m}, \end{array}$$

$$(46)$$

where $\mathcal{E}_{\mathcal{F}}$ is the set of edges affected by the attacker, upper bounded by $\bar{m}$. Similar to (45), the above problem is NP-hard in general.

10.2 Strategic Sensor and Actuator Placement on Graphs

Our focus in the previous sections was to find graph conditions which ensure resiliency to attacks for certain distributed algorithms. However, in many situations, the underlying network topology is sparse and cannot be changed. Furthermore, the number of sensors/actuators is limited. In these cases, an alternative approach is to place those limited number of sensors (or actuators) on specific nodes in the network in order to optimally detect the attack or mitigate its impact.

There is a vast literature on sensor (or actuator) placement to enhance the observability (or controllability) in terms of the rank of the observability (or controllability) matrix, Gramian-based metrics (as discussed previously), or the error variance of the Kalman filters [194, 195]. In all these problems, there is a single decision maker which deploys sensors (or place actuators) on a set of nodes. However, in security problems, the adversary plays as another decision maker which tries to optimize its own cost function, e.g., maximize impact or minimize visibility. This introduces a *strategic* sensor (or actuator) placement problem, taking the attacker's actions into account. In this direction, game theory can be used as a powerful tool to address this set of problems; see [51, 196–199] and references therein.

Strategic sensor placement in the network to detect cyber-attacks has been recently studied [200–203]. In this setting, the attacker seeks to apply attack inputs while being stealthy and the detector tries to detect the attack. Several approaches have been adopted to characterize the equilibria of the security games. Nash equilibrium is used to model simultaneous decision making and Stackelberg game model is used for the case where the defender must act before the attacker. In design problems, the Stackelberg game is a popular approach to defend against cyber-attacks. In particular, the detector acts as the game leader and places sensors on nodes considering the worst case attack strategies. The applicability of each method, based on the nature of the attack and the structure of the cyber physical system, is discussed in [196].

In addition to the strategic attack detection, a defense

mechanism can help mitigate the impact of the attack via certain control actions [51, 204–206]. The attacker's impact can be quantified by either of the methods discussed in Section 9. A comprehensive defense strategy must include both strategic detection and mitigation mechanisms. This is a potential avenue for further research.

10.3 Network Coherence as a Measure of Resiliency

In the network control systems literature, the notion of network coherence is used to quantify the ability of a network to reject communication disturbances while performing a formation control or a consensus algorithm in large scale systems [207]. It is usually described in terms of system $\mathcal{H}_2$ or $\mathcal{H}_\infty$ norms from the disturbance signal to the output of interest, e.g., position (phase) or velocity (frequency) [208–210]. Interpreting the disturbances as attack inputs, some of the approaches and the results in this line of research, e.g., scalability of algorithms and leader selection schemes, can be readily used in resilient distributed algorithms, as we will briefly discuss in the following paragraphs.

10.3.1 Leader Selection

The objective is to choose the optimal nodes in the network as leaders, i.e., nodes which receive the control signal, such that the network coherence is maximized, i.e., system $\mathcal{H}_2$ or $\mathcal{H}_\infty$ norms from the disturbance input to the output of interest are minimized. It is shown in [211] that the optimal leader to minimize the $\mathcal{H}_2$ norm in consensus dynamics is the *information central* node in the network: a node in which the summation of effective resistances to the rest of the nodes in the network is minimized. It is also shown that the leader which optimizes the $\mathcal{H}_\infty$ norm is not necessarily the graph's information center [189]. Moreover, a graph-theoretic condition for the leader to co-optimize both metrics is discussed in [189].

10.3.2 Scalability

Another problem of interest is the scalability of the network coherence in graphs with various structures. In [207], it is discussed that how the network coherence scales with the network size for regular lattices in 1, 2 and higher dimensions. The scalability of these metrics is also discussed for random graphs. In particular, for Erdős-Rényi Random Graphs and random regular graphs, tight characterizations of the network coherence are discussed in [189, 212].

11 Resilience in Certain Classes of Graphs

In this section, we summarize some pertinent topological properties of certain classes of graphs. We focus on
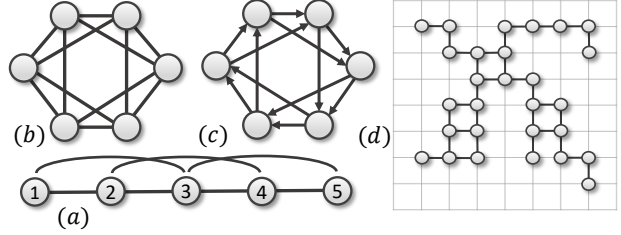

Fig. 12. (a) A 2-nearest neighbor path (b,c) an undirected and a directed circulant graph, (d) A connected formation on a 2-dimensional lattice.

the properties of graph connectivity and graph robustness provided in Section 2.1.2, since they are particularly relevant to the resilience of distributed algorithms against adversarial actions. Our focus is on undirected networks, unless indicated. We start from simpler cases.

11.1 Paths, Cycles, Trees, and Complete Graphs

An undirected path of length n is the simplest connected graph with connectivity $\kappa = 1$ and robustness $r = 1$. A cycle is a path of length n whose start and end nodes are connected. For a cycle graph, the connectivity is $\kappa = 2$ and the robustness is $r = 1$. Trees are connected acyclic graphs with connectivity $\kappa = 1$ and robustness $r = 1$. A complete graph is $(n-1)$-connected and $\lceil \frac{n}{2} \rceil$-robust.

11.2 Circulant Networks, k-Nearest Neighbor Paths, and 1-D Random Geometric Graphs

Definition 10 (k-Nearest Neighbor Paths) *A k-nearest neighbor path, $\mathcal{P}(n, k)$, is a network comprised of n nodes in a path, where the nodes are labeled as $1, 2, ..., n$ from one end of the path to the other and each node i can communicate with its k nearest neighbors behind it and k nearest neighbors ahead of it, i.e., $i - k, i - k + 1, ..., i - 1, i + 1, i + 2, ..., i + k$, for some $k \in \mathbb{N}$. An example of $\mathcal{P}(n, k)$ is shown in Fig. 12 (a).*

Such k-nearest neighbor paths are relevant for modelling vehicle platoons, due to the limited sensing and communication range for each vehicle. It is shown in [213] that a k-nearest neighbor path, $\mathcal{P}(n, k)$, is a k-connected graph. We will discuss the robustness of $\mathcal{P}(n, k)$ later in this section. A similar structure to $\mathcal{P}(n, k)$ is a 1-dimensional geometric random graph which captures edges between nodes that are in close (spatial) proximity to each other.

Definition 11 (Geometric Random Graphs) *A geometric random graph $\mathcal{G}_{n,\rho,l}^d = \{\mathcal{V}, \mathcal{E}\}$ is an undirected graph generated by first placing n nodes (according to some mechanism) in a d-dimensional region $\Omega_d = [0, l]^d$, where $d \in \mathbb{Z}_{\geq 1}$. We denote the position of node $i \in \mathcal{V}$ by $x_i \in \Omega_d$. Nodes $i, j \in \mathcal{V}$ are connected by an edge if and only if $\|x_i - x_j\| \leq \rho$ for some threshold ρ, where $\|.\|$ is some appropriate norm (often taken to be the standard*

Euclidean norm). When the node positions are generated randomly (e.g., uniformly and independently) in the region, one obtains a geometric random graph.

In the more general models of $\mathcal{G}^d_{n,\rho,l}$, the length l is also allowed to increase and the density $\frac{n}{l^d}$ can converge to some constant, making it suitable for capturing both dense and sparse random networks. The following result holds for 1-dimensional geometric random graphs.

Proposition 20 *In $\Omega_1 = [0, l]$ with fixed l, if $\mathcal{G}^1_{n,\rho,l}$ is k-connected, then it is at least $\lfloor \frac{k}{2} \rfloor$-robust.*

Based on Definition 11, the k-nearest neighbor path can be seen as a geometric graph $\mathcal{G}^d_{n,\rho,l}$ with $\rho = \frac{lk}{n-1}$ and placing the nodes as follows: the first node, 1, is placed on one end of the line and the i-th node is placed at distance $\frac{(i-1)l}{n-1}$ from 1. Thus, based on Proposition 20 and the fact that the k-nearest neighbor path $\mathcal{P}(n, k)$ is k-connected, we conclude that it is at least $\lfloor \frac{k}{2} \rfloor$-robust.

Definition 12 (Circulant Graphs) *An undirected graph of n nodes is called circulant if the n vertices of the graph can be numbered from 0 to $n-1$ in such a way that if some two vertices numbered x and $(x \pm d)$ mod n are adjacent, then every two vertices numbered z and $(z \pm d)$ mod n are adjacent. A directed graph is circulant (with the above labeling) if some two vertices numbered x and $(x + d)$ mod n are adjacent, then every two vertices numbered z and $(z + d)$ mod n are adjacent.*

Informally speaking, an undirected k-circulant graph is a k-nearest neighbor cycle graph. Thus, with the same reasoning, an undirected k-circulant graph is $2k$-connected and at least $\lfloor \frac{k}{2} \rfloor$-robust. The result is extended to directed circulant graphs where it is shown that it is at least $\lfloor \frac{k+2}{4} \rfloor$-robust [214]. An example of an undirected and a directed circulant graph is shown in Fig. 12 (b,c).

11.3 Formation Graphs on 2-Dimensional Lattice

A specific type of geometric graph is the two dimensional lattice which has been widely used in formation control of autonomous robots [215,216]. A lattice is a set of linear combinations with integer coefficients of the elements of a basis of $\mathbb{R}^2$. The elements of the set are lattice points. Let $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$ be bases of a 2-dimensional lattice with $\|\boldsymbol{v}_1\| = \|\boldsymbol{v}_2\| = \ell$ where ℓ is the lattice length. Every point x on the lattice can be described by $x = a_i \boldsymbol{v}_1 + b_j \boldsymbol{v}_2$ where $a_i, b_j \in \mathbb{Z}$. In a lattice, two nodes i and j are connected if $\|x_i - x_j\| \leq \ell$ where $\|.\|$ is the Euclidean norm. Given a set of nodes $\mathcal{V}$ and a distance ℓ, the graph $\mathcal{G}_\ell = \{\mathcal{V}, \mathcal{E}_\ell\}$ with edge set $\mathcal{E}_\ell \triangleq \{(i,j)|\|x_i - x_j\| \leq \ell\}$ is called the proximity graph of the set $\mathcal{V}$. We describe the communication range by a function $R : \mathbb{Z}_{\geq 1} \to \mathbb{R}$ that maps the number of robots m to a distance where m robots are ensured to be reached. A formation of n robots is said to be connected if its associated proximity graph $\mathcal{G}_\ell$ is connected. An example of a lattice and a connected formation is shown in Fig. 12 (d). It was shown in [215] that in a connected formation of n robots, every robot has at least $1 \leq m \leq n-1$ robots within a distance $m\ell$. Based on this, one can compute a minimum communication range for the robots in a formation to guarantee resilience in the communication network.

Proposition 21 ([215]) *Given a set $\mathcal{V}$ of $4f+1$ nodes in a connected formation, if the communication range of every node satisfies $R \geq 3f\ell$, then the associated graph of the formation is $(2f + 1)$-robust.*

11.4 Random Graphs

A common approach to modelling complex networks is via the framework of *random graphs*, i.e., by drawing a graph from a certain probability distribution over the set of all possible graphs on a given set of nodes. Such random graph models have diverse applications [217, 218], including in modeling cascading failures in large scale systems [219, 220]. Here, we summarize the connectivity and robustness properties of certain commonly studied random graph models.

11.4.1 Erdős-Rényi Random Graphs

An Erdős-Rényi (ER) random graph $\mathcal{G}(n, p)$ is a graph on n nodes, where each edge between two distinct nodes is presented independently with probability p (which could be a function of n). We say that a graph property holds *asymptotically almost surely* if the probability of drawing a graph with that property goes to 1 as $n \to \infty$. The following theorem shows the probability threshold for which a graph $\mathcal{G} \in \mathcal{G}(n, p)$ is r-connected and r-robust.

Theorem 22 ([149]) *For any constant $r \in \mathbb{Z}_{\geq 1}$,*

$$t(n) = \frac{\ln n + (r - 1) \ln \ln n}{n}$$

is a threshold function for the ER random graph $\mathcal{G}$ to have minimum degree r, to be r-connected, and to be r-robust.

According to the example graph in Fig. 3 (a), graph robustness is a much stronger property than the graph connectivity and the minimum degree. However, Theorem 22 indicates that the above threshold function for r-connectivity (and minimum degree r) is also a threshold function for the stronger property of r-robustness in ER random graphs.

11.4.2 Random Regular Graphs

Let $\Omega_{n,d}$ be the set of all undirected graphs on n nodes where every node has degree d (note that this assumes

that nd is even). A *random d-regular graph* (RRG), denoted $\mathcal{G}_{n,d}$ is a graph drawn uniformly at random from $\Omega_{n,d}$. For $d \geq 3$, it is shown that $\mathcal{G}_{n,d}$ is asymptotically almost surely d-connected [217]. Based on [221], for any $\epsilon > 0$, the algebraic connectivity of a random d-regular graph satisfies

$$\lambda_2(L) \geq d - 2\sqrt{d-1} - \epsilon, \tag{47}$$

asymptotically almost surely. As discussed in [216], if the algebraic connectivity of a graph is bigger than $r - 1$, then the network is at least $\lfloor \frac{r}{2} \rfloor$-robust. Hence, according to (47), an RRG is at least $\lfloor \frac{d-2\sqrt{d-1}}{2} \rfloor$-robust asymptotically almost surely.

11.4.3 Random Interdependent Networks

An interdependent network $\mathcal{G}$ is denoted by a tuple $\mathcal{G} = (\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_k, \mathcal{G}_p)$ where $\mathcal{G}_l = (\mathcal{V}_l, \mathcal{E}_l)$ for $l = 1, 2, \ldots, k$ are called the *subnetworks* of the network $\mathcal{G}$, and $\mathcal{G}_p = (\mathcal{V}_1 \cup \mathcal{V}_2 \cup \ldots \cup \mathcal{V}_k, \mathcal{E}_p)$ is a k-partite network with $\mathcal{E}_p \subseteq \cup_{l \neq t} \mathcal{V}_l \times \mathcal{V}_t$ specifying the interconnection (or inter-network) topology. Applications of interdependent networks in modelling communication networks and power grid are discussed in [222]. Define the sample space Ω_n to consist of all possible interdependent networks $(\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_k, \mathcal{G}_p)$ and the index $n \in \mathbb{N}$ denotes the number of nodes in each subnetwork. A random interdependent network is a network $\mathcal{G} = (\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_k, \mathcal{G}_p)$ drawn from Ω_n according to a given probability distribution.

We assume that $|\mathcal{V}_1| = |\mathcal{V}_2| = \cdots = |\mathcal{V}_k| = n$ and that the number of subnetworks k is at least 2. Similar to Theorem 22 for ER random graphs, there exists a sharp threshold for connectivity and robustness of random interdependent networks.

Theorem 23 ([223]) *Consider a random interdependent network $\mathcal{G} = (\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_k, \mathcal{G}_p)$. Then, for any positive integers r and $k \geq 2$,*

$$t(n) = \frac{\ln n + (r-1)\ln \ln n}{(k-1)n}$$

is a threshold for r-connectivity and r-robustness of $\mathcal{G}$.

11.4.4 Random Intersection Networks

Random intersection graphs belong to class of random graphs for which every node is assigned a set of objects selected by some random mechanism. They have applications in wireless sensor networks, frequency hopping spread spectrum, spread of epidemics, and social networks [224].

Given a node set $\mathcal{V} = \{1, 2, ..., n\}$, each node i is assigned an object set S_i from an object pool $\mathcal{P}$ consisting of P_n distinct objects, where P_n is a function of n. Each object S_i is constructed using the following two-step procedure: (i) The size of S_i, $|S_i|$, is determined according to some probability distribution $\mathcal{D} : \{1, 2, ..., P_n\} \to [0, 1]$ in which $\sum_{x=1}^{P_n} \mathbb{P}(|S_i| = x) = 1$. (ii) Conditioning on $|S_i| = s_i$, set S_i is chosen uniformly among all s_i-size subsets of $\mathcal{P}$. Finally, an undirected edge is assigned between two nodes if and only if their corresponding object sets have at least one object in common. There are variations of the general random intersection graph such as binomial random intersection graphs and uniform random intersection graphs, each of which focuses on a certain probability distribution $\mathcal{D}$. The following theorem discusses the connectivity and robustness of random intersection graphs.

Theorem 24 ([225]) *Consider a general random intersection graph $\mathcal{G}(n, P_n, \mathcal{D})$. Let X be a random variable following probability distribution $\mathcal{D}$. With a sequence α_n for all n defined through $\frac{\{\mathbb{E}[X]\}^2}{P_n} = \frac{\ln n + (r-1)\ln \ln n + \alpha_n}{n}$, if $\mathbb{E}[X] = \Omega(\sqrt{\ln n})$, $Var[X] = o(\frac{\{\mathbb{E}[X]\}^2}{n(\ln n)^2})$ and $\alpha_n = o(\ln n)$, and $\lim_{n \to \infty} \alpha_n = \infty$, then the graph is asymptotically almost surely r-connected and r-robust.*

12 Future Directions

This paper provided an overview of the existing graph-theoretic tools which can be used to analyze the resilience of distributed algorithms in networked control systems. Comparing with system-theoretic approaches to the robustness and fault tolerance of control systems, graph-theoretic approaches are relatively new and demand more development, primarily in the following three directions: (i) Developing graph-theoretic methods to facilitate analyzing a wider range of distributed algorithms and more complex adversarial actions. (ii) Reinterpreting the known system-theoretic notions of resilience and robustness of dynamical systems from a graph-theoretic perspective. (iii) Investigating the resilience of a wider range of distributed algorithms using available graph-theoretic tools.

Here, we propose a few research avenues which are worth investigating in the future.

Spectral Approach to Network Structures: One of the necessary steps towards reconciling system-theoretic approaches and graph theory is to find algebraic interpretations of certain network structures. Algebraic graph theory is an active topic of research in mathematics [81]. However, specific structural properties of networks that are widely used in analyzing the resilience of networked control systems, e.g., network robustness, are quite new notions defined within the field of systems and control. Hence, their algebraic interpretations have not been well studied.

An example is a work on the relation of the algebraic connectivity and the network robustness. Defining the *edge-boundary* of a set of nodes $S \subset \mathcal{V}$ is given by $\partial S = \{(i,j) \in \mathcal{E} \mid i \in S, j \in \mathcal{V} \setminus S\}$. The *isoperimetric constant* of $\mathcal{G}$ is defined as [226]

$$i(\mathcal{G}) \triangleq \min_{S \subset \mathcal{V}, |S| \le \frac{n}{2}} \frac{|\partial S|}{|S|}. \qquad (48)$$

Based on the above definition and the definition of the network robustness we conclude that if $i(\mathcal{G}) > r - 1$, then the graph is at least r-robust. Moreover, we have $\lambda_2(L) \le 2i(\mathcal{G})$ [226]. Based on this, if $\lambda_2(L) > r - 1$, then the network is at least $\lfloor \frac{r}{2} \rfloor$-robust [216, 223]. However, $\lfloor \frac{r}{2} \rfloor$ provides a loose lower bound for the network robustness. An example is a star graph which is 1-robust with $\lambda_2(L) = 1$. Further research is needed to be done in this direction.

Resilience with Minimum Communication: Due to the fact that communications between nodes can be costly in many applications, the problem of reaching a certain level of resilience with minimum communication, i.e., edges between nodes, is worth investigating. In specific dynamical systems, e.g., consensus dynamics, adding edges may also degrade the controllability of the system. Few recent works have focused on minimizing the number of edges while reaching a certain level of security [227] or maximizing the connectivity without violating the controllability of the system [228]. While reaching a certain level of connectivity with the minimum number of edges has been well-studied in the literature, the same problem for network robustness remains open.

Graph-Theoretic Approach to Attack Energy: In Section 9 some graph-theoretic interpretations of the attack's impact was discussed. Among those, the attack energy has not been well studied. The objective of the attacker, other than impact and detectability, can be to access the system with minimum energy. One way to quantify the attack energy is via using the spectra of the controllability Gramian $\mathcal{W}_{\mathcal{F}}$. An interesting research avenue is to design the network to maximize the attack energy using an appropriate spectrum of the controllability Gramian.

Resilience of other Classes of Random Graphs: The study of the connectivity and network robustness of random graphs with the various probability distribution of edge formation is another important future research line. Among them are scale-free networks in which their degree distributions follow a power-law. Many networks have been reported to be scale-free, including preferential attachment and the fitness model [229]. Analyzing the structural properties of these classes of networks enables us to quantify the resilience of distributed algorithms on various natural and man-made large-scale systems.

Resilience of Networks with Nonlinear Interactions: Throughout this survey paper, the focus was on NCSs in which the interactions between agents are linear. In some cases of NCSs, the local interactions are nonlinear, e.g., synchronization of Kuramoto Oscillators with applications to power systems [230] and attraction-repulsion functions in swarm robotics [231]. Extension of the methods discussed in this survey to those classes of nonlinear systems requires further investigations.

Resilience of Multi-Agent Reinforcement Learning Algorithms: Another emerging area of interest pertains to enabling a team of agents to cooperatively learn optimal policies for interacting with their environment, particularly when the dynamics of the environment are not initially known to the agents. Such *multi-agent reinforcement learning* settings share some common features with the distributed consensus, optimization, and estimation problems we described earlier in the paper, in that adversarial agents can provide incorrect information about their observations and costs to the other agents. However, these problems also introduce additional lines of complexity for resilience, in that adversarial agents can also affect the underlying system with their inputs. There have been recent initial explorations of resilient algorithms in such settings [232–235], but much work remains to be done to understand how to mitigate adversaries that can not only send incorrect information but can also take destructive actions.

References

[1] M. J. Newman. Clustering and preferential attachment in growing networks. *Physical review E*, 64(2):025102, 2001.

[2] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[3] R. Albert, H. Jeong, and A. L. Barabási. Error and attack tolerance of complex networks. *Nature*, 406:378–482, 2000.

[4] Z. Biron, S. Dey, and P. Pisu. Real-time detection and estimation of denial of service attack in connected vehicle systems. *IEEE Transaction on Intelligent Transportation Systems*, 19:3893–3902, 2018.

[5] L. Mokdad and A. T. Nguyen J. Ben-Othman. Djavan: Detecting jamming attacks in vehicle ad hoc networks. *Performance Evaluation*, 87:47–59, 2015.

[6] K. Koscher, A. Czeskis, F. Roesner, S. Patel, Tadayoshi Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage. Experimental security analysis of a modern automobile. *Security and Privacy (SP), 2010 IEEE Symposium on*, pages 447–462, 2010.

[7] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive experimental analyses of automotive attack surfaces. *USENIX Security Symposium*, pages 447–462, 2011.

[8] J. A. Stankovic. Research directions for the internet of things. *IEEE Internet of Things Journal*, pages 3–9, 2014.

[9] A. Hooshyar and R. Iravani. Microgrid protection. *Proceedings of the IEEE*, 105:1332–1353, 2017.

[10] S. Sridhar, A. Hahn, and M. Govindarasu. Cyber–physical system security for the electric power grid. *Proceedings of the IEEE*, 100:210–224, 2011.

[11] T. Huang, B. Satchidanandan, P. R. Kumar, and L. Xie. An online detection framework for cyber attacks on automatic generation control. *IEEE Transactions on Power Systems*, 33:6816–6827, 2018.

[12] Z. Li, M. Shahidehpour, and F. Aminifar. Cybersecurity in distributed power systems. ” *Proceedings of the IEEE*, 105:1367–1388, 2017.

[13] L. Xie, Y. Mo, and B. Sinopoli. Integrity data attacks in power market operations. *IEEE Transactions on Smart Grid*, 2:659–666, 2011.

[14] C.W. Ten and G. Manimaran C. C. Liu. Vulnerability assessment of cybersecurity for scada systems. In *IEEE Transactions on Power Delivery*, volume 23, pages 1836–1846, 2008.

[15] G. N. Ericsson. Cyber security and power system communication—essential parts of a smart grid infrastructure. In *IEEE Transactions on Power Delivery*, volume 25, pages 1501–1507, 2010.

[16] I¿ C. Lin and T. C. Liao. A survey of blockchain security issues and challenges. *International Journal of Network Security*, 19:653–659, 2017.

[17] A. Kushleyev, D. Mellinger andC. Powers, and V. Kumar. Towards a swarm of agile micro quadrotors. *Autonomous Robots*, 35:287–300, 2013.

[18] H. Hamann. Swarm robotics: A formal approach. *New York: Springer International Publishing*, 2018.

[19] Jill Slay and Michael Miller. Lessons learned from the maroochy water breach. In *Proceedings of International Conference on Critical Infrastructure Protection*, pages 73–82, 2007.

[20] N. Falliere, L. Murchu, and E. Chien. W32. stuxnet dossier: Symantec security response. *https://www.symantec.com*, 2013.

[21] P. V. Marsden and N. E. Friedkin. *Network studies of social influence.* Sociological Methods & Research, 1993.

[22] M. Jackson. *Social and Economic Networks.* Princeton University Press, 2010.

[23] S. M. Dibaji, M. Pirani, D. Flamholz, A. M. Annaswamy, K. H. Johansson, and A. Chakrabortty. A systems and control perspective of CPS security. *Annual Reviews in Control*, 2019.

[24] M. A. Massoumnia, G. C. Verghese, and A. S. Willsky. Failure detection and identification. *IEEE Transactions on automatic control*, 34:316–321, 1989.

[25] Y. Guan and M. Saif. A novel approach to the design of unknown input observers. *IEEE Transactions on automatic control*, 36:632–635, 1991.

[26] M. Darouach, M. Zasadzinski, and S. J. Xu. Full-order observers for linear systems with unknown inputs. *IEEE Transactions on automatic control*, 39:606–609, 1994.

[27] C. De Persis and A. Isidori. A geometric approach to nonlinear fault detection and isolation. *IEEE transactions on automatic control*, 46:853–865, 2001.

[28] Andre Teixeira, Kin Cheong Sou, Henrik Sandberg, and Karl Henrik Johansson. Secure control systems: A quantitative risk management approach. *IEEE Control Systems Magazine*, 35(1):24–45, 2015.

[29] R. Koetter and M. Médard. An algebraic approach to network coding. *IEEE/ACM transactions on networking*, 11:782–795, 2003.

[30] F. Farokhi, I. Shames, and N. Batterham. Secure and private control using semi-homomorphic encryption. *Control Engineering Practice*, 67:13–20, 2017.

[31] C. Dwork. Differential privacy: A survey of results. *International conference on theory and applications of models of computation*, pages 1–19, 2008.

[32] J. L. Ny and G. J. Pappas. Differentially private filtering. *IEEE Transactions on Automatic Control*, 59:341–354, 2014.

[33] R Motwani and P Raghavan. Randomized algorithms. *ACM Computing Surveys*, 28, 1996.

[34] S. Weerakkody and B. Sinopoli. Detecting integrity attacks on control systems using a moving target approach. In *IEEE Conference on Decision and Control*, pages 5820–5826, 2015.

[35] F. Pasqualetti, F. Dorfler, and F. Bullo. Attack detection and identification in cyber-physical systems. *IEEE transactions on automatic control*, 58:2715–2729, 2013.

[36] R. E. Blahut. Algebraic codes for data transmission. *Cambridge, UK: Cambridge University Press*, 2003.

[37] Y. Mo, S. Weerakkody, and B. Sinopoli. Physical authentication of control systems: Designing watermarked control inputs to detect counterfeit sensor outputs. In *IEEE Control Systems Magazine*, volume 35, pages 93–109, 2015.

[38] P. Garcia-Teodoro, J. Diaz-Verdejo, G. Macia-Fernandez, and E. Vazquez. Anomaly-based network intrusion detection: ´ techniques, systems and challenges. In *Computers & Security*, volume 28, pages 18–28, 2009.

[39] C.-F. Tsai, Y.-F. Hsu, C.-Y. Lin, and W.-Y. Lin. Anomaly-based network intrusion detection: ´ techniques, systems and challenges. In *Intrusion detection by machine learning: A review*, volume 36, page 11994–12000, 2009.

[40] Y. He, G. J. Mendis, and J. Wei. Real-time detection of false data injection attacks in smart grids: A deep learningbased intelligent mechanism. In *IEEE Transactions on Smart Grid*, volume 8, pages 1–12, 2016.

[41] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I Jordan, and S. S Sastry. Kalman filtering with intermittent observations. *IEEE Transactions on Automatic Control*, 49:1453–1464, 2004.

[42] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu. A survey of recent results in networked control systems. *Proceedings of the IEEE*, 95(1):138–162, 2007.

[43] R. Touri C. N. Hadjicostis. Feedback control utilizing packet dropping network links. *Proceedings of the 41st IEEE Conference on Decision and Control*, pages 1205–1210, 2002.

[44] H. Fawzi, P. Tabuada, and S. Diggavi. Secure estimation and control for cyber-physical systems under adversarial attacks. *IEEE Transactions on Automatic Control*, 59:1454–1467, 2014.

[45] E. Y. Chow and A. S. Willsky. Analytical redundancy and the design of robust failure detection systems. In *IEEE Transactions on Automatic Control*, volume 29, pages 603–614, 1984.

[46] J. R. Sklaroff. Redundancy management technique for space shuttle computers. In *IBM Journal of Research and Development*, volume 20, pages 20–28, 1976.

[47] W. Abbas, A. Laszka, and X. Koutsoukos. Improving network connectivity and robustness using trusted nodes with application to resilient consensus. In *IEEE Transactions on Control of Network Systems*, volume 5, pages 2036–2048, 2017.

[48] M. Momani and S. Challa. Survey of trust models in different network domains. In *arXiv preprint arXiv:1010.0168*, 2010.

[49] W. Heemels, K. H. Johansson, and P. Tabuada. An introduction to event-triggered and self-triggered control. In *in Proceedings of IEEE Conference on Decision and Control*, pages 3270–3285, 2012.

[50] C. De Persis and P. Tesi. On resilient control of nonlinear systems under denial-of-service. In *in Proceedings of IEEE Conference on Decision and Control*, pages 5254–5259, 2014.

[51] Q. Zhu and T. Basar. Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: Games-in-games principle for optimal cross-layer resilient control systems. In *IEEE control systems*, volume 35, pages 45–65, 2015.

[52] D. P. Bertsekas and J. N. Tsitsiklis. Parallel and distributed computation: numerical methods. *Prentice hall*, 1989.

[53] J. Tsitsiklis. Problems in decentralized decision making and computation. *Massachusetts Inst of Tech Cambridge Lab for Information and Decision Systems*, 1984.

[54] N. Sandell, P. Varaiya, M. Athans, and M. Safonov. Survey of decentralized control methods for large scale systems. *IEEE Transactions on automatic Control*, 23:108–128, 1978.

[55] V. Borkar and P. Varaiya. Asymptotic agreement in distributed estimation. *IEEE Transactions on Automatic Control*, 27:650–655, 1982.

[56] D. Teneketzis and P. Varaiya. Consensus in distributed estimation with inconsistent beliefs. *Systems & control letters*, 4:217–221, 1984.

[57] H. Witsenhausen. A counterexample in stochastic optimum control. *SIAM J. Control*, 6:131–147, 1968.

[58] A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48:988–1001, 2003.

[59] W. Ren and R. W. Beard. Consensus seeking in multi agent systems under dynamically changing interaction topologies. *IEEE Trans. Autom. Control*, 50:655–661, 2005.

[60] C. Tomlin, G. J. Pappas, and S. Sastry. Conflict resolution for air traffic management: A study in multiagent hybrid systems. *IEEE Transactions on automatic control*, 43:509–521, 1998.

[61] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *IEEE Transactions on Automatic Control*, 95:215–233, 2007.

[62] N. A. Lynch. Distributed algorithms. *San Francisco, California: Morgan Kaufmann Publishers Inc*, 1997.

[63] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32:374–382, 1985.

[64] S. Sundaram and C. N. Hadjicostis. Distributed function calculation via linear iterative strategies in the presence of malicious agents. *IEEE Transactions on Automatic Control*, 56(7):1495–1508, 2011.

[65] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram. Resilient asymptotic consensus in robust networks. *IEEE Journal on Selected Areas in Communications*, 31:766–781, 2013.

[66] F. Pasqualetti, A. Bicchi, and F. Bullo. Consensus computation in unreliable networks: A system theoretic approach. *IEEE Transactions on Automatic Control*, 57:90–104, 2012.

[67] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.

[68] Z. Jiang, A. Balu, C. Hegde, and S. Sarkar. Collaborative deep learning in fixed topology networks. *Advances in Neural Information Processing Systems*, pages 5904–5914, 2017.

[69] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10:1–19, 2019.

[70] Luis Muñoz-González, Kenneth T Co, and Emil C Lupu. Byzantine-robust federated machine learning through adaptive model averaging. *arXiv preprint arXiv:1909.05125*, 2019.

[71] Avishek Ghosh, Justin Hong, Dong Yin, and Kannan Ramchandran. Robust federated learning in a heterogeneous environment. *arXiv preprint arXiv:1906.06629*, 2019.

[72] Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, 70:1142–1154, 2022.

[73] Jairo Giraldo, Esha Sarkar, Alvaro A Cardenas, Michail Maniatakos, and Murat Kantarcioglu. Security and privacy in cyber-physical systems: A survey of surveys. *IEEE Design & Test*, 34(4):7–17, 2017.

[74] H. Ishii, Y. Wang, and S. Feng. An overview on multi-agent consensus under adversarial attacks. *Annual Reviews in Control*, 2022.

[75] Zhixiong Yang, Arpita Gang, and Waheed U Bajwa. Adversary-resilient distributed and decentralized statistical inference and machine learning: An overview of recent advances under the byzantine threat model. *IEEE Signal Processing Magazine*, 37(3):146–159, 2020.

[76] A. Prorok, M. Malencia, L. Carlone, G. S. Sukhatme, B. M. Sadler, and V. Kumar. Beyond robustness: A taxonomy of approaches towards resilient multi-robot systems. *arXiv preprint arXiv:2109.12343*, 2021.

[77] D. B. West. *Introduction to Graph Theory*. Prentice-Hall Inc., Upper Saddle River, New Jersey, 2001.

[78] Haotian Zhang and Shreyas Sundaram. Robustness of information diffusion algorithms to locally bounded adversaries. In *2012 American Control Conference (ACC)*, pages 5855–5861. IEEE, 2012.

[79] W. Chen, J. Liu, Y. Chen, S. Zhen Khong, D. Wang, T. Başar, L. Qiu, and K. H. Johansson. Characterizing the positive semidefiniteness of signed laplacians via effective resistances. *IEEE 55th Conference on Decision and Control (CDC)*, pages 985–990, 2016.

[80] Y Hou, J. Li, and Pan. On the laplacian eigenvalues of signed graphs. *Linear and Multilinear Algebra*, 51:21–30, 2003.

[81] C. Godsil and G. Royle. *Algebraic Graph Theory*. Springer, 2001.

[82] L. Farina and S. Rinaldi. Positive linear systems: theory and applications. *John Wiley & Sons*, 2000.

[83] J. M. Dion, C. Commault, and J. van der woude. Generic properties and control of linear structured systems: a survey. *Automatica*, 39(7):1125–1144, 2003.

[84] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst*, 4(2):382–401, 1982.

[85] Z. Bouzid, M. G. Potop-Butucaru, and S. Tixeuil. Optimal byzantineresilient convergence in uni-dimensional robot networks. *Theoretical Computer Science*, 411(34):3154–3168, 2010.

[86] D. B. West. Introduction to graph theory. *Upper Saddle River, NJ: Prentice-Hall*, 2001.

[87] S. Sundaram and C. N. Hadjicostis. Distributed function calculation and consensus using linear iterative strategies. *IEEE Journal on Selected Areas in Communications*, 26:650–660, 2008.

[88] D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *J. Assoc. Comput. Mach*, 40(1):17–47, 1993.

[89] R. Olfati-Saber and R. M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49:1520–1533, 2004.

[90] S. Sundaram and C. N. Hadjicostis. Distributed function calculation and consensus using linear iterative strategies. *IEEE Journal on Selected Areas in Communications*, 26:650–660, 2008.

[91] M. Pirani, E. M. Shahrivar, and S. Sundaram. Coherence and convergence rate in networked dynamical systems. *Proceedings of CDC 2015, the 54th IEEE Conference on Decision and Control*, pages 968–973, 2015.

[92] J. Ghaderi and R. Srikant. Opinion dynamics in social networks: A local interaction game with stubborn agents. *Automatica*, 50:3209 – 3215, 2014.

[93] E. Yildiz, D. Acemoglu, A. Ozdaglar, A. Saberi, and A. Scaglione. Discrete opinion dynamics with stubborn agents. *Operations Research*, 2011.

[94] M. Pirani and S. Sundaram. Spectral properties of the grounded Laplacian matrix with applications to consensus in the presence of stubborn agents. *American Control Conference*, pages 2160 – 2165, 2014.

[95] D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl. Reaching approximate agreement in the presence of faults. *J. Assoc. Comput. Mach*, 33:499–516, 1986.

[96] Peter J Huber. The 1972 wald lecture robust statistics: A review. *The Annals of Mathematical Statistics*, 43(4):1041–1067, 1972.

[97] R. M. Kieckhafer and M. H. Azadmanesh. Reaching approximate agreement with mixed mode faults. *IEEE Transactions on Parallel and Distributed Systems*, 5:53–63, 1994.

[98] H. LeBlanc and X. Koutsoukos. Consensus in networked multi-agent systems with adversaries. *Proceedings of the 14th international conference on Hybrid systems: computation and control, ser. HSCC*, pages 281–290, 2011.

[99] S.M. Dibaji and H. Ishii. Resilient consensus of second-order agent networks: Asynchronous update rules with delays. *Automatica*, 81:123–132, 2017.

[100] S.M.Dibaji, H. Ishii, and R. Tempo. Resilient randomized quantized consensus. *IEEE Transactions on Automatic Control*, 63:2508–2522, 2018.

[101] K. Saulnier, D. Saldana, A. Prorok, G. J. Pappas, and V. Kumar. Resilient flocking for mobile robot teams. *IEEE Robotics and Automation Letters*, 2:1039–1046, 2017.

[102] S. Gil, C. Baykal, and D. Rus. Resilient multi-agent consensus using wi-fi signals. *IEEE Control Systems Letters*, 3:126–131, 2019.

[103] D. Saldana, A. Prorok, S. Sundaram, M. F. M. Campos, and V. Kumar. Resilient consensus for time-varying networks of dynamic agents. *American Control Conference*, 65:252–258, 2017.

[104] N. H. Vaidya. Iterative byzantine vector consensus in incomplete graphs. *Distributed Computing and Networking*, pages 14–28, 2014.

[105] Z. Xiang and N. H. Vaidya. Brief announcement: Relaxed byzantine vector consensus. *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 401–403, 2016.

[106] H. Park and S. A. Hutchinson. Fault-tolerant rendezvous of multirobot systems. *IEEE Transactions on Robotics*, 33:565–582, 2017.

[107] J. R. Reay. An extension of radon's theorem. *Illinois J. Math*, 12:184–189, 1968.

[108] W. Abbas, M. Shabbir, J. Li, and X. Koutsoukos. Interplay between resilience and accuracy in resilient vector consensus in multi-agent networks. *arXiv:2009.02417v1*, 2020.

[109] I. Bárány and P. Soberón. A generalization of radon's theorem. *Journal of the London Mathematical Society*, 41:123–128, 1966.

[110] I. Bárány and P. Soberón. A "safe kernel" approach for resilient multi-dimensional consensus. *21st IFAC World Congress*, 55:2507–2512, 2020.

[111] W. Mulzer and D. Werner. Approximating tverberg points in linear time for any fixed dimension. *Discrete Comput Geom*, 50:520–535, 2013.

[112] M. Shabbir, J. Li, W. Abbas, and X. Koutsoukos. Resilient vector consensus in multi-agent networks using centerpoints. *American Control Conference*, 33:4387–4392, 2020.

[113] J. Matousek. Lectures on discrete geometry. *Springer*, 2002.

[114] A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54:48–61, 2009.

[115] N. Ravi, A. Scaglione, and A. Nedic. A case of distributed optimization in adversarial environment. *IEEE International Conference on Acoustics, Speech and Signal Processing*, page 5252–5256, 2019.

[116] S. Sundaram and B. Gharesifard. Distributed optimization under adversarial nodes. *IEEE Transactions on Automatic Control*, 64:1063–1076, 2018.

[117] L. Su and N. H. Vaidya. Byzantine-resilient multiagent optimization. *IEEE Transactions on Automatic Control*, 66:2227–2233, 2021.

[118] Kananart Kuwaranancharoen and Shreyas Sundaram. On the location of the minimizer of the sum of two strongly convex functions. In *IEEE Conference on Decision and Control (CDC)*, pages 1769–1774, 2018.

[119] Z. Yang and W. U. Bajwa. Byrdie: Byzantine-resilient distributed coordinate descent for decentralized learning. *IEEE Transactions on Signal and Information Processing over Networks*, 2019.

[120] Kananart Kuwaranancharoen, Lei Xin, and Shreyas Sundaram. Byzantine-resilient distributed optimization of multi-dimensional functions. In *2020 American Control Conference (ACC)*, pages 4399–4404. IEEE, 2020.

[121] Nirupam Gupta, Thinh T Doan, and Nitin H Vaidya. Byzantine fault-tolerance in decentralized optimization under 2f-redundancy. In *2021 American Control Conference (ACC)*, pages 3632–3637. IEEE, 2021.

[122] Nirupam Gupta and Nitin H Vaidya. Fault-tolerance in distributed optimization: The case of redundancy. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 365–374, 2020.

[123] Soummya Kar and José MF Moura. Convergence rate analysis of distributed gossip (linear parameter) estimation: Fundamental limits and tradeoffs. *IEEE Journal of Selected Topics in Signal Processing*, 5(4):674–690, 2011.

[124] Soummya Kar, José MF Moura, and Kavita Ramanan. Distributed parameter estimation in sensor networks: Nonlinear observation models and imperfect communication. *IEEE Transactions on Information Theory*, 58(6):3575–3605, 2012.

[125] Soummya Kar and Jose MF Moura. Consensus+ innovations distributed inference over networks: cooperation and sensing in networked systems. *IEEE Signal Processing Magazine*, 30(3):99–109, 2013.

[126] Le Xie, Dae-Hyun Choi, Soummya Kar, and H Vincent Poor. Fully distributed state estimation for wide-area monitoring systems. *IEEE Transactions on Smart Grid*, 3(3):1154–1169, 2012.

[127] Yuan Chen, Soummya Kar, and Jose MF Moura. Resilient distributed estimation through adversary detection. *IEEE Transactions on Signal Processing*, 66(9):2455–2469, 2018.

[128] Xingkang He, Xiaoqiang Ren, Henrik Sandberg, and Karl H Johansson. How to secure distributed filters under sensor attacks. *IEEE Transactions on Automatic Control*, 2021.

[129] Yuan Chen, Soummya Kar, and José MF Moura. Resilient distributed estimation: Sensor attacks. *IEEE Transactions on Automatic Control*, 64(9):3772–3779, 2018.

[130] Yuan Chen, Soummya Kar, and JM Moura. Topology free resilient distributed estimation. *arXiv preprint arXiv:1812.08902*, 2018.

[131] Yuan Chen, Soummya Kar, and José MF Moura. Resilient distributed parameter estimation with heterogeneous data. *IEEE Transactions on Signal Processing*, 67(19):4918–4933, 2019.

[132] Yuan Chen, Soummya Kar, and José MF Moura. Resilient distributed field estimation. *SIAM Journal on Control and Optimization*, 58(3):1429–1456, 2020.

[133] L. Su and S. Shahrampour. Finite-time guarantees for byzantine-resilient distributed state estimation with noisy measurements. *IEEE Transactions on Automatic Control*, 65:3758–3771, 2020.

[134] Nitin H Vaidya, Lewis Tseng, and Guanfeng Liang. Iterative approximate byzantine consensus in arbitrary directed graphs. In *Proceedings of the 2012 ACM symposium on Principles of distributed computing*, pages 365–374, 2012.

[135] Heath J LeBlanc and Firas Hassan. Resilient distributed parameter estimation in heterogeneous time-varying networks. In *Proceedings of the 3rd international conference on High confidence networked systems*, pages 19–28, 2014.

[136] Liwei An and Guang-Hong Yang. Byzantine-resilient distributed state estimation: A min-switching approach. *Automatica*, 129:109664, 2021.

[137] A. Mitra and S. Sundaram. Byzantine-resilient distributed observers for lti systems. *Automatica*, 108, 2019.

[138] Shinkyu Park and Nuno C Martins. Design of distributed LTI observers for state omniscience. *IEEE Trans. on Autom. Control*, 62(2):561–576, 2017.

[139] Aritra Mitra and Shreyas Sundaram. Distributed observers for LTI systems. *IEEE Trans. on Autom. Control*, 63(11):3689–3704, 2018.

[140] Lili Wang and A Stephen Morse. A distributed observer for a time-invariant linear system. *IEEE Trans. on Autom. Control*, 63(7), 2018.

[141] Weixin Han, Harry L Trentelman, Zhenhua Wang, and Yi Shen. A simple approach to distributed observer design for linear systems. *IEEE Trans. on Autom. Control*, 64(1):329–336, 2019.

[142] Francisco FC Rego, A Pedro Aguiar, António M Pascoal, and Colin N Jones. A design method for distributed Luenberger observers. In *Proc. of the 56th IEEE Conference on Decision and Control*, pages 3374 – 3379, 2017.

[143] Álvaro Rodríguez del Nozal, Pablo Millán, Luis Orihuela, Alexandre Seuret, and Luca Zaccarian. Distributed estimation based on multi-hop subspace decomposition. *Automatica*, 99:213–220, 2019.

[144] Taekyoo Kim, Chanhwa Lee, and Hyungbo Shim. Completely decentralized design of distributed observer for linear systems. *IEEE Trans. on Autom. Control*, 2019.

[145] Michelle S Chong, Masashi Wakaiki, and Joao P Hespanha. Observability of linear systems under adversarial attacks. In *2015 American Control Conference (ACC)*, pages 2439–2444. IEEE, 2015.

[146] Aritra Mitra, John A Richards, Saurabh Bagchi, and Shreyas Sundaram. Resilient distributed state estimation with mobile agents: overcoming byzantine adversaries, communication losses, and intermittent measurements. *Autonomous Robots*, 43(3):743–768, 2019.

[147] Aritra Mitra, John A Richards, Saurabh Bagchi, and Shreyas Sundaram. Distributed state estimation over time-varying graphs: Exploiting the age-of-information. *arXiv preprint arXiv:2001.07006*, 2020.

[148] Aritra Mitra, Faiq Ghawash, Shreyas Sundaram, and Waseem Abbas. On the impacts of redundancy, diversity, and trust in resilient distributed state estimation. *IEEE Transactions on Control of Network Systems*, 2021.

[149] H. Zhang, E. Fata, and S. Sundaram. A notion of robustness in complex networks. *Control of Network Systems, IEEE Transactions on*, 2:310–320, 2015.

[150] M. Deghat, V. Ugrinovskii, I. Shames, and C. Langbort. Detection and mitigation of biasing attacks on distributed estimation networks. *Automatica*, 99:369–381, 2019.

[151] Valery Ugrinovskii. Distributed robust filtering with h_∞ consensus of estimates. *Automatica*, 47(1):1–13, 2011.

[152] Valery Ugrinovskii and Cédric Langbort. Distributed h_∞ consensus-based estimation of uncertain systems via dissipativity theory. *IET Control Theory & Applications*, 5(12):1458–1469, 2011.

[153] V. Ugrinovskii. Distributed H_∞ estimation resilient to biasing attacks. *IEEE Transactions on Control of Network Systems*, 2019.

[154] Yanwen Mao and Paulo Tabuada. Decentralized resilient state-tracking. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 3480–3485. IEEE, 2021.

[155] Yanwen Mao, Suhas Diggavi, Christina Fragouli, and Paulo Tabuada. Secure state-reconstruction over networks subject to attacks. *IEEE Control Systems Letters*, 5(1):157–162, 2020.

[156] Ali Jadbabaie, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi. Non-Bayesian social learning. *Games and Economic Behavior*, 76(1):210–225, 2012.

[157] Ali Jadbabaie, Pooya Molavi, and Alireza Tahbaz-Salehi. Information heterogeneity and the speed of learning in social

networks. *Columbia Business School Research Paper*, pages 13–28, 2013.

[158] Qipeng Liu, Aili Fang, Lin Wang, and Xiaofan Wang. Social learning with time-varying weights. *Journal of Systems Science and Complexity*, 27(3):581–593, 2014.

[159] Shahin Shahrampour, Alexander Rakhlin, and Ali Jadbabaie. Distributed detection: Finite-time analysis and impact of network topology. *IEEE Transactions on Automatic Control*, 61(11):3256–3268, 2016.

[160] Angelia Nedić, Alex Olshevsky, and César A Uribe. Fast convergence rates for distributed Non-Bayesian learning. *IEEE Transactions on Automatic Control*, 62(11):5538–5553, 2017.

[161] Anusha Lalitha, Tara Javidi, and Anand Sarwate. Social learning and distributed hypothesis testing. *IEEE Transactions on Information Theory*, 64(9), 2018.

[162] Lili Su and Nitin H Vaidya. Defending Non-Bayesian learning against adversarial attacks. *Distributed Computing*, pages 1–13, 2016.

[163] César A Uribe, James Z Hare, Lance Kaplan, and Ali Jadbabaie. Non-Bayesian social learning with uncertain models over time-varying directed graphs. In *Proc. of the 58th IEEE Conference on Decision and Control*, pages 3635–3640, 2019.

[164] Aritra Mitra, John A. Richards, and Shreyas Sundaram. A new approach for distributed hypothesis testing with extensions to Byzantine-resilience. In *Proc. of the American Control Conference*, pages 261–266, 2019.

[165] Aritra Mitra, John A Richards, and Shreyas Sundaram. A new approach to distributed hypothesis testing and non-Bayesian learning: Improved learning rate and Byzantine-resilience. *IEEE Transactions on Automatic Control*, 2020.

[166] Aritra Mitra, Hamed Hassani, and George Pappas. Exploiting heterogeneity in robust federated best-arm identification. *arXiv preprint arXiv:2109.05700*, 2021.

[167] Bo Wu, Steven Carr, Suda Bharadwaj, Zhe Xu, and Ufuk Topcu. Resilient distributed hypothesis testing with time-varying network topology. In *2020 American Control Conference (ACC)*, pages 1483–1488. IEEE, 2020.

[168] J. M. Dion, C. Commault, and J. van der woude. The generic number of invariant zeros of a structured linear system. *SIAM J. Control and Optimiz.*, 38(1):1–21, 1999.

[169] Stephanie Gil, Swarun Kumar, Mark Mazumder, Dina Katabi, and Daniela Rus. Guaranteeing spoof-resilient multi-robot networks. *Autonomous Robots*, 41(6):1383–1400, 2017.

[170] Venkatraman Renganathan and Tyler Summers. Spoof resilient coordination for distributed multi-robot systems. In *2017 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, pages 135–141. IEEE, 2017.

[171] A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt. Controllability of multi-agent systems from a graph-theoretic perspective. *SIAM Journal on Control and Optimization*, 48:162–186, 2009.

[172] A. Olshevsky. Minimal controllability problems. *IEEE Transactions on Control of Network Systems*, pages 249–258, 2014.

[173] S. Sundaram and C. Hadjicostis. Structural controllability and observability of linear systems over finite fields with applications to multi-agent systems. *Automatic Control, IEEE Transactions on*, 58:60–73, 2013.

[174] S. Pequito, S. Kar, and A. P. Aguiar. A framework for structural input/output and control configuration selection in large-scale systems. *IEEE Transactions on Automatic Control*, 61:303–318, 2016.

[175] S. Pequito, S. Kar, and A. P. Aguiar. On the complexity of the constrained input selection problem for structural linear systems. *Automatica*, 62:193–199, 2015.

[176] S. Aji. Estimating reachability set sizes in dynamic graphs. *Virginia Polytechnic Institute and State University*, 2014.

[177] A. Y. Yazıcıoglu, W. Abbas, and M. Egerstedt. Graph distances and controllability of networks. *IEEE Transactions on Automatic Control*, 61:4125–4130, 2016.

[178] R. E. Kalman, Y. C. Ho, and S. K. Narendra. Controllability of linear dynamical systems. *Contrib. Differ. Equ*, 1:189–213, 1963.

[179] T. Kailath. Linear systems. *Linear Systems. Englewood Cliffs, NJ, USA: Prentice-Hall*, 1980.

[180] J. Sun and A. E. Motter. Controllability transition and nonlocality in network control. *Contrib. Differ. Equ*, 1:189–213, 1963.

[181] F. Pasqualetti, S. Zampieri, and F. Bullo. Controllability metrics, limitations and algorithms for complex networks. *Control of Network Systems, IEEE Transactions on*, 1:40–52., 2014.

[182] T. H. Summers, F. L. Cortesi, and J. Lygeros. On submodularity and controllability in complex dynamical networks. *IEEE Transactions on Control of Network Systems*, 1:91–101, 2015.

[183] A. Olshevsky. On (non) supermodularity of average control energy. *IEEE Transactions on Control of Network Systems*, 5:1177–1181, 2017.

[184] A. Clark, B. Alomair, L. Bushnell, and R. Poovendran. *Submodularity in Dynamics and Control of Networked Systems.* Springer, 2016.

[185] T. Summers, I. Shames, J. Lygeros, and F. Dorfler. Topology design for optimal network coherence. *arXiv preprint arXiv:1411.4884*, 2014.

[186] M. Pirani and J. Taylor. Controllability of ac power networks with dc lines. *IEEE Transactions on Power Systems*, 2020.

[187] J. Wei and A. J van der Schaft. Load balancing of dynamical distribution networks with flow constraints and unknown in/outflows. *Systems & Control Letters*, 62:1001–1008, 2013.

[188] A. Ghosh, S. Boyd, and A. Saberi. Minimizing effective resistance of a graph. *SIAM review*, 50:37–66, 2008.

[189] M. Pirani, E. Moradi Shahrivar, B. Fidan, and S. Sundaram. Robustness of leader - follower networked dynamical systems. *IEEE Transaction on Control of Network Systems*, 5(4):1752–1763, 2018.

[190] H. Hao and P. Barooah. Stability and robustness of large platoons of vehicles with double-integrator models and nearest neighbor interaction. *International Journal of Robust and Nonlinear Control*, 23:2097–2122, 2013.

[191] K. Tan and K. M. Grigoriadis. Stabilization and H_∞ control of symmetric systems: an explicit solution. *Systems & Control Letters*, 44:57–72, 2001.

[192] X. Liu, S. Pequito, S. Kar Y. Mo, B. Sinopoli, and A. Pedro Aguiar. Minimum robust sensor placement for large scale linear time-invariant systems: A structured systems approach. *4th IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pages 417–424, 2013.

[193] S. Pequito, G. Ramos, S. Kar, A P. Aguiar, and J. Ramos. The robust minimal controllability problem. *Automatica*, 82:261–268, 2017.

[194] H. Zhang, R. Ayoub, and S. Sundaram. Sensor selection for kalman filtering of linear dynamical systems: Complexity, limitations and greedy algorithms. *Automatica*, 78:202–210, 2017.

[195] Lintao Ye, Nathaniel Woodford, Sandip Roy, and Shreyas Sundaram. On the complexity and approximability of optimal sensor selection and attack for Kalman filtering. *IEEE Transactions on Automatic Control*, 66(5):2146–2161, 2020.

[196] M. Manshaei, Q. Zhu, T. Alpcan, T. Basar, and J. P. Hubaux. Game theory meets network security and privacy. *ACM Computing Surveys*, 45:53–73, 2013.

[197] A. Alpcan and S. Buchegger. Security games for vehicular networks. *SIAM Journal on Control and Optimization*, 10:280–290, 2011.

[198] Z. Han, D. Niyato, W. Saad, T. Başar, and A. Hjorungnes. Game theory in wireless and communication networks: theory, models, and applications. *Cambridge university press*, 2012.

[199] Lintao Ye, Sandip Roy, and Shreyas Sundaram. Resilient sensor placement for Kalman filtering in networked systems: Complexity and algorithms. *IEEE Transactions on Control of Network Systems*, 7(4):1870–1881, 2020.

[200] J. Milosevic, M.Dahan, S. Amin, and H. Sandberg. A network monitoring game with heterogeneous component criticality levels. In *arXiv preprint arXiv:1903.07261*, 2019.

[201] M. Dahan, L. Sela, and S. Amin. Network inspection for detecting strategic attacks. *arXiv:1705.00349v3*, 2018.

[202] M. Pirani, E. Nekouie, H. Sandberg, and K.H.Johansson. A game-theoretic framework for security-aware sensor placement problem in networked control systems. *Proceedings of ACC, the 38th American Control Conference*, 2019.

[203] M. Pirani, J. A. Taylor, and B. Sinopoli. Strategic sensor placement on graphs. *Systems and Control Letters*, 148:104855, 2021.

[204] M. Pirani, E. Nekouie, H. Sandberg, and K.H.Johansson. A graph-theoretic equilibrium analysis of attacker-defender game on consensus dynamics under H_2 performance metric. *IEEE Transactions on Network Science and Engineering*, 2020.

[205] Q. He, C. Wang, G. Cuiand B. Li, R. Zhou, Q. Zhou, Y. Xiang anf H. Jin, and Y. Yang. A game-theoretical approach for mitigating edge ddos attack. *IEEE Transactions on Dependable and Secure Computing*, 2021.

[206] A. Gueye and V. Marbukh. A game-theoretic framework for network security vulnerability assessment and mitigation. *International Conference on Decision and Game Theory for Security*, 2012.

[207] B. Bamieh, M. R. Jovanovic, P. Mitra, and S. Patterson. Coherence in large-scale networks: Dimension-dependent limitations of local feedback. *IEEE Transactions on Automatic Control*, 57:2235–2249, 2012.

[208] E. Tegling, B. Bamieh, and D. F. Gayme. The price of synchrony: Evaluating the resistive losses in synchronizing power networks. *IEEE Transactions on Control of Network Systems*, 2:254–266, 2015.

[209] M. Pirani, J. W. Simpson-Porco, and B. Fidan. System-theoretic performance metrics for low-inertia stability of power networks. *56th IEEE Conference on Decision and Control*, pages 5106–5111, 2017.

[210] B. K. Poolla, S. Bolognani, and F. Dörfler. Optimal placement of virtual inertia in power grids. *IEEE Transactions on Automatic Control*, 62:6209–6220, 2015.

[211] K. E. Fitch and N. E. Leonard. Joint centrality distinguishes optimal leaders in noisy networks. *IEEE Transactions on Control of Network Systems*, 3:366–378, 2015.

[212] Mohammad Pirani and Shreyas Sundaram. On the smallest eigenvalue of grounded Laplacian matrices. *IEEE Transactions on Automatic Control*, 61(2):509–514, 2016.

[213] M. Pirani, E. Hashemi, A. Khajepour, B. Fidan, B. Litkouhi, S. K. Chen, and S. Sundaram. Cooperative vehicle speed fault diagnosis and correction. *IEEE Transactions on Intelligent Transportation Systems (to appear)*, 20(2):783–789, 2019.

[214] J. Usevitch and D. Panagou. r-robustness and (r, s)-robustness of circulant graphs. *Annual Conference on Decision and Control*, pages 4416–4421, 2017.

[215] L. Guerrero-Bonilla, D. Saldana, and V. Kumar. Design guarantees for resilient robot formations on lattices. *IEEE Robotics and Automation Letters*, 4:89–96, 2018.

[216] K. Saulnier, D. Saldana, A. Prorok, G. J. Pappas, and V. Kumar. Resilient flocking for mobile robot teams. *IEEE Robotics and Automation Letters*, 2:1039–1046, 2017.

[217] B. Bollobás. Random graphs, 2nd edition,. *Cambridge University Press*, 2001.

[218] M. E.j. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Physical review E*, 64, 2001.

[219] P. Crucitti, V. Latora, and M. Marchiori. Model for cascading failures in complex networks. *Physical Review E*, 69, 2004.

[220] O. Yagan, D. Qian, J. Zhang, and D. Cochran. Optimal allocation of interconnecting links in cyber-physical systems: Interdependence, cascading failures, and robustness. *IEEE Transactions on Parallel and Distributed Systems*, 23:1708–1720, 1998.

[221] J. Friedman. A proof of Alon's second eigenvalue conjecture. *Proceedings of the Thirty Fifth Annual ACM Symposium on Theory of Computing*, pages 720–724, 2003.

[222] M. Parandehgheibi and E. Modiano. Robustness of interdependent networks: The case of communication networks and the power grid. *IEEE Global Communications Conference (GLOBECOM)*, pages 2164–2169, 2013.

[223] E. M. Shahrivar, M. Pirani, and S. Sundaram. Spectral and structural properties of random interdependent networks. *Automatica*, 83:234–242, 2017.

[224] K. Singer-Cohen. Random intersection graphs. *PhD thesis, Department of Mathematical Sciences, The Johns Hopkins University*, 1995.

[225] J. Zhao, O. Yağan, and V. Gligor. On the strengths of connectivity and robustness in general random intersection graphs. *IEEE Conference on Decision and Control*, pages 3661–3668, 2014.

[226] F. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[227] S. Weerakkody, X. Liu, and B. Sinopoli. Robust structural analysis and design of distributed control systems to prevent zero dynamics attacks. *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1356–1361, 2017.

[228] W. Abbas, M. Shabbir, H. Jaleel, and X. Koutsoukos. Improving network robustness through edge augmentation while preserving strong structural controllability. *arXiv preprint arXiv:2003.05490*, 2020.

[229] A. Clauset, C. Rohilla Shalizi, and M. E. Newman. Power-law distributions in empirical data. *SIAM review*, 51:661–703, 2009.

[230] F. Dörfler and F. Bullo. Synchronization in complex networks of phase oscillators: A survey. *Automatica*, 50(6):1539–1564, 2014.

[231] V. Gazi and K. M. Passino. Stability analysis of swarms. *IEEE transactions on automatic control*, 48(4):692–697, 2003.

[232] Yijing Xie, Shaoshuai Mou, and Shreyas Sundaram. Towards resilience for multi-agent qd-learning. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 1250–1255, 2021.

[233] Yixuan Lin, Shripad Gade, Romeil Sandhu, and Ji Liu. Toward resilient multi-agent actor-critic algorithms for distributed reinforcement learning. In *2020 American Control Conference (ACC)*, pages 3953–3958. IEEE, 2020.

[234] Zhaoxian Wu, Han Shen, Tianyi Chen, and Qing Ling. Byzantine-resilient decentralized td learning with linear function approximation. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5040–5044, 2021.

[235] Martin Figura, Krishna Chaitanya Kosaraju, and Vijay Gupta. Adversarial attacks in consensus-based multi-agent reinforcement learning. In *2021 American Control Conference (ACC)*, pages 3050–3055. IEEE, 2021.