

Resilient Distributed State Estimation with Mobile Agents: Overcoming Byzantine Adversaries, Communication Losses, and Intermittent Measurements

Aritra Mitra, John A. Richards, Saurabh Bagchi, and Shreyas Sundaram

Received: date / Accepted: date

Abstract Applications in environmental monitoring, surveillance and patrolling typically require a network of mobile agents to collectively gain information regarding the state of a static or dynamical process evolving over a region. However, these networks of mobile agents also introduce various challenges, including intermittent observations of the dynamical process, loss of communication links due to mobility and packet drops, and the potential for malicious or faulty behavior by some of the agents. The main contribution of this paper is the development of resilient, fully-distributed, and provably correct state estimation algorithms that simultaneously account for each of the above considerations, and in turn, offer a general framework for reasoning about state estimation problems in dynamic, failure-prone and adversarial environments. Specifically, we develop a simple switched linear observer for dealing with the issue of time-varying measurement models, and resilient filtering techniques for dealing with worst-case adversarial behavior subject to time-varying communication patterns among the agents. Our approach considers both communication patterns that recur in a deterministic manner, and patterns that are induced by

random packet drops. For each scenario, we identify conditions on the dynamical system, the patrols, the nominal communication network topology, and the failure models that guarantee applicability of our proposed techniques. Finally, we complement our theoretical results with detailed simulations that illustrate the efficacy of our algorithms in the presence of the technical challenges described above.

Keywords Distributed State Estimation · Byzantine Attacks · Resilient Robotic Teams · Situational Awareness

1 Introduction

Consider a dynamical process evolving over a geographical region. Measurements of this process are available at certain sensing locations distributed over the region. A set of mobile agents is tasked with collectively estimating the state of the dynamical process by executing patrols that visit the various sensing locations, and exchanging information with each other over a communication medium.

There are various challenges that arise in enabling the agents to achieve the above task. The first challenge arises from the fact that the agents are mobile, and hence, do not have continuous access to the measurements from any given sensing location. Thus, even when the monitored dynamical process is described by a time-invariant system, the measurement model for any given mobile agent is time-varying. The second challenge arises from the fact that agents may be assigned to different portions of the overall region, and execute persistent patrols that visit only a subset of the sensing locations. Thus, each agent can only directly estimate

A. Mitra, S. Bagchi, and S. Sundaram are with the School of Electrical and Computer Engineering at Purdue University. Email: {mitra14, sbagchi, sundara2}@purdue.edu. J. A. Richards is with Sandia National Laboratories. Email: jaricha@sandia.gov. This work was supported in part by NSF CAREER award 1653648, and by a grant from Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

a portion of the overall state, and must rely on (carefully crafted) information exchanges with other agents in order to recover the entire state. Such information exchange rules must not only account for time-varying communication links between the agents (due to mobility and communication losses), but also for malicious agents that seek to disrupt the state estimation algorithm. Such malicious behavior can arise, for example, due to some agents being compromised by an attacker who causes the agents to report incorrect information, deviate from their patrols [1], or drop out of the network altogether [2]. Indeed, as we show in Example 1 later in the paper, without accounting for such behavior, a single adversarial agent can potentially disrupt the overall state estimation process.

Given the problem and associated challenges listed above, **the goal of this paper is to formulate resilient distributed state estimation algorithms that allow networks of mobile agents to estimate the state of the monitored dynamical process, despite time-varying measurement models, time-varying communication links, and malicious adversaries.**

Applications: The framework developed in this paper can be employed for the purpose of environmental monitoring [3–6], oceanographic explorations [7–9], and surveillance with civilian [10] and military [11, 12] applications. Essentially, the task of monitoring the state of a changing environment using autonomous mobile agents falls within the purview of our present analysis. For instance, one might be interested in monitoring spatio-temporal processes where a non-negative scalar quantity (e.g., temperature, oil, dirt, salinity or traffic congestion) constitutes the state of interest (see [6–8]). One of the key applications of our framework, however, pertains to mission-critical scenarios where adversarial attacks on the mobile agents can have far-reaching consequences. A specific example of such a scenario involves the use of autonomous mobile robots for estimating radiation concentrations around nuclear plants, following leakages that are either accidental or due to malicious intent [13–15]. Emergency response in such hazardous environments dictates the need for attack-immune distributed approaches, and therein lies the practical motivation of our work.

1.1 Related Work

To highlight the specific contributions made by this paper, we now provide a comprehensive discussion of the similarities and differences existing between our problem formulation and various related domains.

Persistent Monitoring: When monitoring the state of a process that grows over time, it is necessary to persistently visit locations where information regarding the process is available. This leads to the notion of *persistent monitoring*, a problem that has been extensively studied in the robotics community [16–21]. Typically, the persistent monitoring literature aims to design the trajectories of the mobile agents so as to accomplish the persistent task in an optimal manner. In contrast, our main focus is centered around estimating the state of an underlying dynamical process, despite time-varying measurement models, communication losses, and adversarial attacks. In particular, our analysis reveals various conditions to be met by the patrol so as to guarantee stability of the estimation error dynamics (based on our proposed strategy). These conditions are a combination of system-theoretic requirements, network-connectivity requirements, and requirements dictated by the adversarial and communication loss models. In this sense, our work complements the existing literature on persistent monitoring by providing insights into the design of joint patrolling and state estimation schemes in dynamic, failure-prone and adversarial environments.

Sensor Scheduling and Active Information Gathering: Given a dynamical system affected by noise, and a set of sensors measuring the states of the system, the sensor scheduling literature [22–25] aims to design a rule for choosing sensors sequentially (in time) so as to minimize a metric that appropriately captures performance against noise. Most of the sensor scheduling literature [23–25] focuses on the finite-horizon version of the problem, and hence, stability of the error process is not a major concern in these works. A notable exception is the very recent work in [26] that provides exact conditions under which an infinite-horizon sensor schedule leads to a uniformly bounded sequence of error covariance matrices.

There are various similarities between the sensor scheduling problem as described above, and the problem of *active information gathering* in mobile robotics [27–29]. Indeed, the design of a patrol visiting the various sensing locations in the latter is analogous (to a certain extent) to the design of a sensor scheduling policy in the former. The formulations in [27–29] differ from the standard sensor scheduling setup by explicitly accounting for the motion models of the mobile sensors under consideration. However, the focus still remains on finite-horizon settings. In contrast to the sensor scheduling literature and the active information gathering formulations, our primary goal is to identify conditions on the patrols that guarantee *stability* of the estimation error dynamics of each (uncom-

promised) mobile agent. The recent work [29] extends the approach and results in [28] to a scenario where a certain number of mobile sensors are under attack. In addition to various other differences, our formulation involves a distributed setup (where the communication network plays a key role) unlike the decentralized setup considered in [28] and [29].

Distributed State Estimation: The problem of estimating the state of a (linear time-invariant) dynamical process using a network of static sensors has been studied by several researchers over the past decade [30–37]. However, single-time-scale algorithms that solve such problems under the most general conditions on the system and network have been proposed only recently in [38–44]. While the works stated above primarily cater to time-invariant communication graphs, the authors in [45] propose a hybrid observer that accounts for a broad class of time-varying networks. Although these papers provide a rich variety of approaches, none of them deal with the aspect of adversarial agents. Preliminary attempts towards addressing adversarial behavior in the context of distributed state estimation were undertaken in [46, 47], but without any theoretical guarantees. Recently, the authors in [48] developed an H_∞ -based filtering approach for detecting biasing attacks in sensor networks. While the analysis in [48] was limited to a certain class of attack inputs, much more general adversarial models were considered in our prior work [49, 50], albeit for time-invariant networks and measurement models.

Resilient Distributed Algorithms: Recent years have witnessed a significant amount of research dedicated towards the design of resilient distributed algorithms, with applications to consensus [51, 52], optimization [53, 54], hypothesis testing [55], static parameter estimation [56] and broadcasting [57]. Researchers in the robotics community have also looked into the problem of forming and maintaining robust mobile-robot formations that facilitate resilient consensus [58–64]. Thus, a key aspect of such problems is the identification of network topologies that are robust to different adversarial models. Unlike the consensus scenario, the results in [49, 50] indicate that when it comes to estimating the state of an external dynamical system despite adversarial behavior, one needs to incorporate redundancy in not only the network topology, but also the measurement structure of the sensors. However, as mentioned earlier, the analysis in [49, 50] was limited to time-invariant communication networks and static agents. In light of the above discussion, the main contributions of this paper are as follows.

Summary of Contributions: We consider a set of mobile agents tasked with estimating the state of a linear time-invariant dynamical system. Each agent is assumed to have a predefined patrol that visits a subset of the sensing locations. In Section 3, we develop a simple switched linear observer that allows a given mobile agent to recover those states that can be detected based on the measurements of the sensing locations it persistently visits. We establish asymptotic stability of the proposed observer for a class of periodic patrols. In Section 5, we consider a class of deterministic communication loss patterns, and develop a resilient distributed state estimation algorithm that allows each agent to process the information received from other agents to recover the true state, despite arbitrary adversarial behavior. Our algorithm is inspired by recent work that addresses the resilient consensus problem in asynchronous settings [61, 65]. As a byproduct of our analysis, we argue that our proposed algorithm probably works even in the presence of bounded (potentially random, time-varying) communication delays. We also characterize the convergence time of our algorithm in terms of the system instability, the upper bound on the delay, and certain properties of the communication network topology.

In Section 7, we model the communication links among the mobile agents as analog erasure channels that randomly drop packets based on an i.i.d. Bernoulli process. For this model, we propose a simple state estimate update rule, and identify conditions on the dynamical system, the network topology, and the erasure probability that guarantee mean-square-stability of the estimation error process. We show how a notion of network robustness (suitable for the problem under consideration) known as ‘strong-robustness’ allows one to deal with high packet drop probabilities, while still guaranteeing stability. We support our theoretical results via detailed simulations discussed in Section 8. Finally, we emphasize that all our results apply to a sophisticated and worst-case adversarial model (termed Byzantine adversaries) which is typically considered in the literature on resilient distributed algorithms [51, 52, 66]. From an implementation standpoint, the results obtained in this paper provide guidelines for designing patrols that account for each of the technical challenges discussed in this section.

We reported certain preliminary results in [67]. In this paper, we significantly expand upon the content in [67] by considering mobile agents instead of static agents (which leads to the aspect of time-varying measurement models), providing full proofs of all results, and supporting such results with illustrations and detailed simulations.

2 Problem Formulation

In this section, we will first clarify the notation to be used throughout the paper. Subsequently, we will describe each of the constituent models needed to formally define the problem of interest.

Notation: A directed graph is denoted by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \dots, m\}$ is the set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents the edges. An edge from node j to node i , denoted by (j, i) , implies that node j can transmit information to node i . The neighborhood of the i -th node is defined as $\mathcal{N}_i \triangleq \{j \mid (j, i) \in \mathcal{E}\}$. The notation $|\mathcal{V}|$ is used to denote the cardinality of a set \mathcal{V} . Throughout the rest of this paper, we use the terms ‘edges’ and ‘communication links/channels’ interchangeably. The set of all eigenvalues (or modes) of a matrix \mathbf{A} is denoted by $sp(\mathbf{A}) = \{\lambda \in \mathbb{C} \mid \det(\mathbf{A} - \lambda \mathbf{I}) = 0\}$ and the set of all marginally stable and unstable eigenvalues of \mathbf{A} is denoted by $\Lambda_U(\mathbf{A}) = \{\lambda \in sp(\mathbf{A}) \mid |\lambda| \geq 1\}$. The notations \mathbb{N} and \mathbb{N}_+ are used to denote the set of all non-negative integers and positive integers, respectively. For a random variable \mathbb{X} , its expected value is denoted by $E[\mathbb{X}]$.

Dynamical System Model: Throughout this paper, we will focus on a linear time-invariant dynamical process of the form

$$\mathbf{x}[k+1] = \mathbf{A}\mathbf{x}[k], \quad (1)$$

where $k \in \mathbb{N}$ is the discrete-time index, $\mathbf{x}[k] \in \mathbb{R}^n$ is the state vector and $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the system matrix. Measurements of the dynamical process (1) are available at N sensing locations distributed over a geographical region. The notation $\mathcal{Q} = \{1, \dots, N\}$ will be used to denote the set of all sensing locations. At each location $i \in \mathcal{Q}$, measurements of a portion of the state $\mathbf{x}[k]$ are available via the following observation model:

$$\mathbf{y}^{(i)}[k] = \mathbf{C}^{(i)}\mathbf{x}[k], \quad (2)$$

where $\mathbf{y}^{(i)}[k] \in \mathbb{R}^{r_i}$ and $\mathbf{C}^{(i)} \in \mathbb{R}^{r_i \times n}$. We denote $\mathbf{y}[k] = [\mathbf{y}^{(1)}[k]^T \dots \mathbf{y}^{(N)}[k]^T]^T$, and $\mathbf{C} = [\mathbf{C}^{(1)T} \dots \mathbf{C}^{(N)T}]^T$. For a set $\mathcal{S} \subset \{1, \dots, N\}$, $\mathbf{C}^{(\mathcal{S})}$ will be used to denote the collective measurement matrix corresponding to the sensing locations in the set \mathcal{S} . Such linear (in state) dynamical and observational models are standard in the literature on state estimation [27–29].

An eigenvalue $\lambda \in \Lambda_U(\mathbf{A})$ is said to be a detectable eigenvalue w.r.t. the pair $(\mathbf{A}, \mathbf{C}^{(i)})$ if

$$\text{rank} \begin{bmatrix} \mathbf{A} - \lambda \mathbf{I} \\ \mathbf{C}^{(i)} \end{bmatrix} = n. \quad (3)$$

An eigenvalue with magnitude strictly less than one is considered to be detectable w.r.t. any measurement set. Although we consider noiseless dynamics for clarity of exposition (like [35, 36, 38–42, 45]), the techniques developed in this paper guarantee bounded mean square estimation error in the presence of i.i.d. process and measurement noise with bounded second moments.

Mobile Agent Model: A set $\mathcal{V} = \{1, \dots, m\}$ of m mobile agents is tasked with collaboratively estimating the state $\mathbf{x}[k]$ of the process (1) by persistently visiting the N sensing locations. Specifically, each agent $i \in \mathcal{V}$ is assigned a persistent patrol through a subset of the sensing locations. Over the course of its patrol, each agent can communicate with certain other agents (e.g., when the distance between the agents is less than some communication radius). In the absence of any communication losses, a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is used to model the flow of information between the m mobile agents. Specifically, the graph \mathcal{G} captures the set of all possible agent interactions across time. In other words, $(i, j) \in \mathcal{E}$ implies that agent i will be in a position to directly transmit information to agent j infinitely often while executing its patrol. The graph \mathcal{G} will be referred to as the *baseline communication graph*. The loss of communication between agents (due to agent movements or packet drops) is modeled by a time-varying graph $\mathcal{G}[k] = (\mathcal{V}, \mathcal{E}[k])$, where $\mathcal{E}[k] \subseteq \mathcal{E}$ for all $k \in \mathbb{N}$.

Remark 1 The notion of ‘sensing locations’ and ‘mobile agents’ as discussed above can be used to capture the following different scenarios. (1) In the first scenario, one can assume that physical static sensors are located at each of the N sensing locations, and that the mobile agents obtain measurements from such sensors on visiting the corresponding sensing locations. (2) In the second scenario, one can envision sensors installed on the mobile agents themselves. The sensing locations can then be interpreted as informative points in the geographical region where a mobile sensor can obtain non-zero measurements of the state. Specifically, a mobile sensor present at the i -th location would generate a measurement model of the form (2). In either case, the measurements acquired by a mobile agent are a function of its movement pattern (patrol). The mathematical framework developed in this paper applies identically to each of the above physically different scenarios. For the rest of the paper, we will stick to the first interpretation (for the purposes of illustration), i.e., static ground sensors positioned at the sensing locations communicate with mobile agents passing by.

Adversary Model: A subset $\mathcal{A} \subset \mathcal{V}$ of the mobile agents are adversarial. We assume that the adversarial

agents possess complete knowledge of the dynamical system model, the time-varying communication graph topology, the patrolling strategy and any estimation algorithm employed by the non-adversarial agents. Adversarial agents are not only allowed to update and transmit state estimates in an arbitrary manner, but also to deviate from the rules of any patrolling algorithm. Furthermore, following the Byzantine fault model [66], adversaries are allowed to send differing state estimates to different neighbors at the same instant of time. Adversaries can also choose not to transmit any estimates at all to agents within communication radius. This assumption of omniscient adversarial behavior is motivated by the aim of providing theoretical guarantees against “worst-case” adversarial behavior. We point out that such Byzantine models have been commonly studied in the context of distributed consensus and optimization problems in [51–54]. In return for endowing the adversaries with such worst-case capabilities, we assume that there are at most f adversarial agents in the neighborhood of any non-adversarial agent in the baseline communication graph \mathcal{G} , for some constant $f \in \mathbb{N}$. This property will be referred to as the ‘ f -local’ property of the adversarial set. Summarily, the adversary model described thus far will be called an f -local Byzantine adversary model. The non-adversarial mobile agents will be referred to as regular agents and be represented by the set $\mathcal{R} = \mathcal{V} \setminus \mathcal{A}$. Finally, we remark that the number and identities of the adversarial agents are not known to the regular agents. The regular agents are only aware of the upper bound f on the number of adversaries in their neighborhood (in the baseline communication graph \mathcal{G}).

Given the above setup, we can now describe the problem studied in this paper.

Objective: Suppose we are given the LTI system (1), the measurement model (2), a set of m mobile agents with a baseline communication graph \mathcal{G} executing a patrol, and an f -local Byzantine adversary model. Our objectives are (i) to develop distributed state estimation algorithms that account for time-varying measurement models, communication losses and worst-case attacks, and (ii) to analyze under what conditions (on the dynamical model, the baseline communication graph, the patrols and the communication loss patterns) the proposed algorithms provably enable each regular mobile agent to asymptotically estimate the true state of the system (in a deterministic or stochastic sense).

Achieving the above objective is non-trivial, due to the need to simultaneously address the three challenges (time-varying measurement models, time-varying networks, and adversarial agents). In this paper, we take a significant step in this direction.

At this stage, we should clarify the answer to the following important question: **Based on the problem formulated above, what can one expect from the theoretical results in this paper, when it comes to the aspect of designing the motion plan of the mobile agents?** Briefly, our main results (namely, Proposition 1, Theorem 1, and Theorem 2) lay down various rules that need to be met by the patrols so that they effectively complement the estimation techniques developed in the paper. These rules are tailored to meet the specific technical challenges considered in this work, and answer questions such as: (i) How often does a mobile agent need to visit a sensing location that provides critical information regarding the process of interest? (ii) Given that certain agents can be under attack, how many mobile agents should visit each such location? (iii) How often should agents exchange information amongst themselves? (iv) Given that certain agents can be under attack, how can one resiliently diffuse information across the mobile agent network?

While our answers to the above questions provide high-level specifications that significantly inform the process of patrol design, there are various questions that are left open. For instance, consider the following allocation problem. We are given a fixed number of sensing locations and mobile agents. Constraints are placed that limit the sensing regions of the mobile agents, i.e., each agent can only visit a subset of the sensing locations in the region. Given such constraints, how does one allocate mobile agents to sensing locations while meeting the specifications laid down by our theoretical results? Can such a patrol be designed in the first place? What is the minimal number of mobile agents that is needed to achieve the distributed state estimation task?¹ Such questions are inherently of a combinatorial nature, and addressing them comprehensively is beyond the scope of the present paper. In Section 8, we do, however, discuss a simple strategy for designing patrols that meet the required specifications. Finally, note that the rules imposed by our results can either be used as a guideline when synthesizing patrols, or alternatively as a checklist when given predefined patrols.

To avoid cumbersome notation and to clearly present the key ideas, we make the following assumption on the system matrix in (1).

¹ In the absence of any constraints placed on the sensing capabilities or movement patterns of an agent, one can just have each mobile agent patrol all the sensing locations. However, such an assumption would in general be impractical, thereby necessitating inter-agent communication. Note that it is precisely the need for inter-agent communication that makes the issues of communication losses and adversarial attacks studied in this paper relevant.

Assumption 1. *The system matrix \mathbf{A} has real, distinct eigenvalues.*

While we make the above assumption, the results obtained in this paper can be generalized to system matrices with arbitrary spectrum via a more detailed technical approach (e.g., as outlined in [50], which considered the effects of adversarial behavior for networks with time-invariant topologies and measurements, as opposed to the more general setting that we consider in this paper). The assumption of real eigenvalues, in particular, considerably simplifies the structure of the resilient filtering algorithms studied in Sections 5 and 7, and hence, facilitates a better understanding of our core algorithmic ideas. Note that the assumption of a real spectrum applies to various relevant dynamical models including (but not limited to) the discretized double integrator moving target model considered in [29], the methane gas concentration model considered in [27], and the diffusion dynamics models studied in [68–70].

Regarding the observation model (2), we assume that the pair (\mathbf{A}, \mathbf{C}) is detectable. Clearly, this is a basic necessary condition for state estimation even in the absence of time-varying measurement models, packet drops, or attacks. It should be pointed out that for any given location $i \in \mathcal{Q}$, we do not assume detectability of the pair $(\mathbf{A}, \mathbf{C}^{(i)})$. In a similar spirit, we do not assume that the set of sensing locations visited by any agent during its patrol is informative enough to allow that agent to recover the entire state.

Having introduced the main problem and its specific technical challenges, we now proceed to develop a solution that addresses these challenges in the subsequent sections.

3 Periodic Patrols for Estimating Locally Detectable States

There are two main goals that we seek to achieve in this section. First, we will focus on the design of a simple switched linear observer that enables each regular agent to estimate those states that are detectable w.r.t. the measurements of the sensing locations it persistently visits. Second, we will identify conditions on the patrol that guarantee asymptotic stability of the error dynamics induced by the proposed switched linear observer. Once the aforementioned objectives are met, a regular agent can be viewed as a source agent for the states that are detectable via the sensing locations it visits.

3.1 Design of Switched Linear Observers

To achieve the objectives stated above, we first note that based on Assumption 1, one can perform a coordinate transformation $\mathbf{z}[k] \triangleq \mathbf{V}\mathbf{x}[k]$ on (1) with an appropriate non-singular matrix \mathbf{V} to obtain

$$\begin{aligned} \mathbf{z}[k+1] &= \mathbf{M}\mathbf{z}[k] = \text{diag}(\lambda_1, \dots, \lambda_n)\mathbf{z}[k], \\ \mathbf{y}^{(i)}[k] &= \bar{\mathbf{C}}^{(i)}\mathbf{z}[k], \quad \forall i \in \{1, \dots, N\} \end{aligned} \quad (4)$$

where $sp(\mathbf{A}) = \{\lambda_1, \dots, \lambda_n\}$, $\mathbf{M} = \mathbf{V}\mathbf{A}\mathbf{V}^{-1}$ and $\bar{\mathbf{C}}^{(i)} = \mathbf{C}^{(i)}\mathbf{V}^{-1}$. Commensurate with this decomposition, the j -th component of the state vector $\mathbf{z}[k]$ will be denoted by $z^{(j)}[k]$, and will be referred to as the component corresponding to the eigenvalue λ_j . Since a non-singular transformation maps $\mathbf{z}[k]$ to $\mathbf{x}[k]$, we focus on estimating $\mathbf{z}[k]$. Consider any regular agent $i \in \mathcal{R}$, and let the subset of sensing locations it visits be denoted by $\mathcal{P}_i = \{i_1, \dots, i_{|\mathcal{P}_i|}\}$, where $\mathcal{P}_i \subset \mathcal{Q}$. Let $\mathcal{O}^{(i_r)}$ denote the eigenvalues of \mathbf{A} that are detectable w.r.t. the measurements available at location i_r (i.e., $\mathcal{O}^{(i_r)}$ denotes the set of detectable eigenvalues of the pair $(\mathbf{A}, \mathbf{C}^{(i_r)})$). Thus, the set of all eigenvalues that are detectable w.r.t. the set of sensing locations \mathcal{P}_i is given by $\mathcal{O}_i \triangleq \bigcup_{r=1}^{|\mathcal{P}_i|} \mathcal{O}^{(i_r)}$ (in other words, \mathcal{O}_i denotes the set of all eigenvalues of \mathbf{A} that are detectable w.r.t. the collective measurement set $\mathbf{C}^{(\mathcal{P}_i)}$). Our goal is to design an observer that enables agent i to asymptotically estimate all the components of $\mathbf{z}[k]$ corresponding to the eigenvalues in \mathcal{O}_i . To achieve this goal, we will build a partial observer for each location visited by agent i . Specifically, the partial observer at location $i_r \in \mathcal{P}_i$ will be designed to recover the states of $\mathbf{z}[k]$ corresponding to the eigenvalues in $\mathcal{O}^{(i_r)}$. Let such states be denoted by $\mathbf{v}^{(i_r)}[k]$.² We make two simple observations at this point. First, if $\mathbf{J}^{(i_r)}$ represents the diagonal matrix with the eigenvalues in $\mathcal{O}^{(i_r)}$ on its diagonal, then we have

$$\mathbf{v}^{(i_r)}[k+1] = \mathbf{J}^{(i_r)}\mathbf{v}^{(i_r)}[k]. \quad (5)$$

The above equation follows directly from the definitions of each of the components involved in the equation, and the decoupled nature of the dynamics (4). The second observation is as follows:

$$\mathbf{y}^{(i_r)}[k] = \tilde{\mathbf{C}}^{(i_r)}\mathbf{v}^{(i_r)}[k], \quad (6)$$

where $\tilde{\mathbf{C}}^{(i_r)}$ contains the columns of $\bar{\mathbf{C}}^{(i_r)}$ corresponding to the matrix $\mathbf{J}^{(i_r)}$. The second observation follows from the fact that for a system with distinct eigenvalues, a given unstable or marginally stable eigenvalue

² We resort to such a notation here since the superscript on the $z[k]$ states are reserved for eigenvalues, and the subscripts are reserved for mobile agents. Thus, we introduce the notation $\mathbf{v}[k]$, with a superscript on $\mathbf{v}[k]$ pointing to a location number.

is detectable if and only if the column of the measurement matrix corresponding to that eigenvalue is non-zero [71]. Let $\sigma_i : \mathbb{N} \rightarrow \mathcal{J} = \{1, \dots, N, \omega\}$ be a function that records the location of the i -th mobile agent at time-step k . Specifically, for $i \in \{1, \dots, N\}$, $\sigma_i[k] = i_r$ implies that the i -th mobile agent is at location i_r at time-step k , whereas $\sigma_i[k] = \omega$ implies that it is commuting between locations at time-step k . We are now in position to propose the following switched linear partial observer for estimating $\mathbf{v}^{(i_r)}[k]$:

$$\hat{\mathbf{v}}_i^{(i_r)}[k+1] = \mathbf{F}_{\sigma_i[k]}^{(i_r)} \hat{\mathbf{v}}_i^{(i_r)}[k] + \alpha_{\sigma_i[k]}^{(i_r)} \mathbf{L}_i^{(i_r)} \mathbf{y}^{(i_r)}[k], \quad (7)$$

where

$$\mathbf{F}_{\sigma_i[k]}^{(i_r)} = \begin{cases} (\mathbf{J}^{(i_r)} - \mathbf{L}_i^{(i_r)} \tilde{\mathbf{C}}^{(i_r)}) & \text{if } \sigma_i[k] = i_r, \\ \mathbf{J}^{(i_r)} & \text{if } \sigma_i[k] \neq i_r, \end{cases} \quad (8)$$

$$\alpha_{\sigma_i[k]}^{(i_r)} = \begin{cases} 1 & \text{if } \sigma_i[k] = i_r, \\ 0 & \text{if } \sigma_i[k] \neq i_r. \end{cases}$$

In the above equations, $\hat{\mathbf{v}}_i^{(i_r)}[k]$ represents the estimate of $\mathbf{v}^{(i_r)}[k]$ maintained by the i -th mobile agent, and $\mathbf{L}_i^{(i_r)}$ represents an output-injection gain that needs to be designed appropriately to guarantee asymptotic stability of the estimation error dynamics.³ The purpose of the partial observer given by equations (7) and (8) is to allow the i -th mobile agent to recover the states that are detectable w.r.t. the measurements of location i_r , namely, the states aggregated in the vector $\mathbf{v}^{(i_r)}[k]$. From the structure of the observer, we note that the i -th mobile agent switches between a Luenberger-style update rule and an open-loop update rule, depending upon its current position.

3.2 Periodic Patrols and Stability Analysis

As indicated by the above discussion, the stability of the proposed observer depends critically upon the movement patterns of the mobile agents. In what follows, we will restrict our attention to periodic patrols; such patrols are commonly considered in the literature (e.g., [16, 72]), and offer structure that can be leveraged to simplify our analysis. To formally characterize a periodic patrol, recall that $\mathcal{P}_i = \{i_1, \dots, i_{|\mathcal{P}_i|}\}$ represents the set of sensing locations visited by the i -th mobile agent. With each such location $i_r \in \mathcal{P}_i$, we associate a non-negative integer $\tau_i^{(i_r)}$ and a positive integer $T_i^{(i_r)}$ such that $\sigma_i[\tau_i^{(i_r)} + qT_i^{(i_r)}] = i_r, \forall q \in \mathbb{N}$. Here, $\tau_i^{(i_r)}$ represents the first time location i_r is visited by the i -th mobile agent, and $T_i^{(i_r)}$ represents the time-period

with which agent i visits location i_r . We say that the i -th mobile agent executes a **feasible** periodic patrol if: (i) the mobile agent is never at more than one sensing location at any given point in time, (ii) each location in the set \mathcal{P}_i is visited infinitely often, and (iii) a given location $i_r \in \mathcal{P}_i$ is visited at time-step k only if $k = \tau_i^{(i_r)} + qT_i^{(i_r)}$, for some $q \in \mathbb{N}$. Notice that the first two constraints place certain limitations on the values that $\tau_i^{(i_r)}$ and $T_i^{(i_r)}$ can take on. For instance, we must have $T_i^{(i_r)} \neq 1, \forall i_r \in \mathcal{P}_i$ (assuming $|\mathcal{P}_i| > 1$). The third property of a feasible periodic patrol implies that a mobile agent does not stay at any location in \mathcal{P}_i for more than a single time-step.

Let the vector $\mathbf{z}_{\mathcal{O}_i}[k] = [\mathbf{v}^{(i_1)}[k]^T \dots \mathbf{v}^{(i_{|\mathcal{P}_i|})}[k]^T]^T$ aggregate the components of $\mathbf{z}[k]$ that correspond to the set \mathcal{O}_i (recall that \mathcal{O}_i denotes the set of detectable eigenvalues w.r.t. the pair $(\mathbf{A}, \mathbf{C}^{(\mathcal{P}_i)})$). Our objective is to identify conditions on the time-periods $\{T_i^{(i_r)}\}$ that enable the i -th regular mobile agent to asymptotically recover $\mathbf{z}_{\mathcal{O}_i}[k]$. To this end, we need the following result.

Lemma 1 Consider a detectable pair (\mathbf{A}, \mathbf{C}) , where \mathbf{A} satisfies Assumption 1. Then, for any positive odd integer \bar{T} , the pair $(\mathbf{A}^{\bar{T}}, \mathbf{C})$ is also detectable.

Proof. Perform a similarity transformation that brings the pair (\mathbf{A}, \mathbf{C}) to the form $(\mathbf{M}, \bar{\mathbf{C}})$, where \mathbf{M} represents the Jordan canonical form of \mathbf{A} . Detectability of (\mathbf{A}, \mathbf{C}) then implies detectability of $(\mathbf{M}, \bar{\mathbf{C}})$. If \bar{T} is a positive odd integer, then based on Assumption 1, $\mathbf{M}^{\bar{T}}$ is a diagonal matrix with real distinct eigenvalues.⁴ Detectability of $(\mathbf{M}^{\bar{T}}, \bar{\mathbf{C}})$ follows as a consequence of the PBH test [71], and the detectability of $(\mathbf{M}, \bar{\mathbf{C}})$. Since a similarity transformation maps $(\mathbf{M}^{\bar{T}}, \bar{\mathbf{C}})$ back to $(\mathbf{A}^{\bar{T}}, \mathbf{C})$, the pair $(\mathbf{A}^{\bar{T}}, \mathbf{C})$ is also detectable. \square

Proposition 1 Suppose we are given the LTI system (1) and the measurement model (2) such that the system matrix \mathbf{A} is non-singular and satisfies Assumption 1. Let each regular mobile agent i execute a feasible periodic patrol characterized by the parameters $\tau_i^{(i_r)}$ and $T_i^{(i_r)}$, $i_r \in \mathcal{P}_i$, such that each time-period $T_i^{(i_r)}$ is a positive odd integer. Additionally, let each regular mobile agent i implement the observer given by (7) and (8). Then, for each such agent $i \in \mathcal{R}$, there exists a choice of output-injection gains $\{\mathbf{L}_i^{(i_r)}\}$ that guarantees asymptotic convergence of $\hat{\mathbf{z}}_{\mathcal{O}_i}[k]$ to $\mathbf{z}_{\mathcal{O}_i}[k]$, where $\hat{\mathbf{z}}_{\mathcal{O}_i}[k]$ represents the estimate of $\mathbf{z}_{\mathcal{O}_i}[k]$ maintained by agent i .

⁴ Essentially, an odd period ensures that eigenvalues that are equal in magnitude, but opposite in sign in \mathbf{A} , remain so in $\mathbf{A}^{\bar{T}}$. Thus, if the eigenvalues of \mathbf{A} are distinct in magnitude, then clearly no restrictions need to be imposed on the time-period \bar{T} .

³ The gains $\mathbf{L}_i^{(i_r)}$ are agent-specific, since different agents might visit the same location with different frequencies.

Proof. For a given mobile agent $i \in \mathcal{R}$, establishing that $\hat{\mathbf{z}}_{\mathcal{O}_i}[k]$ converges to $\mathbf{z}_{\mathcal{O}_i}[k]$ asymptotically requires us to establish that the estimation error dynamics associated with each location in \mathcal{P}_i converges to zero asymptotically. In other words, our aim is to prove that for each $i_r \in \mathcal{P}_i$, $\lim_{k \rightarrow \infty} \|\hat{\mathbf{v}}_i^{(i_r)}[k] - \mathbf{v}^{(i_r)}[k]\| = 0$. To this end, fix a location i_r , and let $\mathbf{e}_i^{(i_r)}[k] \triangleq \hat{\mathbf{v}}_i^{(i_r)}[k] - \mathbf{v}^{(i_r)}[k]$ denote the estimation error associated with location i_r . Based on (5), (6), (7), and (8), we obtain:

$$\mathbf{e}_i^{(i_r)}[k+1] = \mathbf{F}_{\sigma_i[k]}^{(i_r)} \mathbf{e}_i^{(i_r)}[k]. \quad (9)$$

Recalling that $\tau_i^{(i_r)}$ represents the first time agent i visits location i_r , $T_i^{(i_r)}$ represents the time-period with which agent i visits location i_r , and using (9), we obtain the following periodic error dynamics:

$$\mathbf{e}_i^{(i_r)}[\tau_i^{(i_r)} + (k+1)T_i^{(i_r)} + 1] = \mathbf{M}_i^{(i_r)} \mathbf{e}_i^{(i_r)}[\tau_i^{(i_r)} + kT_i^{(i_r)} + 1], \quad (10)$$

where $k \in \mathbb{N}$, and

$$\begin{aligned} \mathbf{M}_i^{(i_r)} &= (\mathbf{J}^{(i_r)} - \mathbf{L}_i^{(i_r)} \tilde{\mathbf{C}}^{(i_r)}) (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1} \\ &= (\mathbf{J}^{(i_r)})^{T_i^{(i_r)}} - \mathbf{L}_i^{(i_r)} \tilde{\mathbf{C}}^{(i_r)} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1}. \end{aligned} \quad (11)$$

To establish asymptotic stability of the periodic error dynamics (10), we need to argue that $\mathbf{M}_i^{(i_r)}$ can be chosen to make $\mathbf{M}_i^{(i_r)}$ Schur stable. Based on (11), this is equivalent to establishing detectability of the pair $((\mathbf{J}^{(i_r)})^{T_i^{(i_r)}}, \tilde{\mathbf{C}}^{(i_r)} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1})$. In other words, we need to show that

$$\text{rank} \begin{bmatrix} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)}} - s\mathbf{I} \\ \tilde{\mathbf{C}}^{(i_r)} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1} \end{bmatrix} = n^{(i_r)}, \quad \forall s \in \mathbb{C} \text{ s.t. } |s| \geq 1, \quad (12)$$

where $n^{(i_r)}$ represents the dimension of $\mathbf{J}^{(i_r)}$. Based on our construction, $\text{sp}(\mathbf{J}^{(i_r)}) \subseteq \text{sp}(\mathbf{A})$, and hence, $\mathbf{J}^{(i_r)}$ is non-singular since \mathbf{A} is assumed to be non-singular. Thus, the following is true for all $s \in \mathbb{C}$:

$$\begin{aligned} &\text{rank} \begin{bmatrix} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)}} - s\mathbf{I} \\ \tilde{\mathbf{C}}^{(i_r)} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1} \end{bmatrix} \\ &= \text{rank} \left(\begin{bmatrix} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)}} - s\mathbf{I} \\ \tilde{\mathbf{C}}^{(i_r)} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1} \end{bmatrix} \left((\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1} \right)^{-1} \right) \\ &= \text{rank} \begin{bmatrix} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)}} - s\mathbf{I} \\ \tilde{\mathbf{C}}^{(i_r)} \end{bmatrix}. \end{aligned} \quad (13)$$

Since the pair $(\mathbf{J}^{(i_r)}, \tilde{\mathbf{C}}^{(i_r)})$ is detectable by construction, the eigenvalues of $\mathbf{J}^{(i_r)}$ are real and distinct (since $\text{sp}(\mathbf{J}^{(i_r)}) \subseteq \text{sp}(\mathbf{A})$ and \mathbf{A} satisfies Assumption 1), and

the time-period $T_i^{(i_r)}$ is an odd positive integer, we infer that the pair $((\mathbf{J}^{(i_r)})^{T_i^{(i_r)}}, \tilde{\mathbf{C}}^{(i_r)})$ is also detectable by appealing to Lemma 1. Based on the foregoing discussion, referring to equations (12) and (13) reveals detectability of the pair $((\mathbf{J}^{(i_r)})^{T_i^{(i_r)}}, \tilde{\mathbf{C}}^{(i_r)} (\mathbf{J}^{(i_r)})^{T_i^{(i_r)} - 1})$. Thus, the observer gain $\mathbf{L}_i^{(i_r)}$ can indeed be chosen appropriately to stabilize the periodically sampled error dynamics (10). Notice that the quantity $\beta_i^{(i_r)} \triangleq \max\{\|(\mathbf{J}^{(i_r)} - \mathbf{L}_i^{(i_r)} \tilde{\mathbf{C}}^{(i_r)})\|, \|\mathbf{J}^{(i_r)}\|\}$ is finite since all matrices under consideration have finite norm. Since the time-period $T_i^{(i_r)}$ is also finite, the maximum error-norm amplification $(\beta_i^{(i_r)})^{T_i^{(i_r)}}$ of the error dynamics (9), over any time-period, is also finite. Asymptotic stability of the periodic error dynamics (10) then readily implies asymptotic stability of the error dynamics (9). This completes the proof. \square

Remark 2 Based on Proposition 1, we see that a given mobile agent $i \in \mathcal{R}$ is able to asymptotically estimate the portion of the state $\mathbf{z}[k]$ that corresponds to the detectable subspace of the pair $(\mathbf{A}, \mathbf{C}^{(P_i)})$ (namely, the portion that we refer to as $\mathbf{z}_{\mathcal{O}_i}[k]$). Furthermore, agent i is able to achieve this without communicating with any other mobile agent. In this sense, the detectable subspace of $(\mathbf{A}, \mathbf{C}^{(P_i)})$ can be viewed as the locally detectable portion of agent i , and agent i can be viewed as the source of information for all the states that correspond to its locally detectable eigenvalues (namely, the set of eigenvalues \mathcal{O}_i). It is important, however, to make a clear distinction between the notion of ‘local detectability’ used here, and that used in [39, 40, 45]. In these works, the task of distributed state estimation is performed collaboratively by a network of static sensors. As mentioned in the introduction, our present formulation is applicable to more general settings (the generalization arising due to the issue of intermittent observations) where the distributed state estimation task is executed either by a network of mobile agents that visit static sensors, or by a network of mobile sensors. Thus, while the locally detectable portion of a static sensor is simply the portion of the state space detectable via the measurements of that specific sensor, the locally detectable portion of a moving agent is the portion of the state space that is detectable w.r.t. the collective measurements of the sensing locations it persistently visits.

Remark 3 The assumption of non-singularity of \mathbf{A} in Proposition 1 is not restrictive. For a system with distinct eigenvalues, the component of the state corresponding to the zero eigenvalue will stay at zero for all time. Hence, the existence of an eigenvalue at zero does not affect our objective of asymptotic state reconstruction.

Remark 4 Note that it is possible for sensors located at distinct sensing locations to share common detectable eigenvalues. In terms of our observer design, this would cause a mobile agent to maintain multiple estimates of the same state corresponding to different sensing locations. Specifically, for a given mobile agent i , the vectors $\hat{\mathbf{v}}_i^{(i_p)}[k]$ and $\hat{\mathbf{v}}_i^{(i_q)}[k]$ corresponding to two distinct sensing locations $i_p, i_q \in \mathcal{P}_i$ might contain common components. One can readily eliminate this redundancy via a slight modification of the approach presented here. However, this comes at the expense of cluttering the exposition with more notation, and hence, we do not delve into such details in this paper.

4 Preliminaries for Resilient Distributed State Estimation

As pointed out earlier, a given mobile agent will be able to estimate only a portion of the system state by persistently visiting its set of sensing locations. To estimate its locally undetectable portion, it is reliant on the information received from its neighbors in the baseline communication graph \mathcal{G} (note that due to communication losses and agent movements, the neighborhood of a given agent at any time-step will in general only be a subset of its neighborhood in the baseline graph). It is precisely this aspect of the problem that dictates the need for robustness against adversarial attacks coupled with communication losses. The focus of this section will be to introduce some of the key ideas and terminology required to address the above issues. To this end, we introduce the notion of *source mobile agents*.

Definition 1 (Source Mobile Agent) A mobile agent $i \in \mathcal{V}$ is said to be a source mobile agent for an eigenvalue $\lambda_j \in \Lambda_U(\mathbf{A})$, if λ_j is detectable w.r.t. the pair $(\mathbf{A}, \mathbf{C}^{(\mathcal{P}_i)})$, i.e., if $\lambda_j \in \mathcal{O}_i$. The set of all source mobile agents for $\lambda_j \in \Lambda_U(\mathbf{A})$ is denoted by \mathcal{S}_j .

In words, an agent is a source mobile agent for an unstable or marginally stable eigenvalue of the system if such an eigenvalue is detectable w.r.t. the collective measurements available from the agent's sensing locations.⁵ Recall that the set of locally detectable eigenvalues of agent i is denoted by \mathcal{O}_i , and let $\mathcal{UO}_i = \text{sp}(\mathbf{A}) \setminus \mathcal{O}_i$. Our goal in the subsequent sections will be to design resilient state estimation algorithms that allow agent i to estimate the components of $\mathbf{z}[k]$ corresponding to the eigenvalues in \mathcal{UO}_i . Such algorithms,

⁵ Since we are considering system matrices with distinct eigenvalues, an eigenvalue is detectable w.r.t. the pair $(\mathbf{A}, \mathbf{C}^{(\mathcal{P}_i)})$ if and only if it is detectable w.r.t. $(\mathbf{A}, \mathbf{C}^{(i_r)})$, for some $i_r \in \mathcal{P}_i$. The 'only if' part of the statement may not be true for system matrices with repeated eigenvalues.




Fig. 1 A scalar unstable plant is monitored by a clique of $m + 2$ agents, where s_1 and s_2 are the only source agents. A single adversary corrupting either of the two sources can render the distributed state estimation problem impossible, irrespective of the choice of algorithm.

however, need to be complemented by incorporating adequate redundancy in not only the communication network topology, but also in the measurement structure.⁶ A simple illustration of this fact is as follows.

Example 1 Consider a scalar unstable plant monitored by a clique of $m + 2$ agents, as depicted in Figure 1. Agents s_1 and s_2 are the only agents with access to non-zero measurements, i.e., they are the source agents for this system. Although this network is fully connected, the presence of a single adversarial agent makes it impossible for **any** algorithm to guarantee estimation of $x[k]$ for every regular agent. Specifically, if the adversary compromises one of the two source agents, then it can behave in a way that makes it impossible for the non-source agents to distinguish between two different state trajectories of the system, due to the conflicting information from the two source agents.⁷

In our prior work [49], we proposed an algorithm that made use of certain directed acyclic subgraphs in addressing the resilient distributed state estimation problem (using static sensors over time-invariant communication graphs). To understand the properties of such subgraphs, let $\Omega_U(\mathbf{A}) \subseteq \Lambda_U(\mathbf{A})$ denote the set of eigenvalues of \mathbf{A} for which $\mathcal{V} \setminus \mathcal{S}_j$ is non-empty.

Definition 2 (Mode Estimation Directed Acyclic Graph (MEDAG)) Consider a mode $\lambda_j \in \Omega_U(\mathbf{A})$. Suppose there exists a spanning subgraph $\mathcal{G}_j = (\mathcal{V}, \mathcal{E}_j)$ of \mathcal{G} with the following properties for all f -local sets \mathcal{A} in \mathcal{G} (and corresponding $\mathcal{R} = \mathcal{V} \setminus \mathcal{A}$).

⁶ This is one of the key differences of our present formulation with the resilient consensus literature. In the latter setting, there is no external state that needs to be tracked, and [73, 74] have shown that making the network sufficiently connected suffices to facilitate resilient consensus.

⁷ Details of such an attack strategy can be found in [50]. For centralized systems where f sensors are compromised, [75, 76] have shown that for recovering the state of the system asymptotically, the system must remain detectable after the removal of any $2f$ sensors.

- (i) If $i \in \{\mathcal{V} \setminus \mathcal{S}_j\} \cap \mathcal{R}$, then $|\mathcal{N}_i^{(j)}| \geq 2f + 1$, where $\mathcal{N}_i^{(j)} = \{l | (l, i) \in \mathcal{E}_j\}$ represents the neighborhood of agent i in \mathcal{G}_j .
- (ii) There exists a $T_j \in \mathbb{N}_+$ such that \mathcal{R} can be partitioned into the sets $\{\mathcal{L}_0^{(j)}, \dots, \mathcal{L}_{T_j}^{(j)}\}$, where $\mathcal{L}_0^{(j)} = \mathcal{S}_j \cap \mathcal{R}$, and if $i \in \mathcal{L}_q^{(j)}$ (where $1 \leq q \leq T_j$), then $\mathcal{N}_i^{(j)} \cap \mathcal{R} \subseteq \bigcup_{r=0}^{q-1} \mathcal{L}_r^{(j)}$. Furthermore, $\mathcal{N}_i^{(j)} = \emptyset, \forall i \in \mathcal{L}_0^{(j)}$.

Then, we call \mathcal{G}_j a *Mode Estimation Directed Acyclic Graph (MEDAG)* for $\lambda_j \in \Omega_U(\mathbf{A})$.

Although the concept of a MEDAG was originally developed for a network with static nodes, we can instead view the MEDAG as a special information flow structure between the mobile agents in our present context. With this in mind, we elaborate on the key properties of this graph structure. First, it should be noted that T_j and the levels $\mathcal{L}_0^{(j)}$ to $\mathcal{L}_{T_j}^{(j)}$ can vary across different f -local sets. For a given f -local set \mathcal{A} , we say a regular agent $i \in \mathcal{L}_m^{(j)}$ “belongs to level m ”, where the levels indicate the distances of the regular agents from the source set \mathcal{S}_j , in the baseline communication graph \mathcal{G} . Consider a state $z^{(j)}[k]$ that grows exponentially with time. To estimate such a state despite adversarial actions, there must exist a secure medium of information flow from the corresponding source set \mathcal{S}_j to the rest of the mobile agents (who do not patrol regions providing information about $z^{(j)}[k]$). A MEDAG \mathcal{G}_j is a subgraph with properties that fulfill this requirement. Specifically, the first property of a MEDAG indicates that every regular agent $i \in \mathcal{V} \setminus \mathcal{S}_j$ has at least $(2f + 1)$ neighbors in the subgraph \mathcal{G}_j , while the second property indicates that all its regular neighbors in such a subgraph belong to levels strictly preceding its own level. Our estimation scheme (described later) requires an agent i to listen to *only* its neighbors in $\mathcal{N}_i^{(j)}$ for estimating $z^{(j)}[k]$. The second property of a MEDAG then indicates that agents in level m only use estimates of regular agents in levels 0 to $m - 1$ for recovering $z^{(j)}[k]$.

Before proceeding further, we need to understand the properties of the baseline communication graph \mathcal{G} that guarantee the existence of a MEDAG $\mathcal{G}_j, \forall \lambda_j \in \Omega_U(\mathbf{A})$. To this end, we require the following definitions and result from [50].

Definition 3 (r -reachable set) For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a set $\mathcal{S} \subset \mathcal{V}$, and an integer $r \in \mathbb{N}_+$, \mathcal{S} is an r -reachable set if there exists an $i \in \mathcal{S}$ such that $|\mathcal{N}_i \setminus \mathcal{S}| \geq r$.

Definition 4 (strongly r -robust graph w.r.t. \mathcal{S}_j) For $r \in \mathbb{N}_+$ and $\lambda_j \in \Omega_U(\mathbf{A})$, a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is *strongly r -robust w.r.t. to the set of source agents \mathcal{S}_j* , if for any non-empty subset $\mathcal{C} \subseteq \mathcal{V} \setminus \mathcal{S}_j$, \mathcal{C} is r -reachable.




Fig. 2 The above figure depicts an LTI process with 6 states that satisfies Assumption 1. The red crosses indicate sensing locations, and the mobile agents are represented by the blue circles. The dashed rectangles are used to demarcate the patrolling regions of the agents. A directed path from one rectangle to another indicates that every agent in the former can transmit information to every agent in the latter in the baseline communication graph. Modes λ_1 and λ_2 of the system are detectable w.r.t. the measurements available from the left-most rectangular region. Thus, agents 1, 2 and 3 act as the source mobile agents for modes λ_1 and λ_2 . Source agents for the other modes can be described similarly.

Lemma 2 Consider an eigenvalue $\lambda_j \in \Omega_U(\mathbf{A})$. The graph \mathcal{G} contains a MEDAG \mathcal{G}_j if and only if \mathcal{G} is strongly $(2f + 1)$ -robust w.r.t. \mathcal{S}_j .

Given a $\lambda_j \in \Omega_U(\mathbf{A})$, there might be more than one subgraph that satisfies the definition of a MEDAG \mathcal{G}_j . In [49], we proposed a distributed algorithm that allowed each node i to identify the sets $\mathcal{N}_i^{(j)}, \forall \lambda_j \in \mathcal{UO}_i$, by explicitly constructing a *specific* MEDAG \mathcal{G}_j for each $\lambda_j \in \mathcal{UO}_i$. In this paper, we assume that these MEDAGs have already been constructed during a design phase using such an algorithm. In other words, we work under the assumption that each agent i is informed of the set $\mathcal{N}_i^{(j)}, \forall \lambda_j \in \mathcal{UO}_i$. It will be important to keep in mind that the sets $\mathcal{N}_i^{(j)}$ are time-invariant as they correspond to specific MEDAGs in the time-invariant baseline communication graph \mathcal{G} ; however, we will allow for the possibility that each regular agent $i \in \mathcal{R}$ can only communicate with a subset of the agents in $\mathcal{N}_i^{(j)}$ at each time-step, due to communication losses and agent mobility.

5 Resilient Distributed State Estimation over Time-varying Networks

In this section, we develop an algorithm that enables each regular mobile agent to estimate its locally undetectable portion subject to arbitrary adversarial attacks and intermittent communication losses. We will focus our attention on communication losses that satisfy the following criterion.




Fig. 3 Consider the scenario described by Figure 2, and a 1-local adversarial model (i.e., $f=1$). The MEDAG \mathcal{G}_1 corresponding to mode λ_1 is depicted by (a). The subsequent figures show one example sequence of how the communication pattern evolves over time. Agents 1, 2 and 3 are the source agents for mode λ_1 , and are at level 0 of \mathcal{G}_1 , agents 4, 5 and 6 are at level 1 of \mathcal{G}_1 , and agents 7, 8 and 9 are at level 2 of \mathcal{G}_1 . With $T = 2$, the figure illustrates a communication pattern satisfying Assumption 2 for mode λ_1 . Specifically, $\forall k \geq 2$, the union graph $\bigcup_{\tau=0}^2 \mathcal{G}[k - \tau]$ contains the MEDAG \mathcal{G}_1 .

Assumption 2. *There exists $T \in \mathbb{N}_+$ such that $\forall k \geq T$, $\bigcup_{\tau=0}^T \mathcal{G}[k - \tau]$ contains the MEDAG \mathcal{G}_j for each $\lambda_j \in \Omega_U(\mathbf{A})$.*

Assumption 2 places certain design constraints on the patrols of the agents. In particular, based on the definition of a MEDAG in Definition 2, the following observations can be made regarding Assumption 2. (i) The patrols should be sufficiently informative, i.e., there should exist at least $(2f + 1)$ source mobile agents for each $\lambda_j \in \Omega_U(\mathbf{A})$. Coupled with Proposition 1, this requires at least $(2f + 1)$ mobile agents to periodically visit each informative location (i.e., locations providing information regarding unstable modes of the system) in the geographical region. (ii) The patrols should ensure that the mobile agents are able to exchange information sufficiently often, and along sufficiently many different paths (i.e., as specified by the MEDAG structure). Thus, loosely speaking, periodic patrols that lead to densely connected communication networks over time are key to our subsequent development. In a latter section (namely Section 7), we will consider packet dropping scenarios that do not necessarily satisfy Assumption 2.

An illustration of some of the concepts introduced in the previous section is shown in Figure 2. Based on the scenario described in Figure 2, a communication loss pattern satisfying Assumption 2 is illustrated in Figure 3. From Figure 3, we notice that $\mathcal{G}[k]$ may not contain the specific MEDAGs constructed during the design phase for some (or all) k , thereby precluding direct use of the technique developed in [49]. However, such MEDAGs will be preserved in the union graph over the interval $[k - T, k]$, $\forall k \geq T$. For our subsequent development, we assume that all estimates being transmitted by regular agents are properly time-stamped. We now propose the following algorithm.

Let $\hat{z}_i^{(j)}[k]$ denote the estimate of $z^{(j)}[k]$ maintained by agent i at time-step k . Then, for each $\lambda_j \in \mathcal{UO}_i$, a regular agent i updates $\hat{z}_i^{(j)}[k]$ in the following manner.

1. At every time-step k , agent i collects the *most recent* estimate of $z^{(j)}[k]$ received from each agent $l \in \mathcal{N}_i^{(j)}$, along with the corresponding time-stamp $\phi_{il}[k] \in \mathbb{N}$. It then evaluates the delay $\tau_{il}[k] = k - \phi_{il}[k]$ and computes the quantity $\bar{z}_{il}^{(j)}[k] \triangleq \lambda_j^{\tau_{il}[k]} \hat{z}_l^{(j)}[k - \tau_{il}[k]]$.⁸ Prior to receiving the first estimate from an agent $l \in \mathcal{N}_i^{(j)}$, the value $\bar{z}_{il}^{(j)}[k]$ is maintained at 0 by agent i .⁹
2. The values $\bar{z}_{il}^{(j)}[k]$ are sorted from largest to smallest; subsequently, the largest f and the smallest f of such values are discarded (i.e., $2f$ values are discarded in all) and $\hat{z}_i^{(j)}[k]$ is updated as

$$\hat{z}_i^{(j)}[k + 1] = \lambda_j \left(\sum_{l \in \mathcal{M}_i^{(j)}[k]} w_{il}^{(j)}[k] \bar{z}_{il}^{(j)}[k] \right), \quad (14)$$

where $\mathcal{M}_i^{(j)}[k] \subset \mathcal{N}_i^{(j)}$ represents the set of agents whose (potentially) delayed estimates are used by agent i at time-step k after the removal of the $2f$ aforementioned values. Agent i assigns the consensus weight $w_{il}^{(j)}[k]$ to agent l at time-step k for estimating the component of the state corresponding to the eigenvalue λ_j . The weights $w_{il}^{(j)}[k]$ are non-negative and satisfy $\sum_{l \in \mathcal{M}_i^{(j)}[k]} w_{il}^{(j)}[k] = 1$, $\forall \lambda_j \in \mathcal{UO}_i$, and $\forall k \in \mathbb{N}$.

We refer to the above algorithm as the Sliding Window Local-Filtering based Resilient Estimation (SW-LFRE) algorithm. We comment on certain features of

⁸ For notational simplicity, while considering the eigenvalue λ_j , we drop the superscript ‘ j ’ on the time-stamp $\phi_{il}[k]$ and the delay $\tau_{il}[k]$.

⁹ If agent i receives an estimate without a time-stamp from some agent in $\mathcal{N}_i^{(j)} \cap \mathcal{A}$, it simply assigns a value of 0 to such an estimate (without loss of generality). Note that based on Assumption 2, agent i is guaranteed to receive a time-stamped estimate from every regular agent l in $\mathcal{N}_i^{(j)}$ at least once over every interval of the form $[k - T, k]$, $\forall k \geq T$, i.e., for each $l \in \mathcal{N}_i^{(j)} \cap \mathcal{R}$, $\bar{z}_{il}^{(j)}[k]$ will necessarily be of the form $\lambda_j^{\tau_{il}[k]} \hat{z}_l^{(j)}[k - \tau_{il}[k]]$, $\forall k \geq T$.

this algorithm and then proceed to analyze its convergence properties.

Remark 5 Like the LFRE algorithm in [50], the SW-LFRE algorithm also relies on a two-stage filtering strategy. Specifically, the first stage of filtering corresponds to a regular agent $i \in \mathcal{V} \setminus \mathcal{S}_j$ listening to only its neighbors $\mathcal{N}_i^{(j)} \subseteq \mathcal{N}_i$ in the MEDAG \mathcal{G}_j . This operation ensures a uni-directional flow of information from the source agents \mathcal{S}_j (some of whom might also be adversarial) to the rest of the network. The second stage of filtering requires agent i to discard certain extreme values received from agents in $\mathcal{N}_i^{(j)}$. Whereas the first stage of filtering is specific to our distributed state estimation approach, the second stage of filtering is similar to the W-MSR algorithm employed in the resilient consensus literature [51], [52]. A key point of difference between the LFRE and SW-LFRE algorithms is that in the latter algorithm, at each time-step k , agent i needs to process the most recent (potentially) delayed state estimate received from each neighbor in $\mathcal{N}_i^{(j)}$. Accounting for such delayed state estimates (of an unstable dynamics) requires us to make careful modifications to the design and analysis of the LFRE algorithm. Also, unlike the LFRE algorithm, implementing the SW-LFRE algorithm requires the agents to possess adequate memory, for reasons that follow from the above discussion.

Remark 6 Our approach does not require the agents to possess a priori knowledge of the value of T in Assumption 2.

Remark 7 Our results will continue to hold if in step 2 of the SW-LFRE algorithm, agent i simply uses the median value of $\bar{z}_{il}^{(j)}[k]$, $l \in \mathcal{N}_i^{(j)}$, in the update rule (14). Although this can reduce computation, the present approach offers a degree of freedom in choosing the weights $w_{il}^{(j)}[k]$, that can be potentially leveraged to account for issues like noise.

Remark 8 As alluded to earlier in the introduction, this communication-loss model offers the adversaries the additional opportunity of sending false information regarding the time-stamps of their estimates.¹⁰ Nevertheless, as we establish in the next section, our proposed algorithm is immune to such misbehavior.

6 Analysis of the SW-LFRE Algorithm

The following is the main result of this section.

¹⁰ In other words, due to false time-stamp information, the quantity $\hat{z}_i^{(j)}[k - \tau_{il}[k]]$ may not represent the true estimate of an adversarial agent l at time $(k - \tau_{il}[k])$. Thus, we resort to a slight abuse of notation here.

Theorem 1 *Given an LTI system (1) and a measurement model (2), suppose all the conditions stated in Proposition 1 are met. Additionally, let the baseline communication graph \mathcal{G} be strongly $(2f+1)$ -robust w.r.t. \mathcal{S}_j , $\forall \lambda_j \in \Omega_U(\mathbf{A})$, and let the communication patterns satisfy Assumption 2. Then, the proposed SW-LFRE algorithm guarantees the following despite the actions of any set of f -local Byzantine adversaries.*

- (**Asymptotic Stability**) *Each regular agent $i \in \mathcal{R}$ can asymptotically estimate the state of the plant, i.e., $\lim_{k \rightarrow \infty} \|\hat{\mathbf{x}}_i[k] - \mathbf{x}[k]\| = 0$, $\forall i \in \mathcal{R}$, where $\hat{\mathbf{x}}_i[k]$ is the estimate of $\mathbf{x}[k]$ maintained by agent i .*
- (**Rate of Convergence**) *Let $e_i^{(j)}[k] = \hat{z}_i^{(j)}[k] - z^{(j)}[k]$ denote the error in estimation of the component $z^{(j)}[k]$ by a regular agent $i \in \mathcal{V} \setminus \mathcal{S}_j$. Suppose agent i belongs to level q of the MEDAG \mathcal{G}_j . Then, there exist constants $\beta^{(j)} > 0$ and $\gamma^{(j)} \in (0, 1)$, such that the estimation error $e_i^{(j)}[k]$ can be bounded as follows $\forall k \geq (T + 1)q$:*

$$|e_i^{(j)}[k]| \leq \beta^{(j)} \left(\frac{|\lambda_j|}{\gamma^{(j)}} \right)^{q(T+1)} (\gamma^{(j)})^k. \quad (15)$$

Proof. For each regular agent i , the state vector $\mathbf{z}[k]$ can be partitioned into the components $\mathbf{z}_{\mathcal{O}_i}[k]$ and $\mathbf{z}_{\mathcal{U}\mathcal{O}_i}[k]$ that correspond to the locally detectable and locally undetectable eigenvalues, respectively, of agent i . Since the conditions stated in Proposition 1 are met, agent i can asymptotically recover $\mathbf{z}_{\mathcal{O}_i}[k]$ via persistent patrolling and by implementing the observer given by (7) and (8). It remains to show that agent i can recover $\mathbf{z}_{\mathcal{U}\mathcal{O}_i}[k]$, or in other words, for each $\lambda_j \in \mathcal{U}\mathcal{O}_i$, we need to prove that $\lim_{k \rightarrow \infty} |\hat{z}_i^{(j)}[k] - z^{(j)}[k]| = 0$. Equivalently, we show that for each $\lambda_j \in \Omega_U(\mathbf{A})$, every regular agent $i \in \mathcal{V} \setminus \mathcal{S}_j$ can asymptotically recover $z^{(j)}[k]$.

Consider any f -local adversarial set \mathcal{A} and let $\mathcal{R} = \mathcal{V} \setminus \mathcal{A}$. Consider an eigenvalue $\lambda_j \in \Omega_U(\mathbf{A})$. Since $\mathcal{E}[k] \subseteq \mathcal{E}$ for all k , Assumption 2 can hold only if the baseline graph \mathcal{G} contains \mathcal{G}_j . The latter follows from the conditions of the theorem and Lemma 2. Next, based on Assumption 2, note that for all $k \geq T$, the union of the graphs over the interval $[k - T, k]$ contains the MEDAG \mathcal{G}_j . Recall that the sets $\{\mathcal{L}_0^{(j)}, \mathcal{L}_1^{(j)}, \dots, \mathcal{L}_q^{(j)}, \dots, \mathcal{L}_{T_j}^{(j)}\}$ form a partition of the set of regular agents \mathcal{R} in such a MEDAG. We prove the desired result by inducting on the level number q . For $q = 0$, $\mathcal{L}_0^{(j)} = \mathcal{S}_j \cap \mathcal{R}$ by definition, and hence all agents in level 0 can estimate $z^{(j)}[k]$ asymptotically by virtue of Proposition 1. Next, consider a regular agent i in $\mathcal{L}_1^{(j)}$ and let $e_i^{(j)}[k] \triangleq \hat{z}_i^{(j)}[k] - z^{(j)}[k]$. We first analyze the SW-LFRE update rule (14). To this end, at each time-step k , let the

neighbor set $\mathcal{N}_i^{(j)}$ of agent i be partitioned into the sets $\mathcal{U}_i^{(j)}[k]$, $\mathcal{M}_i^{(j)}[k]$ and $\mathcal{J}_i^{(j)}[k]$, where $\mathcal{U}_i^{(j)}[k]$ and $\mathcal{J}_i^{(j)}[k]$ contain f agents each, with the highest and the lowest values of $\bar{z}_{il}^{(j)}[k]$ respectively, and $\mathcal{M}_i^{(j)}[k]$ contains the remaining agents in $\mathcal{N}_i^{(j)}$. At any instant k , we can either have $\mathcal{M}_i^{(j)}[k] \cap \mathcal{A} = \emptyset$ or $\mathcal{M}_i^{(j)}[k] \cap \mathcal{A} \neq \emptyset$. In the former case, all agents in $\mathcal{M}_i^{(j)}[k]$ belong to $\mathcal{L}_0^{(j)} = \mathcal{S}_j \cap \mathcal{R}$. Now consider the latter case when agent i uses values transmitted by adversarial agents in its update rule. It follows from the SW-LFRE algorithm, the f -locality of the adversary model, and the fact that $|\mathcal{N}_i^{(j)}| \geq (2f+1)$, that for each $l \in \mathcal{M}_i^{(j)}[k] \cap \mathcal{A}$, there exists an agent $u \in \mathcal{U}_i^{(j)}[k]$ and an agent $v \in \mathcal{J}_i^{(j)}[k]$ such that both $u, v \in \mathcal{L}_0^{(j)}$, and $\bar{z}_{iv}^{(j)}[k] \leq \bar{z}_{il}^{(j)}[k] \leq \bar{z}_{iu}^{(j)}[k]$, i.e., $\bar{z}_{il}^{(j)}[k]$ can be expressed as a convex combination of $\bar{z}_{iu}^{(j)}[k]$ and $\bar{z}_{iv}^{(j)}[k]$.¹¹ Based on the above discussion and (14), it follows that for all k , $\hat{z}_i^{(j)}[k+1]$ belongs to the convex hull formed by $\lambda_j \bar{z}_{il}^{(j)}[k]$, $l \in \mathcal{L}_0^{(j)}$. Specifically, there exist weights $\bar{w}_{il}^{(j)}[k]$ such that $\sum_{l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \bar{w}_{il}^{(j)}[k] = 1$, and

$$\hat{z}_i^{(j)}[k+1] = \lambda_j \left(\sum_{l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \bar{w}_{il}^{(j)}[k] \bar{z}_{il}^{(j)}[k] \right). \quad (16)$$

Since $\sum_{l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \bar{w}_{il}^{(j)}[k] = 1$, and $z^{(j)}[k+1] = \lambda_j z^{(j)}[k]$ based on (4), simple manipulations imply

$$z^{(j)}[k+1] = \lambda_j \left(\sum_{l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \bar{w}_{il}^{(j)}[k] \lambda_j^{\tau_{il}[k]} z^{(j)}[k - \tau_{il}[k]] \right). \quad (17)$$

Based on Assumption 2 and step 1 of the SW-LFRE update rule, we have that for all $k \geq T$, $\bar{z}_{il}^{(j)}[k] = \lambda_j^{\tau_{il}[k]} \hat{z}_i^{(j)}[k - \tau_{il}[k]]$, $l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}$. Subtracting (17) from (16), we then obtain the following error dynamics for all $k \geq T$:

$$e_i^{(j)}[k+1] = \lambda_j \left(\sum_{l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \bar{w}_{il}^{(j)}[k] \lambda_j^{\tau_{il}[k]} e_l^{(j)}[k - \tau_{il}[k]] \right). \quad (18)$$

Noting that the weights $\bar{w}_{il}^{(j)}[k]$ are non-negative, the delay terms $\tau_{il}[k]$ are upper bounded by T for $l \in$

$\mathcal{N}_i^{(j)} \cap \mathcal{R}$, λ_j satisfies $|\lambda_j| \geq 1$, and using the triangle inequality, we obtain the following based on (18) for all $k \geq T$:

$$|e_i^{(j)}[k+1]| \leq |\lambda_j|^{(T+1)} \left(\sum_{l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \bar{w}_{il}^{(j)}[k] |e_l^{(j)}[k - \tau_{il}[k]]| \right). \quad (19)$$

For every $l \in \mathcal{L}_0^{(j)}$, since $e_l^{(j)}[k]$ converges exponentially¹² based on Proposition 1, there exist constants $c_l^{(j)} > 0$ and $\gamma_l^{(j)} \in (0, 1)$ such that $|e_l^{(j)}[k]| \leq c_l^{(j)} (\gamma_l^{(j)})^k$, for all $k \in \mathbb{N}$. Let $\beta^{(j)} \triangleq \max_{l \in \mathcal{L}_0^{(j)}} c_l^{(j)}$ and $\gamma^{(j)} \triangleq \max_{l \in \mathcal{L}_0^{(j)}} \gamma_l^{(j)}$. Then, we obtain the following inequality based on (19) for all $k \geq T$:

$$|e_i^{(j)}[k+1]| \leq |\lambda_j|^{(T+1)} \beta^{(j)} (\gamma^{(j)})^{(k-T)}, \quad (20)$$

where we have used the fact that $\sum_{l \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \bar{w}_{il}^{(j)}[k] = 1$. Thus, we obtain (15) for $q = 1$, implying exponential stability of the error dynamics (18) for all agents in level 1, since $\gamma^{(j)} \in (0, 1)$.

Suppose exponential stability holds for agents in all levels from 0 to q (where $1 \leq q \leq T_j - 1$). It is easy to see that the result holds for all agents in $\mathcal{L}_{q+1}^{(j)}$ as well, by noting that (i) a regular agent $i \in \mathcal{L}_{q+1}^{(j)}$ has $\mathcal{N}_i^{(j)} \cap \mathcal{R} \subseteq \bigcup_{r=0}^q \mathcal{L}_r^{(j)}$, and (ii) any value $\bar{z}_{il}^{(j)}[k]$ used by agent i in the update rule (14) lies in the convex hull formed by $\bar{z}_{iu}^{(j)}[k]$, $u \in \bigcup_{r=0}^q \mathcal{L}_r^{(j)}$. Based on the induction hypothesis, exponential stability can then be argued using the same reasoning as the $q = 1$ case. Verifying (15) is a matter of straightforward algebra. \square

We now focus on the impact of bounded communication delays between mobile agents when the communication graph among them remains unchanged over time. Here, by a bounded communication delay we imply that if $(i, j) \in \mathcal{E}[k]$ and $i, j \in \mathcal{R}$, then any estimate transmitted by agent i to agent j at time-step k is received by agent j no later than time-step $k + T$, for some $T \in \mathbb{N}_+$. It turns out that the arguments used in the proof of Theorem 1 can be used almost identically to analyze the impact of bounded communication delays (potentially random, time-varying) in the presence of adversaries, for time-invariant communication networks. We formalize this observation below.

Corollary 1 *Given an LTI system (1) and a measurement model (2), suppose all the conditions stated in*

¹¹ Explicit dependence of u, v on the parameters represented by i, j, l and k is not shown to avoid cluttering of the exposition.

¹² Although we only establish asymptotic stability of the error dynamics in Proposition 1, verifying exponential stability is fairly straightforward, and hence, not explicitly proven.

Proposition 1 are met. Additionally, let $\mathcal{G}[k] = \mathcal{G} \forall k$, where \mathcal{G} is strongly $(2f + 1)$ -robust w.r.t. $\mathcal{S}_j, \forall \lambda_j \in \Omega_U(\mathbf{A})$. Furthermore, let communication delays between any pair of regular agents in \mathcal{G} be bounded by some $T \in \mathbb{N}_+$. Then, the proposed SW-LFRE algorithm provides identical guarantees as in *Theorem 1*.

We summarize the implications of *Theorem 1* and *Corollary 1* in the following remarks.

Remark 9 For a given mode $\lambda_j \in \Omega_U(\mathbf{A})$, the constants $\beta^{(j)}$ and $\gamma^{(j)}$ are indicative of the time of convergence of the estimation errors (corresponding to the state $z^{(j)}[k]$) of the source agents \mathcal{S}_j . In other words, these constants are dictated by the observer equations (7) and (8). While $\beta^{(j)}$ encompasses the effects of the patrol time-periods and the initial state estimation errors, $\gamma^{(j)}$ essentially represents the slowest rate of convergence among the source agents \mathcal{S}_j . In this context, *Theorem 1* relates the time of convergence of the non-source agents to that of the source agents, and shows how the instability of the mode under consideration, the maximum delay T , and distances from the source set (captured by the different levels of the MEDAG) feature in such a relation.

Remark 10 When it comes to addressing the effect of network-induced delays in the context of distributed state estimation (using static sensors and in the absence of any adversarial attacks), there is limited literature that provides any theoretical guarantees. Approaches such as the one outlined in [77] typically seek to account for delays and packet-drops by formulating LMI-based conditions that do not in general provide any graph-theoretic insights. In contrast, the proposed SW-LFRE algorithm allows one to deal with bounded delays in a much simpler manner by exploiting the uni-directional flow of information that is inherent to our approach. We conjecture that our results pertaining to bounded delays will carry over to more general system and measurement models, such as those considered in [38, 40, 41].

7 Resilient Distributed State Estimation over Analog Erasure Channels

In the previous section, we analyzed a communication failure model where the patterns described by *Assumption 2* recurred in a deterministic manner. In this sense, the time-varying communication patterns considered in the previous section can be attributed to agent movements. In contrast, the focus of this section will be to analyze time-varying communication patterns that are a consequence of imperfections in the communication channel. We model such imperfections as random

packet drops that can potentially lead to a violation of the conditions stated in *Assumption 2*. To isolate the impact of random packet drops, we will assume for the remainder of this section that the patrols have been designed to ensure that the baseline communication topology is retained at every time-step in the absence of packet drops. In other words, in this section, we assume that such packet drops are the sole cause of communication losses. The task of analyzing time-varying communication patterns that are a consequence of both agent movements and random packet drops is left as future work.

With these points in mind, we now explore a scenario where each communication link between the mobile agents is modeled as an analog erasure channel as defined in [78]. In particular, the transmission of information across any link $(i, j) \in \mathcal{E}$ is governed by a random process $\xi_{ij}[k]$ that is memoryless, i.e., $\xi_{ij}[k]$ is i.i.d. over time. Furthermore, across space, the packet dropping processes over different links are independent. For any k , the random variable $\xi_{ij}[k]$ follows a Bernoulli fading distribution, i.e., $\xi_{ij}[k] = 0$ with erasure probability p and $\xi_{ij}[k] = 1$ with probability $(1 - p)$; the implications of $\xi_{ij}[k]$ assuming the values 0 and 1 will be discussed shortly.

Our **objective** in this section will be to design an estimation protocol that guarantees mean-square stability of the estimation error dynamics for each regular agent, in the following sense.

Definition 5 (Mean-Square Stability (MSS)) The estimation error dynamics of the regular agents is said to be mean-square stable if $\lim_{k \rightarrow \infty} E[\|\mathbf{e}_i[k]\|^2] = 0, \forall i \in \mathcal{R}$, where $\mathbf{e}_i[k] = \hat{\mathbf{x}}_i[k] - \mathbf{x}[k]$, and the expectation is taken with respect to the packet dropping processes $\xi_{ij}[k], (i, j) \in \mathcal{E}$.

7.1 Channels with no delay

We first consider the case where $\xi_{ij}[k] = 1$ implies that any data packet transmitted by agent i at time k is received perfectly by agent j at time k , and when $\xi_{ij}[k] = 0$, such a packet is dropped completely. For this model, we propose a simple algorithm described as follows.

For each $\lambda_j \in \mathcal{UO}_i$, a regular agent i updates its estimate of $z^{(j)}[k]$ in the following manner.¹³

¹³ Unlike the SW-LFRE algorithm developed in Section 5, the algorithm we propose here is memoryless, i.e., at each time-step, an agent acts only on the information that it acquires (via measurements and from neighboring agents) at that time-step. We do this primarily to simplify the analysis.

- At each time-step k , if it receives estimates from at least $(2f+1)$ agents in $\mathcal{N}_i^{(j)}$, it runs the LFRE algorithm, i.e., it removes the largest f and the smallest f estimates $\hat{z}_l^{(j)}[k]$, $l \in \mathcal{N}_i^{(j)}$ and updates $\hat{z}_i^{(j)}[k]$ as

$$\hat{z}_i^{(j)}[k+1] = \lambda_j \left(\sum_{l \in \mathcal{M}_i^{(j)}[k]} w_{il}^{(j)}[k] \hat{z}_l^{(j)}[k] \right), \quad (21)$$

where the set $\mathcal{M}_i^{(j)}[k]$ and the weights $w_{il}^{(j)}[k]$ are defined as in the description of the SW-LFRE algorithm in Section 5. Otherwise, it runs open-loop as follows:

$$\hat{z}_i^{(j)}[k+1] = \lambda_j \hat{z}_i^{(j)}[k]. \quad (22)$$

The above algorithm provides the following guarantees.

Theorem 2 *Given an LTI system (1), and a measurement model (2), suppose all the conditions stated in Proposition 1 are met. Let the baseline communication graph \mathcal{G} be strongly $(mf+1)$ -robust w.r.t. $\mathcal{S}_j, \forall \lambda_j \in \Omega_U(\mathbf{A})$, where $m \in \mathbb{N}_+$ is a constant. For each $(i, j) \in \mathcal{E}$, let $\xi_{ij}[k]$ be a Bernoulli packet dropping process with erasure probability p , that is i.i.d. over time and independent of packet dropping processes over other links. Suppose $m \geq 3$ and that the following is true.¹⁴*

$$\rho^2 \bar{p} < 1, \quad (23)$$

where ρ is the spectral radius of \mathbf{A} , and

$$\bar{p} \triangleq 1 - \sum_{l=(2f+1)}^{(m-1)f+1} \binom{(m-1)f+1}{l} (1-p)^l p^{(m-1)f+(1-l)}. \quad (24)$$

Then, the estimation algorithm described by the update rules (21) and (22) guarantees mean-square stability in the sense of Definition 5, despite the actions of any f -local set of Byzantine adversaries.

Proof. Note that the packet dropping processes do not affect the estimation of the locally detectable portions of the state, i.e., each regular mobile agent i can recover $\mathbf{z}_{\mathcal{O}_i}[k]$ asymptotically since the conditions stated in Proposition 1 are satisfied. Consider any f -local adversarial set \mathcal{A} and let $\mathcal{R} = \mathcal{V} \setminus \mathcal{A}$. Consider an eigenvalue $\lambda_j \in \Omega_U(\mathbf{A})$. Since \mathcal{G} is strongly $(mf+1)$ -robust w.r.t. \mathcal{S}_j , a trivial extension of Lemma 2 implies that there exists a MEDAG \mathcal{G}_j with $|\mathcal{N}_i^{(j)}| \geq (mf+1), \forall i \in \{\mathcal{V} \setminus \mathcal{S}_j\} \cap \mathcal{R}$. We induct on the level numbers q of such a MEDAG \mathcal{G}_j present in the baseline communication graph \mathcal{G} . Let i be an agent in level 1. Let $\mathcal{I}_i[k]$ be an

indicator random variable¹⁵ such that $\mathcal{I}_i[k] = 1$ if agent i uses the update rule (22) and $\mathcal{I}_i[k] = 0$ if agent i uses the update rule (21). To make the presentation clear, we make the following assumption. Suppose agent i receives estimates from more than $(2f+1)$ agents in $\mathcal{N}_i^{(j)}$ at a certain time-step k . Then, after removing $2f$ estimates based on the LFRE algorithm, it listens to *only a single agent* l picked arbitrarily from $\mathcal{M}_i^{(j)}[k]$, while running (21).¹⁶ Combining (21) and (22), we obtain

$$\hat{z}_i^{(j)}[k+1] = \lambda_j \left(\mathcal{I}_i[k] \hat{z}_i^{(j)}[k] + (1 - \mathcal{I}_i[k]) \hat{z}_l^{(j)}[k] \right), \quad (25)$$

where $l \in \mathcal{M}_i^{(j)}[k]$.¹⁷ It is easy to see that the error $e_i^{(j)}[k] = \hat{z}_i^{(j)}[k] - z^{(j)}[k]$ follows the dynamics:

$$e_i^{(j)}[k+1] = \lambda_j \left(\mathcal{I}_i[k] e_i^{(j)}[k] + (1 - \mathcal{I}_i[k]) e_l^{(j)}[k] \right). \quad (26)$$

Defining $\sigma_i^{(j)}[k] \triangleq E[(e_i^{(j)}[k])^2]$, and using (26), we obtain:

$$\begin{aligned} \sigma_i^{(j)}[k+1] &= \lambda_j^2 E[\mathcal{I}_i^2[k]] \sigma_i^{(j)}[k] + \lambda_j^2 E[(1 - \mathcal{I}_i[k])^2] (e_l^{(j)}[k])^2 \\ &\quad + \underbrace{2\lambda_j^2 E[(\mathcal{I}_i[k] - \mathcal{I}_i^2[k]) (e_i^{(j)}[k]) (e_l^{(j)}[k])]}_{g^{[k]}}, \\ &\leq \lambda_j^2 p_i^{(j)}[k] \sigma_i^{(j)}[k] + \lambda_j^2 (1 - p_i^{(j)}[k]) \max_{r \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \sigma_r^{(j)}[k], \\ &\leq (\lambda_j^2 \bar{p}) \sigma_i^{(j)}[k] + \lambda_j^2 \left(\max_{r \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \sigma_r^{(j)}[k] \right), \end{aligned} \quad (27)$$

where $p_i^{(j)}[k]$ is the probability that $\mathcal{I}_i[k] = 1$. We now justify each of the above steps. For arriving at the first equality, we used the fact that $e_i^{(j)}[k]$ is independent of $\mathcal{I}_i[k]$ for any $i \in \mathcal{R}$, based on the update rules (21) and (22), and the nature of the packet dropping processes. Notice that when agent l is adversarial, it may have precise knowledge of the number of packets received by agent i at time-step k ; the estimate $\hat{z}_l^{(j)}[k]$ it transmits to agent i might then be influenced by such knowledge. Regardless of this fact, whenever $l \in \mathcal{M}_i^{(j)}[k]$, based on the LFRE update rule (21) and the f -locality of the adversarial model, it follows from arguments identical to those in Theorem 1 that $e_l^{(j)}[k]$ can be expressed as a convex combination of $e_u^{(j)}[k]$ and $e_v^{(j)}[k]$, for some $u, v \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}$. Since such agents are regular, their errors at time k are independent of $\mathcal{I}_i[k]$. The above discussion combined with the

¹⁵ To avoid cluttering the exposition, we drop the superscript ‘ j ’ on $\mathcal{I}_i[k]$ and certain other terms throughout the proof, since they can be inferred from context.

¹⁶ The result continues to hold for the general update rule (21).

¹⁷ The set $\mathcal{M}_i^{(j)}[k]$ is not well-defined when $\mathcal{I}_i[k] = 1$. For such a case, l can be taken to be any node in the set $\mathcal{N}_i^{(j)} \cap \mathcal{R}$.

¹⁴ The choice of $m \geq 3$ is justified later in Remark 13.

fact that $g[k] = 0$ (since $\mathcal{I}_i[k]$ is an indicator random variable) leads to the second inequality in (27). Observe that since $\mathcal{L}_0^{(j)} = \mathcal{S}_j \cap \mathcal{R}$, it follows from Proposition 1 that $\lim_{k \rightarrow \infty} \max_{r \in \mathcal{N}_i^{(j)} \cap \mathcal{L}_0^{(j)}} \sigma_r^{(j)}[k] = 0$.

For arriving at the final inequality, we first note that $p_i^{(j)}[k]$ can potentially vary over time and across different agents since the adversarial agents are allowed to behave arbitrarily. In particular, a compromised agent may choose not to transmit estimates even if all outgoing communication links from such an agent are intact. Thus, since it is impossible to exactly compute $p_i^{(j)}[k]$, we instead seek to upper-bound it. To this end, note that the probability that $\mathcal{I}_i[k] = 0$, i.e., the probability that agent i receives estimates from at least $(2f + 1)$ agents in $\mathcal{N}_i^{(j)}$ at time k , is lower bounded by the probability that it receives estimates from at least $(2f + 1)$ agents in $\mathcal{N}_i^{(j)} \cap \mathcal{R}$ at time k . The latter probability can be further lower bounded by $(1 - \bar{p})$ (where \bar{p} is given by (24)) by noting that $|\mathcal{N}_i^{(j)} \cap \mathcal{R}| \geq ((m-1)f + 1)$ based on the f -locality of the fault model. In light of the above discussion, we have $p_i^{(j)}[k] \leq \bar{p}$, leading to the last inequality in (27). Finally, equation (23) implies that $\lambda_j^2 \bar{p} < 1$, and in turn guarantees that $\lim_{k \rightarrow \infty} \sigma_i^{(j)}[k] = 0$, based on Input to State Stability (ISS) and the foregoing discussion.

Suppose $\lim_{k \rightarrow \infty} \sigma_i^{(j)}[k] = 0$ for all agents in levels 0 to q . Consider an agent $i \in \mathcal{L}_{q+1}^{(j)}[k]$. Its error dynamics can be bounded as in (27), with $g[k] = 0$ for reasons discussed above, and $e_l^{(j)}[k] = \alpha_{il}^{(j)}[k]e_u^{(j)}[k] + (1 - \alpha_{il}^{(j)}[k])e_v^{(j)}[k]$, for some $\alpha_{il}^{(j)}[k] \in [0, 1]$, and some $u, v \in \bigcup_{r=0}^q \mathcal{L}_r^{(j)}$. The last argument follows from the LFRE update rule (21). Consider the term $E[(e_u^{(j)}[k])(e_v^{(j)}[k])]$ appearing in $\sigma_l^{(j)}[k]$. Since $\sigma_u^{(j)}[k]$ and $\sigma_v^{(j)}[k]$ converge to 0 based on the induction hypothesis, we can use the Cauchy-Schwartz inequality to bound $E[(e_u^{(j)}[k])(e_v^{(j)}[k])]$ as follows:

$$E[(e_u^{(j)}[k])(e_v^{(j)}[k])] \leq \sqrt{\sigma_u^{(j)}[k]\sigma_v^{(j)}[k]}. \quad (28)$$

This implies $\lim_{k \rightarrow \infty} \sigma_l^{(j)}[k] = 0, \forall l \in \mathcal{M}_i^{(j)}[k]$. The rest of the proof can be completed following similar arguments as the $q = 1$ case. \square

The term \bar{p} appearing in (23) and (24) can be interpreted as the effective packet drop/erasure probability for the problem under study. With this in mind, the implications of the above result are described as follows.

Remark 11 (Increasing ‘network robustness’ reduces ‘effective packet drop probability’) Given




Fig. 4 Plot illustrating how the effective packet drop probability \bar{p} can be reduced by increasing the level of robustness m . For this example, $\rho = 2$ and $f = 3$.

knowledge of the spectral radius ρ of \mathbf{A} , an upper-bound f on the number of adversaries in the neighborhood of any regular agent, and the erasure probability p of the communication medium, suppose we are faced with the problem of designing a communication topology that guarantees mean-square stability in the sense of Definition 5. Theorem 2 provides an answer to this problem by quantitatively relating our notion of ‘strong-robustness’ in Definition 4 to the effective packet drop probability \bar{p} . For instance, as shown in Figure 4, given the parameters ρ, f and p , one can generate a plot for $\rho^2 \bar{p}$ offline, and choose m to satisfy the MSS criterion $\rho^2 \bar{p} < 1$. Subsequently, one can proceed to design a network that is strongly $(mf + 1)$ -robust w.r.t. $\mathcal{S}_j, \lambda_j \in \Omega_U(\mathbf{A})$. It is easy to verify that \bar{p} is monotonically increasing in p , and monotonically decreasing in m . In other words, for a fixed ρ and f , one can tolerate higher erasure probabilities p by increasing the robustness parameter m . In the context of the mobile agents that we are considering in this paper, this corresponds to changing the patrols of the individual agents so that they encounter the other agents in such a way that the baseline communication network (containing the set of all agent interactions over time) is sufficiently robust.

Remark 12 Note that when $f = 0$, i.e., in the absence of adversaries, equation (23) reduces to $\rho^2 p < 1$. This condition is reminiscent of the MSS criterion for remote stabilization of an LTI system over a Bernoulli packet dropping channel [79]. This observation can be explained by viewing the contribution due to the LFRE update (21) (that helps stabilize the error dynamics (26)) as an analogue of the stabilizing input in the remote stabilization problem.

Remark 13 We now justify the need for $m \geq 3$ in Theorem 2. Suppose the network is strongly $(mf + 1)$ -robust with $m \leq 2$. In this case, each adversarial agent

may follow the simple strategy of never transmitting its estimate. If the adversaries compromise f agents in some set $\mathcal{N}_i^{(j)}$, where $i \in \mathcal{R}$ and $\lambda_j \in \mathcal{U}\mathcal{O}_i$, then such a strategy might cause the regular agent i to run open-loop forever based on the algorithm described by the update rules (21) and (22). Instead of running open-loop, suppose that if a regular agent i does not hear from some neighbor in $\mathcal{N}_i^{(j)}$ at time k , it assigns a value of 0 to the corresponding estimate, and then employs the LFRE update rule (21). Such an approach will in general not work either, due to the following reason. Unlike the communication loss model studied in Section 5, where each regular agent was **guaranteed** to eventually receive estimates from ‘enough’ regular neighbors, no such guarantees can be claimed for the analog erasure channel model studied here. Thus, while strongly $(2f+1)$ -robust networks sufficed in Section 5, the choice of $m \geq 3$ is in fact necessary in the present context for achieving MSS based on our specific approach. However, $m = 2$ does suffice for certain variants of the analog erasure channel model, as we discuss next.

7.2 Channels with erasure and delay

In this section, we consider a variant of the analog erasure channel that accounts for the presence of random delays. To this end, let $(i, j) \in \mathcal{E}$, and let $\mathbf{v}[k]$ be a message transmitted by agent i to agent j at time-step k . Then, a channel with delay and erasure causes agent j to receive the following message:

$$\mathbf{r}[k] = \xi_{ij}[k]\mathbf{v}[k] + (1 - \xi_{ij}[k])\mathbf{v}[k - \tau_{ij}[k]], \quad (29)$$

where $\xi_{ij}[k]$ is the memory-less packet dropping process described earlier, $\tau_{ij}[k] \in \mathbb{N}_+$ is a random delay satisfying $1 \leq \tau_{ij}[k] \leq T$, and $T \in \mathbb{Z}_{>0}$. In words, the channel output $\mathbf{r}[k]$ is either equal to the present channel input $\mathbf{v}[k]$ with probability $(1 - p)$, or equal to a delayed channel input with probability p , where the delay is upper bounded by some positive constant T . It should be noted that the erasure channel model considered here is a generalization of the erasure channel with delay in [78], where the delays are constant. For this model, we have the following result.

Proposition 2 *Given an LTI system (1) and a measurement model (2), suppose all the conditions stated in Proposition 1 are met. Let the baseline communication graph \mathcal{G} be strongly $(2f+1)$ -robust w.r.t. $\mathcal{S}_j, \forall \lambda_j \in \Omega_U(\mathbf{A})$. Let each communication link of \mathcal{G} be modeled as a channel with delay and erasure as described by equation (29). Then, the SW-LFRE algorithm provides identical guarantees as Theorem 1, with probability 1.*

Proof. The proof follows from the following simple observation. Based on the channel model (29), note that for each $\lambda_j \in \Omega_U(\mathbf{A})$, every regular agent $i \in \mathcal{V} \setminus \mathcal{S}_j$ is guaranteed to receive a state estimate that is at most T time-steps delayed, from each of its regular neighbors in $\mathcal{N}_i^{(j)}$, at every time-step $k, \forall k \geq T$. This corresponds to a special case of the bounded delay model in Corollary 1, and the result thus follows. \square

8 Simulation Study

In this section, we substantiate our theoretical results via a detailed simulation study. To this end, we first describe our general simulation setup, and then propose a simple patrolling strategy that meets the design specifications laid down by the theoretical results developed in this paper.

Simulation Model and a Simple Patrolling Strategy:

Our general setup is as follows. We consider a geographical region partitioned into K cells, with r sensing locations distributed within each cell. A dynamical process evolves within each cell, and the processes across different cells are assumed to be decoupled. A network of mobile agents is deployed over this region; the task of each agent is to gain global situational awareness by estimating the state of the dynamical process in every cell (i.e., not only the state of the process within the cell that it is patrolling). Suppose (i) the sensing capabilities and movement patterns of any given agent are limited to a single cell, i.e., an agent can persistently visit all the sensing locations within its own cell, but not cross over to adjacent cells, and (ii) the communication radius is such that each agent in a given cell can nominally communicate with all agents in each adjacent cell. However, such communication is subject to random packet drops based on the model described in Section 7.1.

Based on the above setting, and the development in Section 7, it is clear that we need at least $(3f+1)$ mobile agents patrolling each cell, and hence, a total of $(3f+1)K$ mobile agents monitoring the overall region. The nominal communication patterns can be viewed as arising from a preferential attachment type mechanism (where each new agent attaches itself to $(3f+1)$ existing agents), and ensure that the baseline communication graph is strongly $(3f+1)$ -robust w.r.t. every relevant source set [50].¹⁸ Thus, for the scenario described above, the overall patrolling strategy simply boils down

¹⁸ The need for strong $(3f+1)$ -robustness in the baseline network was provided in Remark 13, and will also be justified explicitly via simulations.




Fig. 5 This figure illustrates the general simulation setup comprising of a region partitioned into 4 cells. Each cell has 10 randomly distributed measurement locations represented by the red crosses, and 4 mobile agents patrolling within the cell. Cell 1 is patrolled by agents 1–4, Cell 2 is patrolled by agents 5–8, Cell 3 is patrolled by agents 9–12, and Cell 4 is patrolled by agents 13–16. The edges depict the baseline communication graph, representing the fact that each agent in a given cell can nominally communicate with all agents in adjacent cells. However, all communication links are subject to random packet drops.

to persistent periodic intra-cell patrols (with the periods chosen appropriately based on Proposition 1). In what follows, we demonstrate how such patrols complement the estimation techniques developed in this paper by considering a specific instance of the above setup.

Let $K=4$, and $r=10$. In each of the 4 cells, the 10 sensing locations are randomly distributed. For all our simulations, we consider a 1-local Byzantine adversary model. There are $(3f+1) = 4$ mobile agents within each cell. As detailed above, each agent executes a periodic patrol that persistently visits each of the sensing locations within its own cell, without crossing over to adjacent cells. This scenario is depicted in Figure 5, where Cell 1 is patrolled by agents 1–4, Cell 2 is patrolled by agents 5–8, Cell 3 is patrolled by agents 9–12, and Cell 4 is patrolled by agents 13–16. Figure 5 depicts the baseline communication graph pattern where each agent in Cell 1 can communicate with all agents in Cells 2 and 4 (adjacent cells), but with no agent in Cell 3; the communication patterns for the other cells can be described similarly. Finally, we emphasize that the baseline communication graph in Fig. 5 may not be retained entirely at each time-step, due to random packet drops.

Information Flow Patterns: Based on the discussion in Section 3, since each mobile agent persistently visits every sensing location within its own cell, it acts




Fig. 6 This figure illustrates the MEDAG (in the baseline communication graph) that dictates the flow of information for estimating the states of the process evolving in Cell 1. The agents in Cell 1 act as source agents (level 0), the agents in Cells 2 and 4 are at level 1, while agents in Cell 3 are at level 2 of this MEDAG. A directed edge from Cell i to Cell j indicates that each agent in Cell i has a directed edge to each agent in Cell j in the baseline graph.




Fig. 7 This figure illustrates the information flow patterns for the estimation of the different inter-cell processes. Each cell is represented by a single square node, and the dashed edges indicate how agents communicate among themselves for estimating the various states. Specifically, figures (a), (b), (c) and (d) represent the way estimates are processed for recovering the states in Cells 1, 2, 3 and 4, respectively.

as a source of information for all the states associated with its own cell. Figure 6 shows the MEDAG that dictates the flow of information among the mobile agents for estimating the states evolving in Cell 1. Specifically, this MEDAG will have all agents in Cell 1 at level 0 (source agents), all agents in Cells 2 and 4 at level 1, and all agents in Cell 3 at level 2. This information flow is compactly represented in Figure 7(a). In a similar manner, MEDAGs for Cells 2, 3 and 4 are compactly represented via Figures 7(b)-(d), respectively.

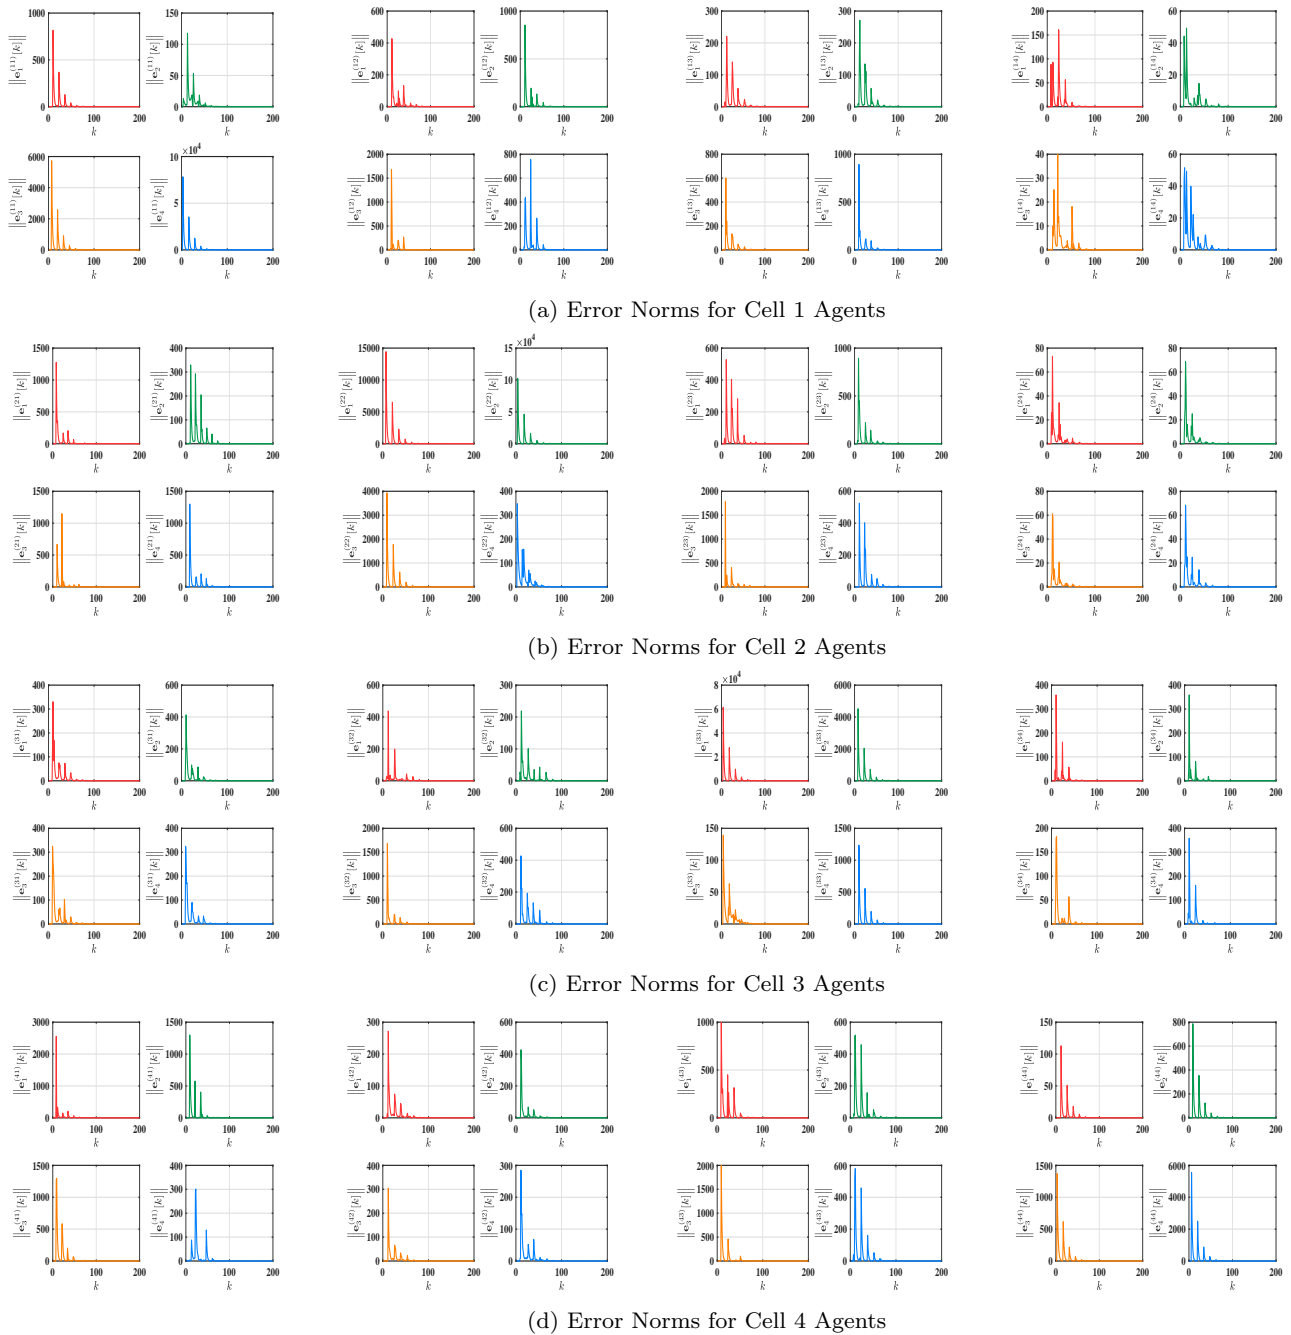


Fig. 8 Plots of the estimation error norms for the scenario described in Section 8.1. The erasure probability is set to its maximum allowable value 0.32 (to ensure MSS) and agent 5 in Cell 2 acts as an adversarial agent. The notation $\|e_q^{(ij)}[k]\|$ is used to indicate the estimation error norm of agent q in Cell i , w.r.t. the dynamics in Cell j , at time-step k .

8.1 Tracking Unstable Dynamical Processes

In our first simulation study based on the setup described above, we consider identical unstable dynamical processes evolving in each cell. Processes in different cells start out from different initial conditions and hence evolve differently across cells. For each cell, we

consider a 20-dimensional LTI system with a diagonal system matrix (any system matrix satisfying Assumption 1 can be diagonalized). The system has real, distinct eigenvalues distributed uniformly from 0.5 to 1.2, i.e., the spectral radius of the system is $\rho = 1.2$. Since all eigenvalues are of the same sign and distinct (i.e., of distinct magnitude), Lemma 1, Proposition 1 and Foot-

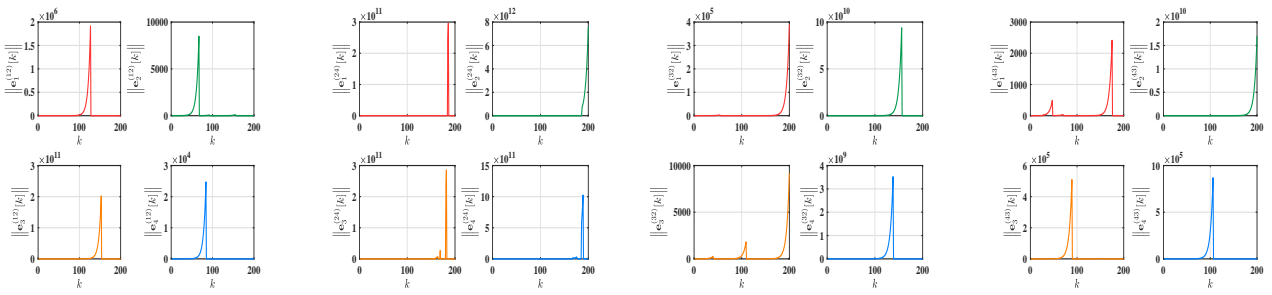


Fig. 9 Illustration of certain inter-cell estimation error plots for the scenario described in Section 8.1, corresponding to high erasure probability. For this example, the erasure probability is $p = 0.8$. However, all agents are regular.

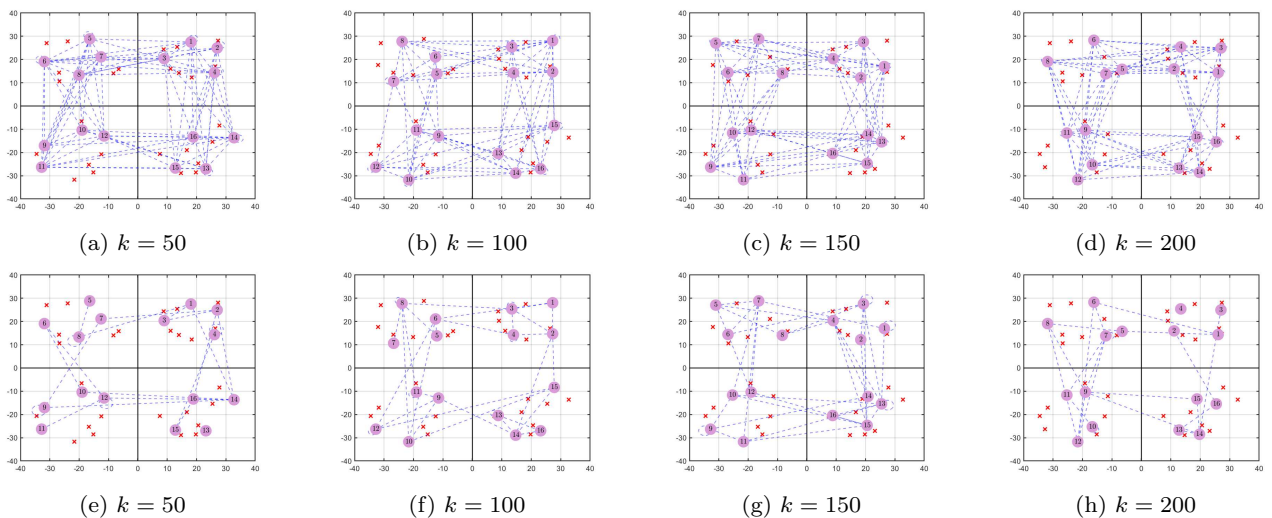


Fig. 10 Illustration of dynamically changing communication links at different time-steps k . Figures (a)-(d) represent the case when the erasure probability p is 0.32. Figures (e)-(h) represent the case when the erasure probability p is 0.8.

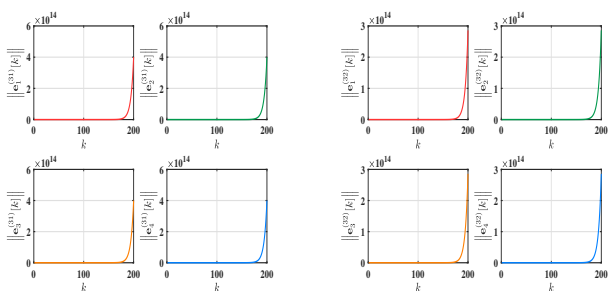


Fig. 11 Illustration of the need for strong- $(3f + 1)$ robustness, as pointed out in Remark 13. Adversarial agents choosing not to transmit estimates (omissive attacks) can cause the estimation errors to grow unbounded with time.

note 4 indicate that the time-period with which each location is visited can be chosen arbitrarily (as long as they induce feasible periodic patrols as defined in Section 3). The 10 measurement locations within each cell

are numbered from 1 to 10, and the measurement vector at the i -th location is taken to be a row vector with entries of 1 at the $(2i - 1)$ -th and $2i$ -th positions. Thus, precisely two eigenvalues of the system are observable from each location within each cell. Locations in Cells 1, 2, 3 and 4 are visited with time-periods 13, 14, 14, and 14, respectively (each location within a given cell is visited by all mobile agents patrolling that cell with the same time-period). The observer gains are designed based on the procedure outlined in the proof of Proposition 1. The initial conditions for states in Cells 1, 2, 3 and 4 are 0.07, 0.05, 0.03 and 0.01, respectively (all the states in the same cell start at the same value). All state estimates are initialized from zero. We now study various aspects of the problem.

Effect of Random Packet Drops: We first focus on the impact of random packet drops coupled with adversarial attacks. For this case, since $\rho = 1.2$, $f = 1$, and the baseline communication graph is strongly $(mf + 1)$ -

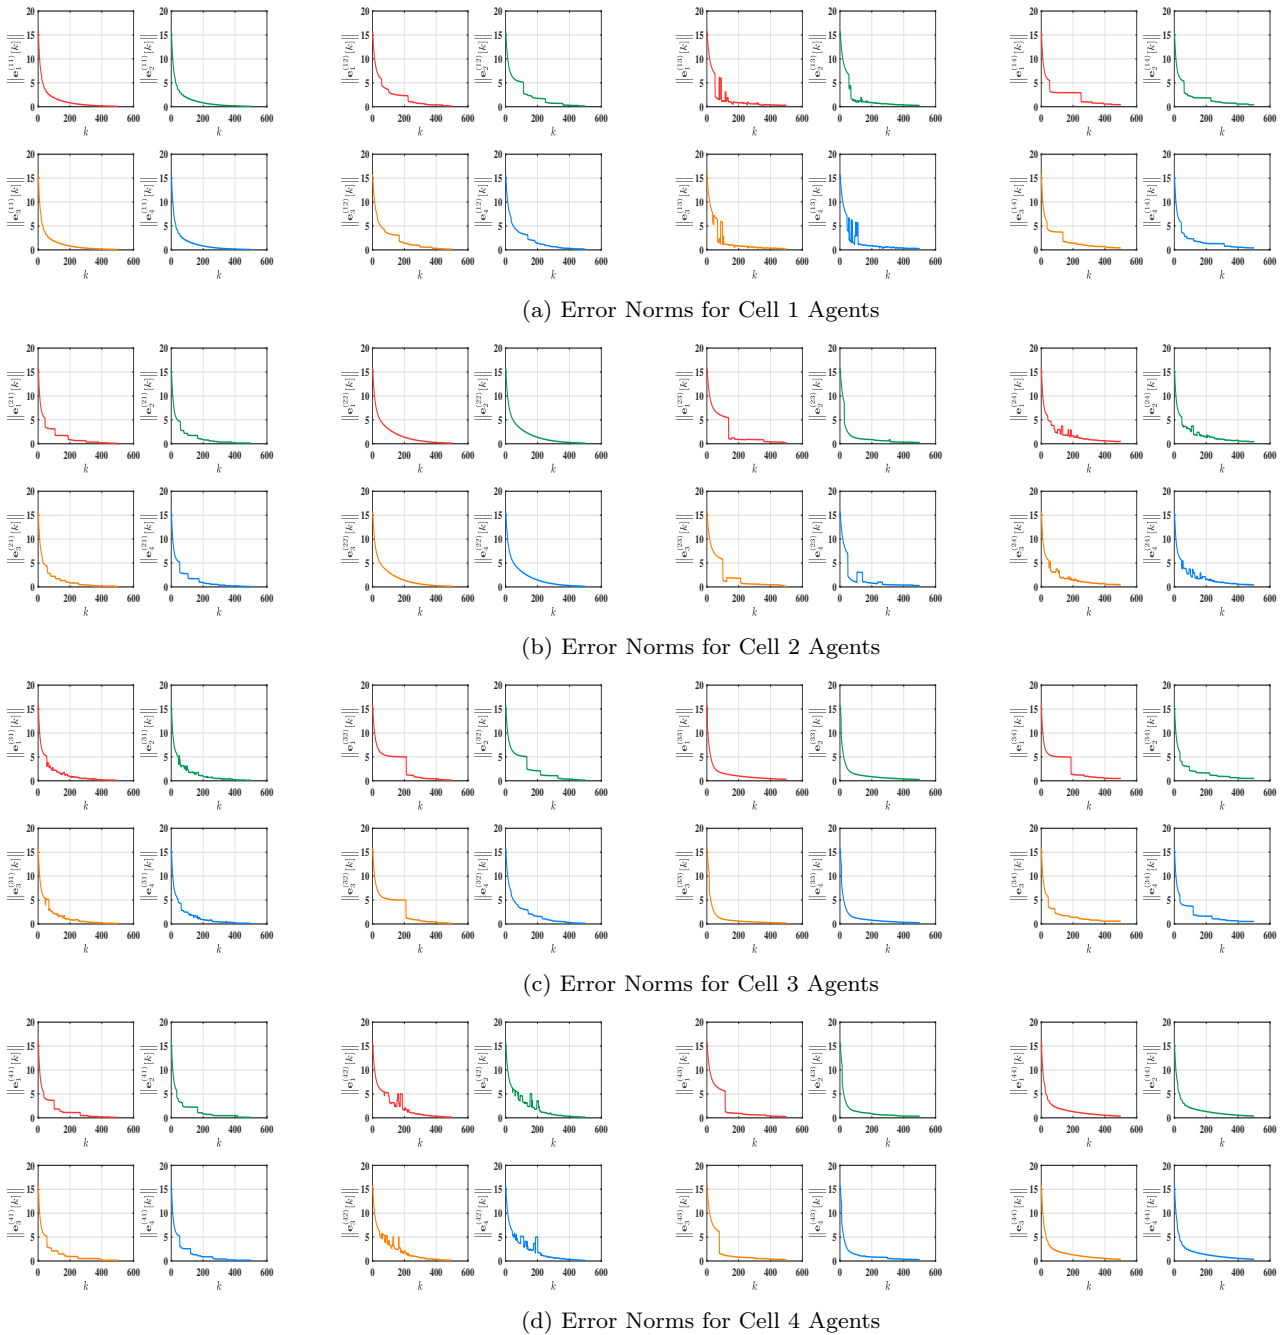


Fig. 12 Plots of the estimation error norms for the scenario described in Section 8.2. The erasure probability is set to $p = 0.8$, and agent 5 in Cell 2 acts as an adversarial agent.

robust with $m = 3$, the condition for MSS stated in Theorem 2 indicates that the erasure probability can be at most 0.32. The erasure probability is set to this maximum value, and agent 5 in Cell 2 is considered to be adversarial. Specifically, at any given time-step, if the communication link from agent 5 to some other agent i is intact, then agent 5 does the following. It adjusts its state estimate transmitted to agent i to be

equal and opposite to the sum of the other state estimates being used by agent i at that time-step. This action is intended to keep the state estimates of agent i static. Note that none of the other agents know that agent 5 is adversarial. We will discuss the specific repercussions of such an adversarial attack in the example discussed in Section 8.2. For now, we focus on the effect of the erasure probability p .




Fig. 13 Illustration of a Byzantine attack for the diffusion dynamics model considered in Section 8.2. Agent 5 in Cell 2 is a Byzantine adversary, and transmits estimates that keep the estimates of the regular agents static. Although the network is strongly $(3f + 1)$ -robust w.r.t. every source set and there are no packet drops, a non-resilient state estimation algorithm fails to achieve asymptotic stability of the error dynamics (i.e., the errors do not converge to zero).

The simulation results for the case described above are shown in Figure 8. The notation $\|e_q^{(ij)}[k]\|$ is used to indicate the estimation error norm of agent q in Cell i , w.r.t. the dynamics in Cell j , at time-step k . Despite adversarial attacks and random packet drops, we see that the error plots corroborate the theory developed in this paper. The decaying spikes in the intra-cell error norm plots are a consequence of the periodic motions of the mobile agents. The inter-cell error norm plots inherit this trend coupled with the effect of random packet drops and adversarial injections.

The effect of a high erasure probability is shown in Figure 9. For this case, the erasure probability is $p = 0.8$. To isolate the effect of random packet drops, we assume that all agents are regular for this specific illustration. Even so, some of the inter-cell estimation error norms grow unbounded with time, as shown in Figure 9. Figure 10 illustrates the dynamically changing communication links between the different agents, for both high and low erasure probabilities.

Effect of Omissive Attacks: We illustrate the need for $(3f + 1)$ -strong robustness, as discussed in Remark 13. Consider a scenario where agents in Cell 3 cannot communicate with agents in Cell 4. Furthermore, suppose each agent in Cell 3 can communicate with only agents 5–7 in Cell 2. This communication pattern leads to a $(2f + 1)$ -strongly robust network w.r.t. the agents in Cell 1. Suppose agent 5 is compromised and follows the simple strategy of never transmitting estimates to agents in Cell 3. For this scenario, the estimation error plots for agents in Cell 3 are shown in Figure 11. These plots justify our claim that $(3f + 1)$ -strong robustness is necessary for achieving MSS based on the algorithm described in Section 7.1.

8.2 Tracking a Diffusion Process

As pointed out in the Introduction, one of the main applications of the theory developed in this paper is environmental monitoring. In particular, one might be interested in monitoring the concentration of a physical quantity (such as a gas) that evolves based on a spatio-temporal process. Such processes are commonly described by the Laplacian dynamics in continuous-time [68–70]. For our purpose, we consider a discrete-time version of the Laplacian dynamics for which the system matrix is of the form $\mathbf{I} - \epsilon\mathbf{L}$, where ϵ is a small number that is indicative of the sampling period, and \mathbf{L} is the graph Laplacian matrix induced by the sensing locations. For our simulations, we take $\epsilon = 0.01$. We consider decoupled diffusion processes evolving in each cell. For each cell, the Laplacian matrix induced by the sensing locations within the cell is constructed as follows. Locations that are within a certain Euclidean distance (taken to be 15 distance units) are considered to be connected. This connectivity pattern defines the adjacency matrix between the sensing locations, and in turn defines the Laplacian dynamics. Based on our choice of the threshold distance, the sensing locations within each cell induce an undirected connected graph. Since the Laplacian matrix corresponding to an undirected connected graph has precisely one eigenvalue at zero, it follows that the dynamics matrix $\mathbf{I} - \epsilon\mathbf{L}$ associated with each cell has precisely one eigenvalue at 1, and all other eigenvalues non-negative with magnitude strictly less than 1 (the latter statement follows from basic properties of a Laplacian matrix). In other words, the dynamics in each cell are marginally stable. Within each cell, we assume that the i -th component of the state is measured by the i -th location. All other parameters (the number and positions of the sensing locations, the time-periods of the patrols, the adversarial model etc.) are the same as the first simulation example. Based on our analysis in Section 7.1, for marginally stable systems, MSS is guaranteed as long as the erasure probability is strictly less than 1. To validate this claim, we consider an erasure probability as high as 0.8. The simulation results for this case as shown in Figure 12 corroborate the developed theory.

Effect of Byzantine Attacks: Finally, to emphasize the importance of the resilient filtering techniques developed in the paper, we consider a scenario when there are no packet drops, and the baseline communication graph is retained at every time-step. Our goal will be to illustrate that although the communication network satisfies the robustness conditions needed for countering adversarial behavior, the absence of a resilient state es-

timization algorithm prevents some of the regular agents from tracking the true state. To this end, we need to consider a state estimation algorithm that does not account for adversarial behavior, such as those proposed in [38, 40, 41, 45]. For the sake of illustration, we consider the algorithm developed in [40]. Unlike the two-stage filtering techniques described in Sections 5 and 7.1, the approach in [40] focuses only on maintaining a uni-directional flow of information from the source mobile agents to the non-source mobile agents. In other words, for some $\lambda_j \in \mathcal{UC}_i$, an agent $i \in \mathcal{R}$ simply takes a convex combination of the state estimates of its neighbors belonging to the set $\mathcal{N}_i^{(j)}$, for updating $\hat{z}_i^{(j)}[k]$. In the absence of adversarial attacks or communication losses, convergence of such an update rule can be established using similar arguments as in [40]. For the present scenario, suppose agent 5 in Cell 2 is adversarial. Since adversarial agents are assumed to be omniscient, agent 5 is aware of the state estimates being transmitted to each agent at every time-step, and the weights placed on such estimates. For any regular agent i , if agent 5 belongs to the set $\mathcal{N}_i^{(j)}$, it simply transmits an estimate to agent i that cancels out the effect of the other estimates corresponding to the set $\mathcal{N}_i^{(j)}$. This attack has the effect of keeping the state estimate $\hat{z}_i^{(j)}[k]$ static. Illustration of such an attack is shown in Figure 13 where asymptotic stability of the error dynamics is not achieved.

9 The Path to Implementation

The framework for distributed state estimation with mobile agents that we have established in this paper encompasses a wide range of real-world considerations, including time-varying measurement models (due to agent mobility), time-varying communication links between agents (due to probabilistic packet drops, agent mobility, and limited communication ranges), and Byzantine agents (which are capable of capturing both benign failures and malicious compromises by attackers). The detailed simulations that we provided in the previous section incorporate all of these features, and demonstrate that the theoretical guarantees that we provide do, in fact, hold under the stated assumptions on the underlying dynamical processes and the multi-agent system. The generality of our framework, coupled with the validation provided by our simulations, indicates that our approach holds promise for real-world implementations.

To achieve a full real-world implementation, the first step (and challenge) will be to deploy a set of mobile robots, each executing a patrol in such a way that the baseline graph satisfies the conditions that we have identified for our resilient distributed state estimation

algorithms (for instance, $(3f+1)$ -strong robustness with respect to each set of source agents for the random packet drop scenario). The patrolling strategy described in our simulation section is a promising (and relatively simple) candidate for implementation, as it only requires the entire region to be partitioned into disjoint cells, and to have $(3f+1)$ agents executing periodic patrols within each cell, with the ability to communicate with all agents in adjacent cells (notwithstanding intermittent packet drops). Given an upper bound on the transmission range of the individual robots, the sizes of the cells can be scaled down appropriately in order to satisfy this communication constraint.

The second step will be to create a dynamical process for the mobile agents to monitor. A simple diffusion process (e.g., a gas spreading over a region starting from a given point) like the one studied in Section 8.2 would be a promising candidate, as this process can be approximated via an LTI system of the form (1), where the system matrix \mathbf{A} is a graph Laplacian [69, 70]. The concentration of gas at different points in the region can be sensed by ground-based sensors [80], which transmit their information to the mobile agents that pass by during their patrols. Several disjoint diffusion processes can be instantiated, one for each cell in the region. Once these processes are instantiated and the mobile agent patrols are implemented, the switched linear observers described in Section 3 can be implemented to test the ability of each agent to asymptotically track the state of the system within its cell (i.e., the concentrations of gas at various points in the cell).

Once each agent is able to estimate the gas concentrations in its own cell, the third step will be to enable the agents to exchange information with agents in neighboring cells. This can be done by constructing the MEDAGs described in Section 4, and having the agents run the SW-LFRE algorithm for the deterministic failure model considered in Section 5, or the algorithm described in Section 7.1 for the random packet drop scenario. The MEDAGs and the estimation algorithms can be programmed into the agents before deployment. As a first test, the ability of the algorithm to enable global state estimation can be verified in the absence of any malicious agents (i.e., by having all agents participate in the algorithm as programmed). Once that has been verified, one of the agents can then be programmed to deviate from the algorithm in an arbitrary manner (e.g., by broadcasting random or large values to its neighbors at certain time-steps). The ability of the algorithm to provide resilience to such malicious behavior can then be verified.

Aside from additional hardware challenges that would be inherent to any real-world implementation, we an-

anticipate that the above pathway (and associated milestones) will lead to a successful demonstration of the resilient distributed state estimation framework that we have established in this paper.

10 Summary and Future Work

In this paper, we studied the problem of estimating the state of a dynamical process evolving over a certain region with a team of mobile agents. We assumed that each agent visits a subset of sensing locations via a periodic patrol; at each sensing location, the agent obtains a measurement of a portion of the state of the system. We showed how to construct observers for each agent to asymptotically recover the locally detectable parts of the system state, and formulated state exchange and update rules for agents to recover the locally undetectable parts of the state. Our algorithms provide resilience to a certain number of worst-case (Byzantine) agents under certain conditions on the baseline network topology. Our framework encompasses intermittent observations due to agent patrols, time-varying communication networks due to packet drops and agent mobility, and Byzantine behavior by the agents. We illustrated the efficacy of our approach via detailed simulations, and described a path to a real-world implementation.

In addition to implementing our algorithm, a natural next step would be to formulate patrol strategies that provide the required robustness conditions on the baseline network while also allowing the individual agents to recover the state. We provided one example of such a patrol strategy in our simulations, and anticipate that the insights from that strategy (along with results about robustness of different networks provided in [58, 59, 62, 81, 82]) can be leveraged to develop additional classes of patrolling strategies. We plan on studying the effect of incorporating memory (like the SW-LFRE approach) in the state estimation algorithm described in Section 7.1 for the random packet drop scenario. While we established MSS based on the memoryless algorithm described in Section 7.1, it would be interesting to investigate if other forms of stochastic stability (such as almost sure convergence) can be established for the analog erasure channel model considered in Section 7.1, subject to adversarial attacks. For more general system and observation models than the ones considered in this paper, the problem of designing a persistent patrol that guarantees stability of the estimation error dynamics is challenging. Recent results along this direction are available in [83]. Finally, while the delays considered in this paper were network-induced, one can also investigate the impact of measurement delays (in a distributed setup) leveraging recent results on this topic [84, 85].

References

1. Dan Goodin. There is a new way to take down drones, and it doesn't involve shotguns. *arsTechnica*, October 2016.
2. Courtney Kube. Russia has figured out how to jam U.S. drones in Syria, officials say. *NBC News*, Apr. 2018.
3. Lev Semenovich Gandin. Objective analysis of meteorological fields. *Israel Program for Scientific Translations*, 242, 1963.
4. Noel Cressie. The origins of kriging. *Mathematical Geology*, 22(3):239–252, 1990.
5. Mohammed Abazeed, Norshiela Faisal, Suleiman Zubair, and Adel Ali. Routing protocols for wireless multimedia sensor network: a survey. *Journal of Sensors*, 2013.
6. Lijun Xie and Xi Zhang. 3D clustering-based camera wireless sensor networks for maximizing lifespan with minimum coverage rate constraint. In *Proc. of the IEEE Global Communications Conference (GLOBECOM)*, pages 298–303, 2013.
7. Ryan N Smith, Mac Schwager, Stephen L Smith, Burton H Jones, Daniela Rus, and Gaurav S Sukhatme. Persistent ocean monitoring with underwater gliders: Adapting sampling resolution. *Journal of Field Robotics*, 28(5):714–741, 2011.
8. Matthew Dunbabin, Jonathan M Roberts, Kane Usher, and Peter Corke. A new robot for environmental monitoring on the Great Barrier Reef. In *Proceedings of the Australasian Conference on Robotics & Automation*. Australian Robotics & Automation Association, 2004.
9. David Higdon. A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environmental and Ecological Statistics*, 5(2):173–190, 1998.
10. Suman Srinivasan, Haniph Latchman, John Shea, Tan Wong, and Janice McNair. Airborne traffic surveillance systems: video surveillance of highway traffic. In *Proceedings of the ACM 2nd International Workshop on Video Surveillance & Sensor Networks*, pages 131–135. ACM, 2004.
11. Michael J Artelli and Richard F Deckro. Modeling the Lanchester laws with system dynamics. *The Journal of Defense Modeling and Simulation*, 5(1):1–20, 2008.
12. Tarunpreet Kaur and Dilip Kumar. Wireless multifunctional robot for military applications. In *Proc. of the 2nd IEEE International Conference on Recent Advances in Engineering & Computational Sciences (RAECS)*, pages 1–5, 2015.
13. Kui Qian, Aiguo Song, Jiatong Bao, and Huatao Zhang. Small teleoperated robot for nuclear radiation and chemical leak detection. *International Journal of Advanced Robotic Systems*, 9(3):70, 2012.
14. Abd Hafiz Zakaria, Yasir M Mustafah, Jaafar Abdullah, Nahrul Khair, and Taufiq Abdullah. Development of autonomous radiation mapping robot. *Procedia Computer Science*, 105:81–86, 2017.
15. Taylor Moore. Robots for nuclear power plants. *IAEA Bulletin*, 27(3):31–38, 1985.
16. Stephen L Smith, Mac Schwager, and Daniela Rus. Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*, 28(2):410–426, 2012.
17. Kevin M Lynch, Ira B Schwartz, Peng Yang, and Randy A Freeman. Decentralized environmental modeling by mobile sensor networks. *IEEE Transactions on Robotics*, 24(3):710–724, 2008.
18. Peng Yang, Randy A Freeman, and Kevin M Lynch. Multi-agent coordination by decentralized estimation and

- control. *IEEE Transactions on Automatic Control*, 53(11):2480–2496, 2008.
19. Rishi Graham and Jorge Cortés. Adaptive information collection by robotic sensor networks for spatial estimation. *IEEE Transactions on Automatic Control*, 57(6):1404–1419, 2012.
 20. Sonia Martínez. Distributed interpolation schemes for field estimation by mobile sensor networks. *IEEE Transactions on Control Systems Technology*, 18(2):491–500, 2010.
 21. Petter Ogren, Edward Fiorelli, and Naomi Ehrich Leonard. Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment. *IEEE Transactions on Automatic Control*, 49(8):1292–1302, 2004.
 22. Vijay Gupta, Timothy H Chung, Babak Hassibi, and Richard M Murray. On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage. *Automatica*, 42(2):251–260, 2006.
 23. Yilin Mo, Roberto Ambrosino, and Bruno Sinopoli. Sensor selection strategies for state estimation in energy constrained wireless sensor networks. *Automatica*, 47(7):1330–1338, 2011.
 24. Michael P Vitus, Wei Zhang, Alessandro Abate, Jianghai Hu, and Claire J Tomlin. On efficient sensor scheduling for linear dynamical systems. *Automatica*, 48(10):2482–2493, 2012.
 25. Syed Talha Jawaid and Stephen L Smith. Submodularity and greedy algorithms in sensor scheduling for linear dynamical systems. *Automatica*, 61:282–288, 2015.
 26. Ahmad Bilal Asghar, Syed Talha Jawaid, and Stephen L Smith. A complete greedy algorithm for infinite-horizon sensor scheduling. *Automatica*, 81:335–341, 2017.
 27. Nikolay Atanasov, Jerome Le Ny, Kostas Daniilidis, and George J Pappas. Information acquisition with sensing robots: Algorithms and error bounds. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 6447–6454, 2014.
 28. Nikolay Atanasov, Jerome Le Ny, Kostas Daniilidis, and George J Pappas. Decentralized active information acquisition: Theory and application to multi-robot SLAM. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4775–4782, 2015.
 29. Brent Schlotfeldt, Vasileios Tzoumas, Dinesh Thakur, and George J Pappas. Resilient active information gathering with mobile robots. *arXiv preprint arXiv:1803.09730*, 2018.
 30. Alberto Speranzon, Carlo Fischione, and Karl Henrik Johansson. Distributed and collaborative estimation over wireless sensor networks. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 1025–1030, 2006.
 31. Reza Olfati-Saber. Kalman-consensus filter: Optimality, stability, and performance. In *Proceedings of the 48th IEEE Conference on Decision and Control held jointly with the 28th Chinese Control Conference*, pages 7036–7042, 2009.
 32. Usman Khan and José MF Moura. Distributing the Kalman filter for large-scale systems. *IEEE Transactions on Signal Processing*, 56(10):4919–4935, 2008.
 33. Ion Matei and John S Baras. Consensus-based linear distributed filtering. *Automatica*, 48(8):1776–1782, 2012.
 34. Usman Khan, Soumya Kar, Ali Jadbabaie, and José MF Moura. On connectivity, observability, and stability in distributed estimation. In *Proceedings of the 49th IEEE Conference on Decision and Control*, pages 6639–6644, 2010.
 35. Usman A Khan and Ali Jadbabaie. Collaborative scalar-gain estimators for potentially unstable social dynamics with limited communication. *Automatica*, 50(7):1909–1914, 2014.
 36. Valery Ugrinovskii. Distributed robust estimation over randomly switching networks using H_∞ consensus. *Automatica*, 49(1):160–168, 2013.
 37. Mohammadreza Doostmohammadian and Usman A Khan. On the genericity properties in distributed estimation: Topology design and sensor placement. *IEEE Journal of Selected Topics in Signal Processing*, 7(2):195–204, 2013.
 38. Shinkyu Park and Nuno C Martins. Design of distributed LTI observers for state omniscience. *IEEE Transactions on Automatic Control*, 62(2):561–576, 2017.
 39. Aritra Mitra and Shreyas Sundaram. An approach for distributed state estimation of LTI systems. In *Proceedings of the 54th Annual Allerton Conference on Communication, Control, and Computing*, pages 1088–1093, 2016.
 40. A. Mitra and S. Sundaram. Distributed observers for LTI systems. *IEEE Transactions on Automatic Control*, 63(9):3689–3704, 2018.
 41. Lili Wang and A Stephen Morse. A distributed observer for a time-invariant linear system. *IEEE Transactions on Automatic Control*, 63(7):2123–2130, 2018.
 42. Weixin Han, Harry L Trentelman, Zhenhua Wang, and Yi Shen. A simple approach to distributed observer design for linear systems. *IEEE Transactions on Automatic Control*, 2018.
 43. Francisco FC Rego, A Pedro Aguiar, António M Pascoal, and Colin N Jones. A design method for distributed Luenberger observers. In *Proceedings of the 56th IEEE Conference on Decision and Control*, pages 3374 – 3379, 2017.
 44. AR del Nozal, L Orihuela, and P Millán. Distributed consensus-based Kalman filtering considering subspace decomposition. *IFAC-PapersOnLine*, 50(1):2494–2499, 2017.
 45. L Wang, AS Morse, D Fullmer, and J Liu. A hybrid observer for a distributed linear system with a changing neighbor graph. In *Proceedings of the 2017 56th IEEE Conference on Decision and Control*, pages 1024–1029, 2017.
 46. Ion Matei, John S Baras, and Vijay Srinivasan. Trust-based multi-agent filtering for increased smart grid security. In *Proceedings of the Mediterranean Conference on Control & Automation*, pages 716–721, 2012.
 47. Usman Khan and Aleksandar M Stankovic. Secure distributed estimation in cyber-physical systems. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 5209–5213, 2013.
 48. Mohammad Deghat, Valery Ugrinovskii, Iman Shames, and Cédric Langbort. Detection of biasing attacks on distributed estimation networks. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2134–2139, 2016.
 49. Aritra Mitra and Shreyas Sundaram. Secure distributed observers for a class of linear time invariant systems in the presence of Byzantine adversaries. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2709–2714, 2016.
 50. Aritra Mitra and Shreyas Sundaram. Byzantine-resilient distributed observers for LTI systems. *arXiv preprint arXiv:1802.09651*, 2018.

51. Nitin H Vaidya, Lewis Tseng, and Guanfeng Liang. Iterative approximate Byzantine consensus in arbitrary directed graphs. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, pages 365–374, 2012.
52. Heath J LeBlanc, Haotian Zhang, Xenofon Koutsoukos, and Shreyas Sundaram. Resilient asymptotic consensus in robust networks. *IEEE Journal on Selected Areas in Communications*, 31(4):766–781, 2013.
53. Shreyas Sundaram and Bahman Gharesifard. Consensus-based distributed optimization with malicious nodes. In *Proceedings of the Annual Allerton Conference on Communication, Control and Computing*, pages 244–249, 2015.
54. Lili Su and Nitin H Vaidya. Fault-tolerant multi-agent optimization: optimal iterative distributed algorithms. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 425–434. ACM, 2016.
55. Lili Su and Nitin H Vaidya. Non-Bayesian learning in the presence of Byzantine agents. In *International Symposium on Distributed Computing*, pages 414–427. Springer, 2016.
56. Yuan Chen, Soumya Kar, and Jose MF Moura. Resilient distributed estimation through adversary detection. *IEEE Transactions on Signal Processing*, 66(9):2455–2469, 2018.
57. Lewis Tseng, Nitin Vaidya, and Vartika Bhandari. Broadcast using certified propagation algorithm in presence of Byzantine faults. *Information Processing Letters*, 115(4):512–514, 2015.
58. Kelsey Saulnier, David Saldana, Amanda Prorok, George J Pappas, and Vijay Kumar. Resilient flocking for mobile robot teams. *IEEE Robotics and Automation Letters*, 2(2):1039–1046, 2017.
59. Luis Guerrero-Bonilla, Amanda Prorok, and Vijay Kumar. Formations for resilient robot teams. *IEEE Robotics and Automation Letters*, 2(2):841–848, 2017.
60. A Yasin Yazıcıoğlu, Magnus Egerstedt, and Jeff S Shamma. Formation of robust multi-agent networks through self-organizing random regular graphs. *IEEE Transactions on Network Science and Engineering*, 2(4):139–151, 2015.
61. David Saldana, Amanda Prorok, Shreyas Sundaram, Mario FM Campos, and Vijay Kumar. Resilient consensus for time-varying networks of dynamic agents. In *Proceedings of the American Control Conference*, pages 252–258, 2017.
62. James Usevitch and Dimitra Panagou. r -robustness and (r, s) -robustness of circulant graphs. In *Proceedings of the 56th IEEE Conference on Decision and Control*, pages 4416–4421, 2017.
63. Hyongju Park and Seth A Hutchinson. Fault-tolerant rendezvous of multirobot systems. *IEEE Transactions on Robotics*, 33(3):565–582, 2017.
64. Hyongju Park and Seth Hutchinson. Robust rendezvous for multi-robot system with random node failures: an optimization approach. *Autonomous Robots*, pages 1–12, 2018.
65. Seyed Mehran Dibaji and Hideaki Ishii. Resilient consensus of second-order agent networks: Asynchronous update rules with delays. *Automatica*, 81:123–132, 2017.
66. Danny Dolev, Nancy A Lynch, Shlomit S Pinter, Eugene W Stark, and William E Weihl. Reaching approximate agreement in the presence of faults. *Journal of the ACM (JACM)*, 33(3):499–516, 1986.
67. Aritra Mitra and Shreyas Sundaram. Secure distributed state estimation of an LTI system over time-varying networks and analog erasure channels. In *Proceedings of the American Control Conference*, pages 6578–6583, 2018.
68. Sandip Roy and Rahul Dhal. Situational awareness for dynamical network processes using incidental measurements. *IEEE Journal of Selected Topics in Signal Processing*, 9(2):304–316, 2015.
69. Dorina Thanou, Xiaowen Dong, Daniel Kressner, and Pascal Frossard. Learning heat diffusion graphs. *IEEE Transactions on Signal and Information Processing over Networks*, 3(3):484–499, 2017.
70. Fan Chung. The heat kernel as the pagerank of a graph. *Proceedings of the National Academy of Sciences*, 104(50):19735–19740, 2007.
71. Chi-Tsong Chen. *Linear system theory and design*. Oxford University Press, Inc., 1998.
72. Soroush Alamdari, Elaheh Fata, and Stephen L Smith. Persistent monitoring in discrete environments: Minimizing the maximum weighted latency between observations. *The International Journal of Robotics Research*, 33(1):138–154, 2014.
73. Shreyas Sundaram and Christoforos 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.
74. Fabio Pasqualetti, Antonio Bicchi, and Francesco Bullo. Consensus computation in unreliable networks: A system theoretic approach. *IEEE Transactions on Automatic Control*, 57(1):90–104, 2012.
75. Hamza Fawzi, Paulo Tabuada, and Suhas Diggavi. Secure estimation and control for cyber-physical systems under adversarial attacks. *IEEE Transactions on Automatic Control*, 59(6):1454–1467, 2014.
76. Michelle S Chong, Masashi Wakaiki, and Joao P Hespanha. Observability of linear systems under adversarial attacks. In *Proceedings of the American Control Conference*, pages 2439–2444, 2015.
77. Pablo Millán, Luis Orihuela, Carlos Vivas, and Francisco R Rubio. Distributed consensus-based estimation considering network induced delays and dropouts. *Automatica*, 48(10):2726–2729, 2012.
78. Nicola Elia. Remote stabilization over fading channels. *Systems & Control Letters*, 54(3):237–249, 2005.
79. Joao P Hespanha, Payam Naghshtabrizi, and Yonggang Xu. A survey of recent results in networked control systems. *Proceedings of the IEEE*, 95(1):138–162, 2007.
80. Xiao Liu, Sitian Cheng, Hong Liu, Sha Hu, Daqiang Zhang, and Huansheng Ning. A survey on gas sensing technology. *Sensors*, 12(7):9635–9665, 2012.
81. Haotian Zhang, Elaheh Fata, and Shreyas Sundaram. A notion of robustness in complex networks. *IEEE Transactions on Control of Network Systems*, 2(3):310–320, 2015.
82. James Usevitch and Dimitra Panagou. Resilient leader-follower consensus to arbitrary reference values. In *Proceedings of the American Control Conference*, pages 1292–1298, 2018.
83. Aritra Mitra and Shreyas Sundaram. A novel switched linear observer for estimating the state of a dynamical process with a mobile agent. In *Proceedings of the 57th IEEE Conference on Decision and Control*, 2018.
84. Ankush Chakrabarty, Raid Ayoub, Stanislaw H Żak, and Shreyas Sundaram. Delayed unknown input observers for discrete-time linear systems with guaranteed performance. *Systems & Control Letters*, 103:9–15, 2017.

-
85. Ankush Chakrabarty, Emilia Fridman, Stanisław H Żak, and Gregory T Buzzard. State and unknown input observers for nonlinear systems with delayed measurements. *Automatica*, 95:246–253, 2018.