CERIAS

The Center for Education and Research in Information Assurance and Security

Resilient Distributed State Estimation and Hypothesis Testing in Dynamic Environments

PURDUE UNIVERSITY

Aritra Mitra[1], John A. Richards[2], Saurabh Bagchi[1] and Shreyas Sundaram[1]

[1]School of Electrical and Computer Engineering, Purdue University, [2]Sandia National Laboratories

Sandia National Laboratories

1. Problem Description

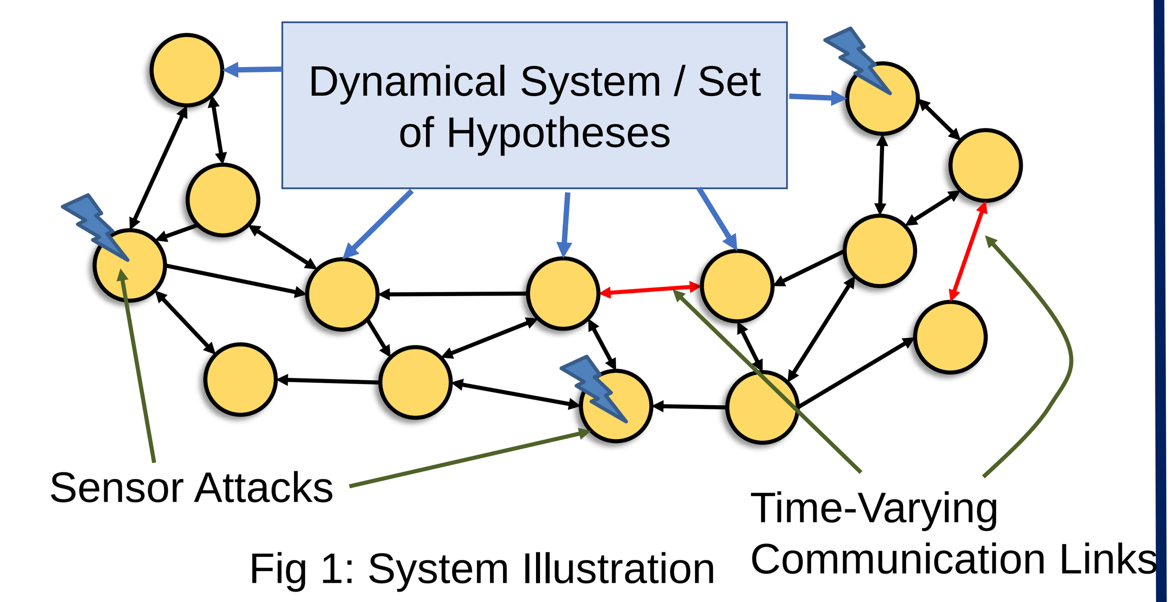

Fig 1: System Illustration

Problem 1: Distributed State Estimation

- A dynamical process $x(k+1) = Ax(k)$ evolves over a region; $x(k)$ is the state of the process at time k.
- A network of sensors monitor the state of the process. Sensor i has measurement $y_i(k) = C_i x(k)$.
- Sensor i maintains an estimate of the state $\hat{x}_i[k]$.
- Goal: The estimate of each sensor should converge to the true state asymptotically.
- Applications: Environmental monitoring of diffusive processes, power systems, smart factories etc.

Problem 2: Distributed Hypothesis Testing/ Non-Bayesian Social Learning

- A group of agents (sensors or humans) aim to learn the true state of the world $\theta^\star \in \{\theta_1, \dots \theta_m\}$.
- Agent i receives i.i.d. private observations $\{s_{i,k}\}$.
- Agent i maintains a belief vector $\mu_{i,k}$ over the set of possible hypotheses, denoted $\Theta = \{\theta_1, \dots \theta_m\}$.
- Goal: The belief of each agent should asymptotically concentrate on the true state $\theta^\star$.
- Engineering Applications: Detection problems (e.g., detecting radiation leakage), object classification, target recognition etc.
- Applications in Social Networks: Deciding which product to buy, which candidate to vote for, whether a news item is true or fake etc.

Contribution: For each of the two problems, we develop novel distributed algorithms that are *light-weight, scalable, robust to various types of communication losses, resilient against worst-case adversarial attacks on certain sensors/agents, and that guarantee exponentially fast convergence to the true state.*

2. Resilient Distributed State Estimation

Basic Estimation Strategy

- Each node can potentially estimate certain portions of the state based on local measurements.
- For remaining states, it communicates with neighbors.
- Identify "source nodes" or "leaders" that can track the unstable (or critical) states.
- Disseminate information from source nodes to the rest of the network.

Worst-Case Byzantine Adversary Model

- Each adversary has complete knowledge of the system model, and can act arbitrarily.
- There are at most f adversaries in the neighborhood of any good node.

Mode Estimation Directed Acyclic Graph

- MEDAGs provide sufficient number of redundant paths for transmitting info from source nodes to rest of the graph.
- Each non-source node has at least $(2f+1)$ parents in a MEDAG.

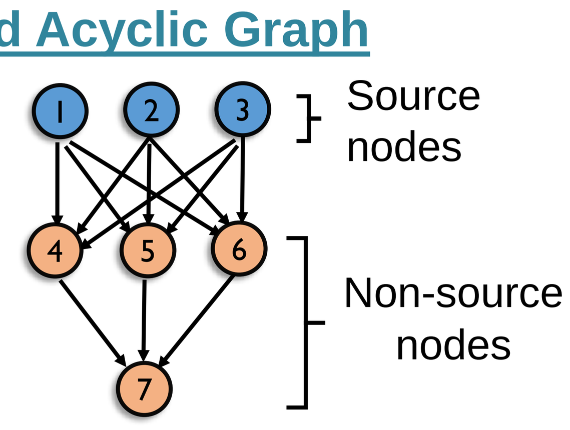

Fig 2: An Illustration of a MEDAG

Local-Filtering based Resilient Estimation

- Let $\mathcal{S}$ denote the nodes that can estimate $x(k)$.
- A good node $i \notin \mathcal{S}$, updates $x(k)$ as follows:
- Step 1: At each time-step k, node i collects the estimates of its parents in the MEDAG.
- Step 2: It rejects the f highest and f lowest estimates (*i.e., rejects extreme estimates*), and updates $\hat{x}_i[k]$ as

$$\hat{x}_i[k+1] = a\left(\sum_{j \in \mathcal{M}_i[k]} w_{ij}[k]\hat{x}_j[k]\right)$$

state transition model

convex weights

parents after removing extreme neighbors

Main Result:

Theorem: $(3f+1)$ strong-robustness of $\mathcal{G}$ w.r.t. the source set $\mathcal{S}$ ⟹ Each good node can track the state exponentially fast, despite the actions of any f-local adversarial set.

Simulation Example

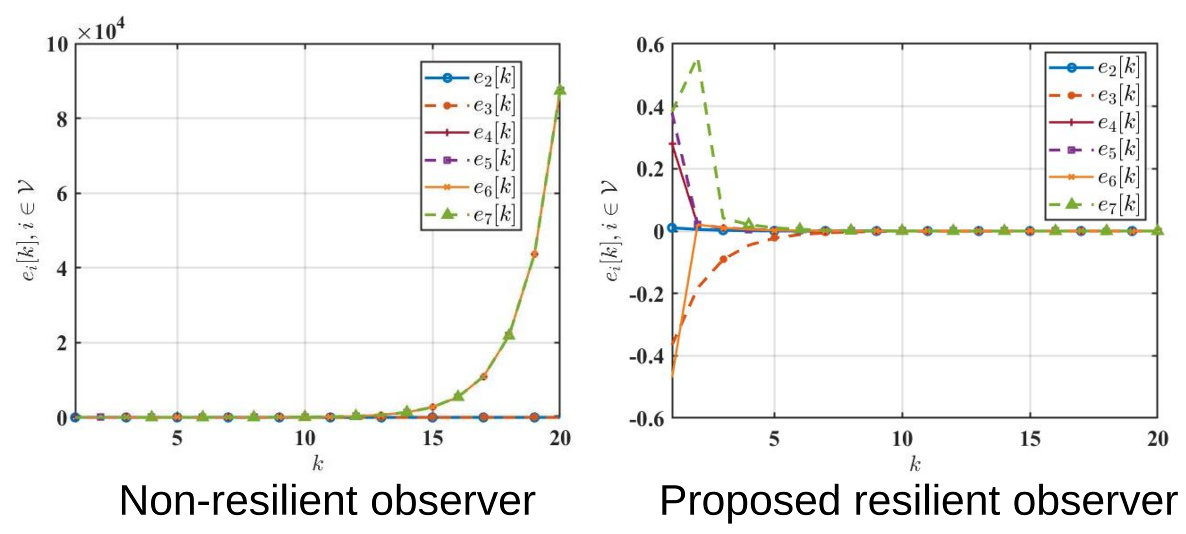

Fig 3. Error plots for network in Figure 2. Node 1 is adversarial.

References:

- *Byzantine-Resilient Distributed Observers for LTI Systems,* Mitra and Sundaram, CDC 16, arXiv 19.
- *Resilient Distributed State Estimation with Mobile Agents*, Mitra et al., ACC 16, Autonomous Robots 19.

3. Resilient Distributed Hypothesis Testing

Basic Detection Strategy

- Each agent can potentially eliminate certain false hypotheses based on private observations.
- To eliminate other false hypotheses, it communicates with neighbors.
- Let $\mathcal{S}(\theta^\star, \theta)$ denote the set of agents that can distinguish between the true state $\theta^\star$ and $\theta \neq \theta^\star$.
- $\mathcal{S}(\theta^\star, \theta)$ can be viewed as the set of "source agents" for the pair $(\theta^\star, \theta)$.
- Disseminate information from $\mathcal{S}(\theta^\star, \theta)$ to other agents to help them eliminate the false hypothesis θ.
- We consider an f-local Byzantine adversary model.

Local-Filtering Based Resilient Hypothesis Elimination

- Each agent maintains a local belief vector $\pi_{i,k}$, and an actual belief vector $\mu_{i,k}$.
- Step 1: The local belief-vector of agent i is updated in a standard Bayesian way:

$$\pi_{i,k+1}(\theta) = \frac{l_i(s_{i,k+1}|\theta)\pi_{i,k}(\theta)}{\sum_{p=1}^{m} l_i(s_{i,k+1}|\theta_p)\pi_{i,k}(\theta_p)}$$

- Step 2: If $|\mathcal{N}_i| < (2f+1)$, then agent i updates its actual beliefs as:

$$\mu_{i,k+1}(\theta) = \pi_{i,k+1}(\theta)$$

- Step 3: If $|\mathcal{N}_i| \geq (2f+1)$, then agent i collects the actual beliefs $\mu_{j,k}(\theta)$ of its neighbors, rejects the highest f and lowest f of them, and updates $\mu_{i,k}(\theta)$ as

$$\mu_{i,k+1}(\theta) = \frac{\min\{\{\mu_{j,k}(\theta)\}_{j \in \mathcal{M}^{\theta}_{i,k}}, \pi_{i,k+1}(\theta)\}}{\sum_{p=1}^{m} \min\{\{\mu_{j,k}(\theta_p)\}_{j \in \mathcal{M}^{\theta_p}_{i,k}}, \pi_{i,k+1}(\theta_p)\}}$$

neighbors of agent i with moderate beliefs on θ_p at time k.

Main Result

Theorem: (i) Strong $(2f+1)$-robustness of $\mathcal{G}$ w.r.t. every source set $\mathcal{S}(\theta_p, \theta_q), \theta_p, \theta_q \in \Theta$, and (ii) non-zero priors of good agents on each hypothesis ⟹ Each good agent can rule out every false hypothesis exponentially fast, despite the actions of any f-local adversarial set.

Simulation Example

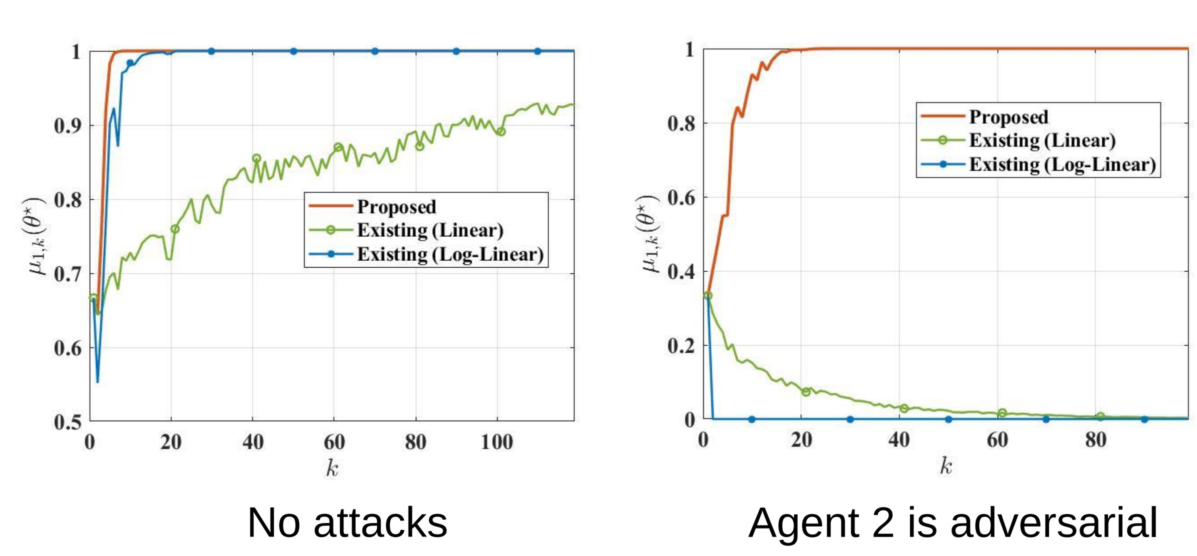

Fig 4. Plots of beliefs on true state for undirected version of network in Figure 2.

References:

- *A New Approach for Distributed Hypothesis Testing with Extensions to Byzantine-Resilience*, Mitra et al., ACC 19.

Email: {mitra14,sbagchi,sundara2}@purdue.edu, jaricha@sandia.gov