Distributed Resilience for Swarms of Autonomous Agents

Shaoshuai Mou, Shreyas Sundaram

PURDUE UNIVERSITY

CRISP, Nov. 8, 2019

Why Multi-Agent Systems (MAS)?

MAS usually offers better **autonomy**, **robustness**, **flexibility… and** are able to achieve sophisticated missions that are well beyond individual systems' capabilities

Features of MAS

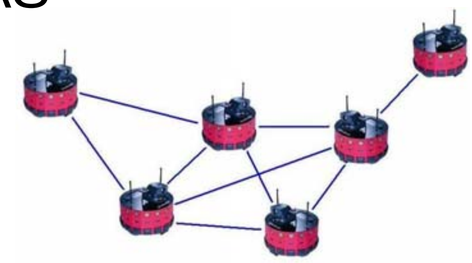

- Each individual: mobile | autonomous | low-cost | local accessibility

- The whole swarm: large-scale | heterogeneous | communication constraints | no centralized controller

- Key enabler: **local** coordination (sensing/communications among nearby neighbors) → **Distributed** algorithms for control/optimization/coordination → **global** objectives (exploration of the unknown, formation flight, cooperative sensing, search/rescue)

- Advantages: scalability, flexibility, **robust** under individual node/link failures.

- Disadvantage: The whole network is **vulnerable to sophisticated cyber-attacks**. (strong dependence on local coordination)

multi-robot rendezvous problem

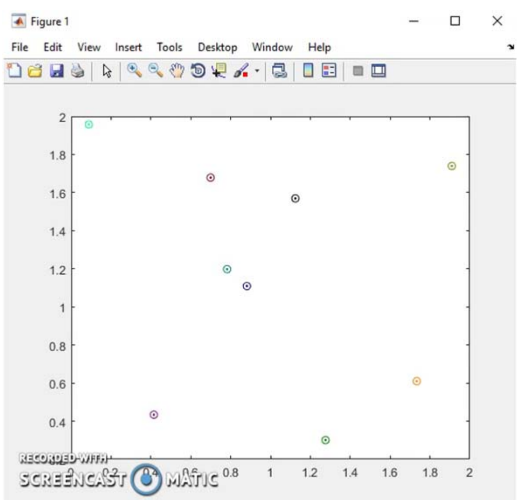

Ideal Case

One robot is under attack

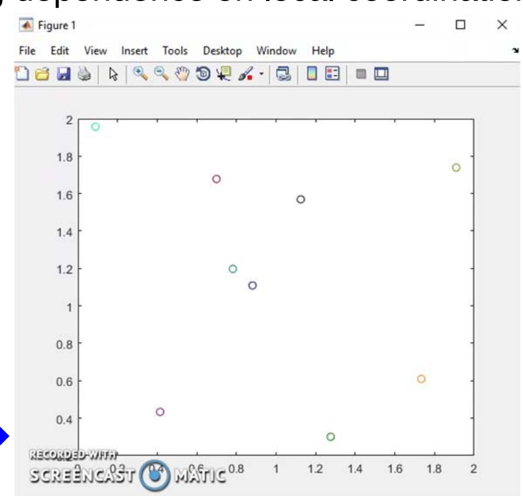

Research Challenges towards Resilience for MAS

Agents

- **low-cost (limited sensing/processing)**
- only **local information** available

Cyber-Attacks

- **highly mobile**
- **sophisticated** (fully control the agent)
- launched **massively** from multiple locations

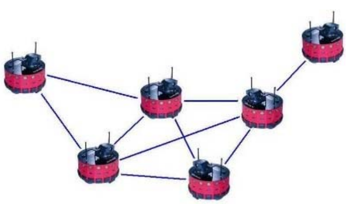

Underlying Networks

- **time-varying**

"the ability to prepare and plan for, absorb, recover from, and more successfully adapt to adverse events."

Techniques for resilience in MAS must be **automated**, **adaptive** to time-varying networks and mobile malicious agents, works in a **fully distributed** scenario.

Resilience for Consensus-based Distributed Algorithms

All agents reach an agreement on some quantity of interest

- Consensus is the key enabler for swarm of agents to work as a **cohesive whole**
  - **Unconstrained Consensus:** All agents reach the **same** value.
  - **Constrained Consensus:** All agents reach the **same** and **specific** value.
  - **Optimal Constrained Consensus:** All agents reach the **same** and **specific** value, which minimizes a **global cost function**.

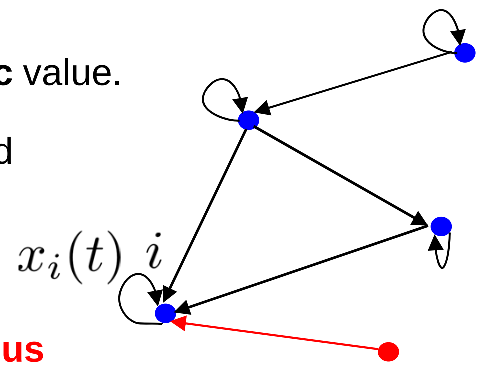

- Common Structure for Consensus:

$$x_i(t+1) = f_i(v_i)$$

failure of consensus

$$v_i = \sum_{j \in \mathcal{N}_i(t)} w_{ij} x_j(t)$$ injects information from malicious agents**.**

How to exclude malicious information without knowing which one is malicious?

- Example:

An normal agent *i* receives some states in the plane.

Suppose it is known that at most one of them is malicious.

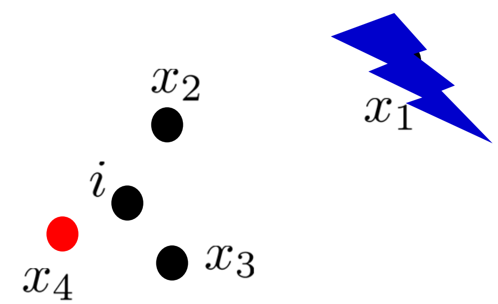

majority rule? works for unconstrained consensus

is **not applicable** here because of local constraints

Any Ideas?

- Achieve an ***resilient average*** $\bar{v}_i = \sum_{j \in \bar{\mathcal{N}}_i} \bar{w}_{ij} x_j$, which
  - is **always** a weighted average of normal agents' states
  - could be achieved **only by local information** plus little knowledge about the malicious agents (an upper bound of how many malicious agents in agents' neighbors)

- **Key Idea**: **Intersection of convex hulls.**

$$\mathcal{M}_i = \{x_i, x_1, x_2, x_3, x_4\} \qquad d_i = 5$$

Suppose we know that at most $r_i = 1$ are malicious.

(1.) List all subsets of $\mathcal{M}_i$ with size $d_i - r_i = 4$

(2.) Choose $\bar{v}_i \bigcap_{k \in r} \mathcal{H}(\mathcal{M}_{ik})$ is always a resilient average $\bar{v}_i = \sum_{j \in \bar{\mathcal{N}}_i} \bar{w}_{ij} x_j$

Advantage: Utilizes all received information but eliminate impacts of malicious ones.

- **Simulations on** consensus-based distributed optimizations.

Time-varying networks; **Byzantine** attack; **Time-varying** locations; Only **local** information available.

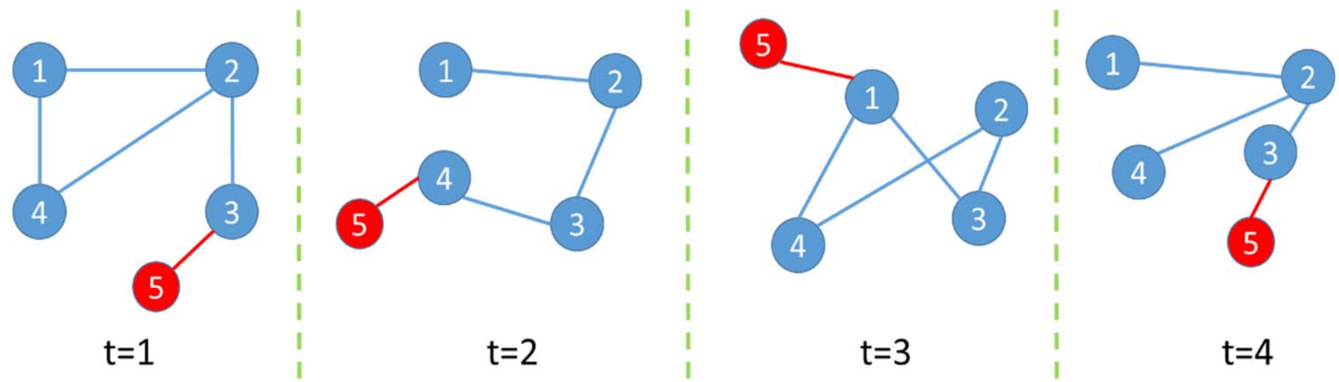

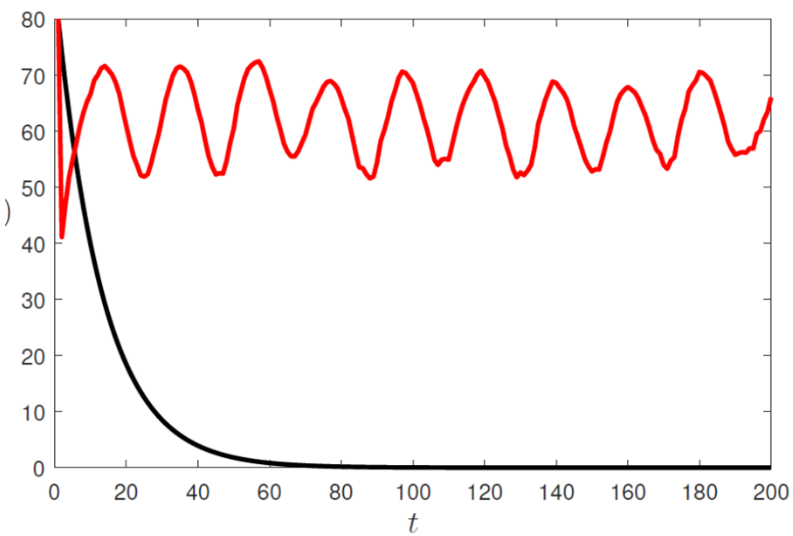

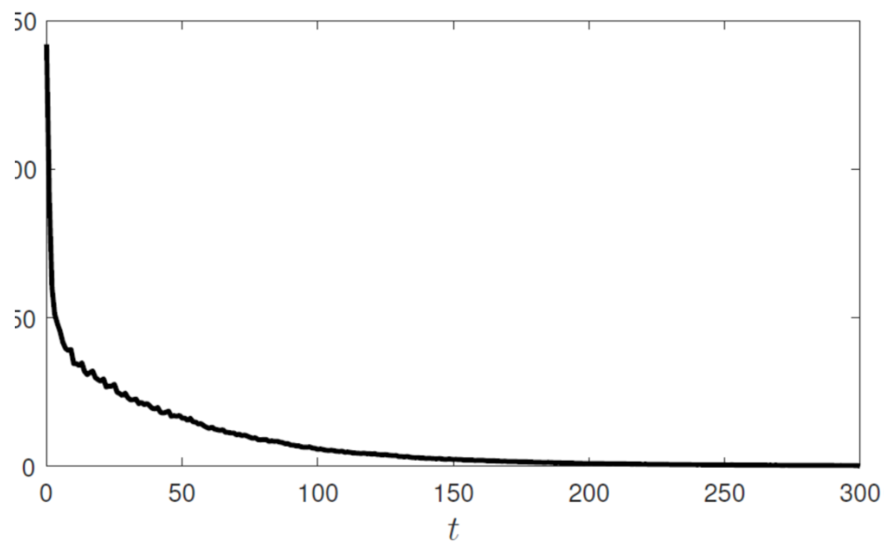

- *X. Wang, S. Mou, S. Sundaram. Numerical Algebra, Control and Optimizations, 9(3), 269-281, 2019*
- *Presentation at the 8th Midwest Workshop on Control and Game Theory, April, 2019*
- *Presentation at the 57th Annual Allerton Conference on Communication, Control and Computing, Oct. 2019*

Ongoing Project: AI-assisted Multi-Agent Systems with Human-in-the-Loop

- We aim to develop an AI-assisted multi-agent platform, which is able to provide **distributed environment perception/situation awareness**; perform **real-time, dynamic and distributed control/management** of assets, planning/decision making/task assignment; include **human-in-the-loop**; **improve performance** as time evolves.

This ongoing project relies on collaboration of multiple universities, with funding from **Northrop Grumman (2019-2021).*

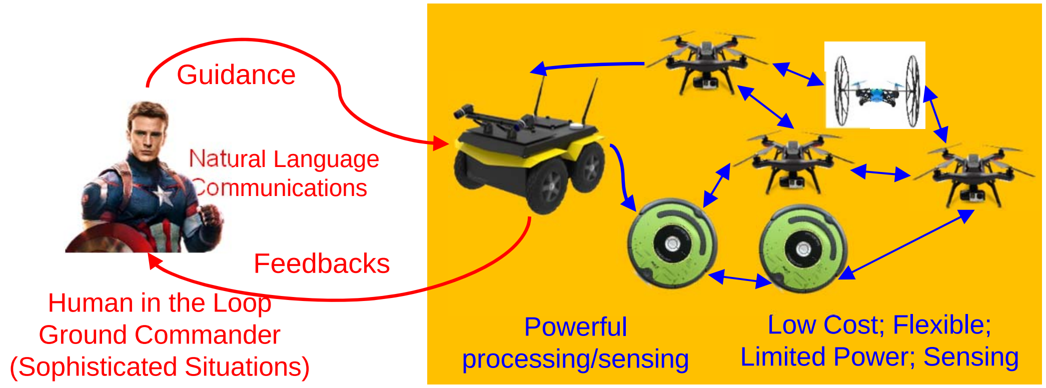

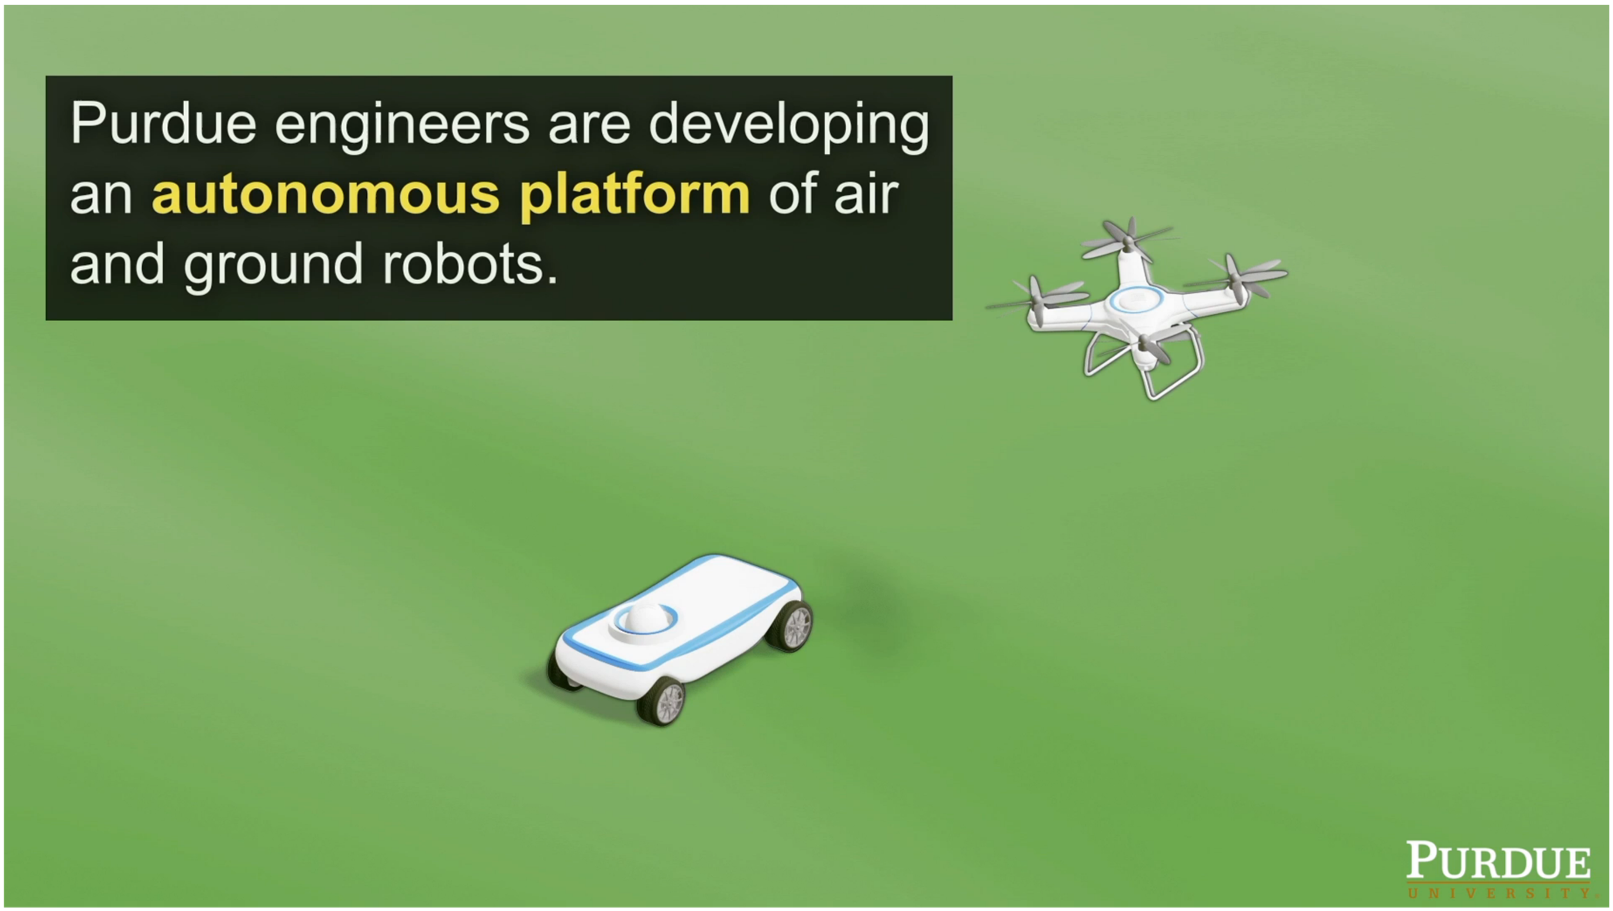

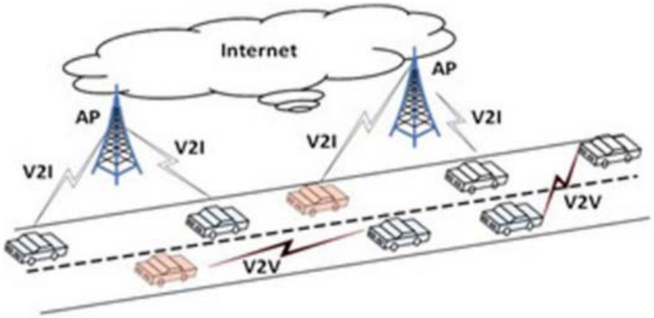

Large vehicular networks

top-down
analysis

Social Networks

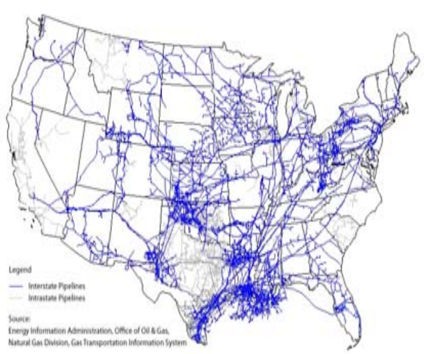

Transportation Networks

Resilient,
Distributed algorithms
for control/optimization/coordination

Thank You!

bottom-up
synthetics

?

For our experimental results, please refer to
Autonomous & Intelligent Multi-agent Systems
(AIMS) Lab
https://engineering.purdue.edu/AIMS