Resilient Distributed Optimization for Cyberphysical Systems

Angelia.Nedich@asu.edu

Arizona State University

Joint work with

Michal Yemini (Bar Ilan University)

Stephanie Gil, Orhan Akgun, Kerem Dayi (Harvard University)

Andrea Goldsmith, Stony Brook University

Cyberphysical Systems

- Interconnected complex system of various AI devices (sensing, computing, information), and human decisions impacting the system.
- Critical technologies:
  - Provision of various virtual services through CPSs
  - Models for humans and their interactions with CPSs
  - Safety
  - **Reliability, Trustworthiness**
- We focus on
  - Agreement (consensus) as a simple model to understand the impact of untrustworthy (noncooperative or malicious) agents
  - Distributed consensus-based methods in presence of untrustworthy agents
  - Resilient coordination in networked multi-robot teams (Stephanie Gil's Lab)

Classic Agreement Problem

There are m agents connected over a communication network.

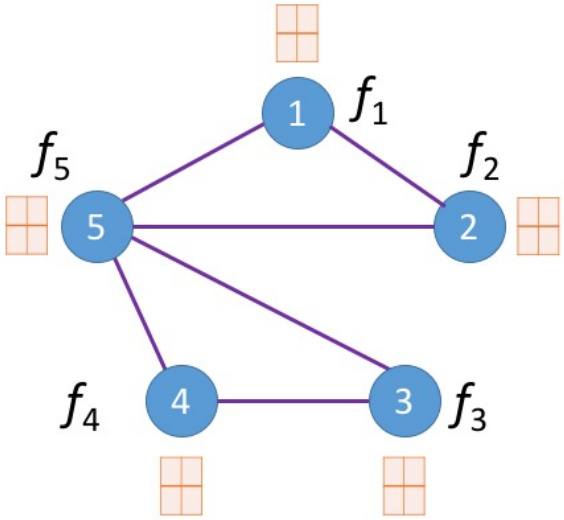

- Agents want to reach an agreement on a common value
- Agents communicate and exchange information with their immediate neighbors only
- There is no central entity that coordinates agents' computations
- There is no common (shared) memory space
- Each agent has local knowledge of the connectivity graph (its neighbors only)
- No agent knows even the total number m of the agents in the system

DeGroot Consensus/Agreement Model

- Consider a set of m agents where every agent i has a value $x_i(0) \in \mathbb{R}$ (opinion).
- The graph representing who-knows-whom is a strongly connected directed graph $\mathcal{G} = ([m], \mathcal{E})$
- Agent i chooses positive trust weights $w_{ij} > 0$ for its in-neighbors $j \in N_i^{\text{in}}$; these weights sum to 1, $\sum_{j \in N_i^{\text{in}}} w_{ij} = 1$
- Over time, the agents communicate with their neighbors and share their values Specifically, at each time t, each agent $i \in [m]$
  - Receives values $x_j(t)$ from its in-neighbors N_i^{in} and
  - Sends its own value $x_i(t)$ to its out-neighbors N_i^{out}

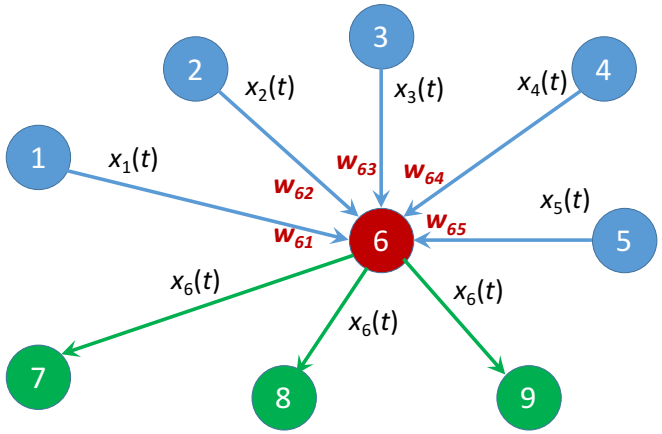

- The *agents obtain no other information based on which they can update their values*
- Upon sharing their values, the agents update using the trust weights they have selected,

$$x_i(t+1) = \sum_{j \in N_i^{\text{in}}} w_{ij} x_j(t) \qquad \text{for all } i \in [m]$$

- To compactly write the evolution of opinions, define

$$w_{ij} = 0 \qquad \text{for all } j \notin N_i^{\text{in}} \text{ and for all } i \in [m]$$

and let $W = [w_{ij}]$. Define $x(t)$ as the column vector with entries $x_i(t)$, $i \in [m]$. Then, we have

$$x(t+1) = Wx(t) \qquad \text{for all } t \geq 0$$

- Thus, the evolution of $x(t)$ is linear

$$x(t) = W^t x(0) \qquad \text{for all } t \geq 0$$

- The trust matrix W is stochastic, i.e., it is a non-negative matrix and the entries sum to 1 in each row

$$W \geq 0, \qquad W\mathbf{1} = \mathbf{1}$$

where $\mathbf{1}$ is the m-dimensional vector with all entries equal to 1.

Existence and Characterization of the Limit

- The matrix W **is compliant with graph** $\mathcal{G}$: there is an edge from j to i iff $W_{ij} > 0$.
- **We assume that the graph** $\mathcal{G}$ **is strongly connected**
- Analysis using Markov Chain theory
  - View W as a transition matrix of a homogeneous Markov Chain: chain is ergodic

$$\lim_{t\to\infty} W^t = \mathbf{1}\pi',$$

    where $\pi = [\pi_1, \ldots, \pi_m]'$ is a positive stochastic vector, i.e., $\pi > 0$ and $\mathbf{1}'\pi = 1$.
  - The vector π is the vector of steady-state distributions of the chain leading to

$$\lim_{t\to\infty} x_i(t) = \langle \pi, x(0)\rangle \qquad \text{for all } i \in [m]$$

- The agents reach a consensus and the consensus value is $\pi' x(0)$
- Convergence rate to consensus is geometric

$$\sum_{i=1}^{m} \pi_i \left(x_i(t+1) - \langle \pi, x(0)\rangle\right)^2 \leq \rho_W \sum_{i=1}^{m} \pi_i \left(x_i(t) - \langle \pi, x(0)\rangle\right)^2$$

  where $\rho_W \in (0,1)$ is the second largest (in modulus) eigenvalue of W
- **Critical assumption**: **Every agent is honest in reporting its value** $x_i(t)$ **and executes the prescribed update rule**. All agents are trustworthy!

Vulnerability to Untrustworthy/Unreliable Agents

- **Classic Result:** In the presence of faulty agents, the agreement among the agents can be reached if and only if *the number of faulty agents is less than 1/3 of the total number of agents and less than 1/2 of the connectivity k of the graph*[*]; Dolev 1982
- **Detection of Malicious Agents**
  - Consensus can be reached under $k/2$ connectivity assumption in undirected graphs ([Sundaram, Hadjicostis 2008], [Pasqualetti, Bicchi, Bullo 2012], [LeBlank et al. 2013]). The consensus value reached among legitimate agents is the same value that would have been reached in the absence of untrustworthy agents - requires identification of untrustworthy agents.
- **Robust Consensus**
  - Regardless whether untrustworthy agents exist or not, the basic consensus method is adjusted so that every agent discards k of its neighbors' values (by cutting out those whose values "maximally" exceed its own value) [LeBlanc, Zhang, Koutsoukos, Sundaram 2013]

[*]Every agent is connected to every other agent via k disjoint paths

  - The consensus is reached on a value that differs from the value that would be reached by trustworthy agents in the absence of untrustworthy agents
  - The achieved consensus value lies in the convex hull of the values $x_i(0)$, $i \in [m]$

- The existing approaches are at the two extremes. Can we allow agents to adjust by learning whom to trust over time? - Yes, if agents have side information to learn from. This motivates our approach: **Consensus with Learning whom to Trust**[†]

[†]M. Yemini, A. Nedić, A. J. Goldsmith and S. Gil, "Characterizing trust and resilience in distributed consensus for cyberphysical systems," IEEE Transactions on Robotics, vol. 38, no. 1, pp. 71-91, 2022.

Learning Trustworthiness - Yemini et al. 2022

- We allow each legitimate agent ℓ to learn and adapt the trust values $W_{\ell j}(t)$ for its neighbors $j \in N_\ell$
- We refer to untrustworthy neighbors as "malicious"
- The learning process is based on agents receiving a random signal $\alpha_{\ell j} \in [0, 1]$, $j \in N_\ell$
- At time t, each agent ℓ receives $x_j(t)$ from neighbor j and observes a random realization $\alpha_{\ell j}(t)$ as a private signal that j is trustworthy (signal can come from communication channel[‡])
- $\alpha_{\ell j}(t) > 1/2$ indicates that agent j is more likely trustworthy, while $\alpha_{\ell j}(t) < 1/2$ indicates that agent j is less likely trustworthy
- We assume that these private signals $\alpha_{\ell j}(t)$ are informative of the true agent's nature on average:
  $\mathbb{E}[\alpha_{\ell j}(t)] > 1/2$ for all legitimate neighbors $j \in N_\ell$ and
  $\mathbb{E}[\alpha_{\ell j}(t)] < 1/2$ for all malicious neighbors $j \in N_\ell$

[‡]S. Gil, S. Kumar, M. Mazumder, D. Katabi, D. Rus, "Guaranteeing spoof-resilient multi-robot networks," Autonomous Robots, 2017.

- The trust observations $\{\alpha_{\ell j}(t), j \in N_\ell, \ell \in \mathcal{L}\}$ are independent across time t
- At time t, each legitimate agent ℓ computes $\beta_{\ell j}(t) = \sum_{k=0}^{t}(\alpha_{\ell j}(k) - 1/2)$ and decides to trust all neighbors j such that $\beta_{\ell j}(t) \geq 0$, i.e., it defines its neighbor set as:

$$N_\ell(t) = \{j \in N_\ell \mid \beta_{\ell j}(t) \geq 0\}.$$

- A trustworthy agent ℓ gives positive trust values $w_{\ell j}(t)$ only to neighbors $j \in N_\ell(t)$, for example, equal-weights:

$$w_{\ell\ell}(t) = \frac{1}{|N_\ell(t)| + 1}, \qquad w_{\ell j}(t) = \frac{1}{|N_\ell(t)| + 1}, \qquad j \in N_\ell(t)$$

and $w_{\ell i}(t) = 0$ for all other neighbors i, or

$$w_{\ell j}(t) = \frac{1}{\max\{\kappa, |N_\ell(t)| + 1\}}, \quad j \in N_\ell(t), \qquad w_{\ell\ell}(t) = 1 - \sum_{j \in N_\ell(t)} w_{\ell j}(t),$$

where $\kappa > 0$ can be thought of as a parameter limiting the maximum influence that the neighbors of agent ℓ can have

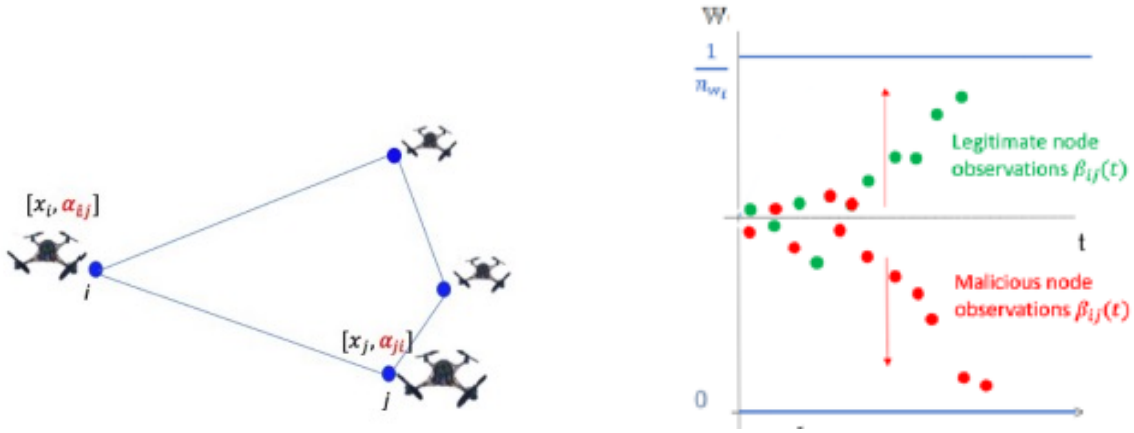

Figure 1: On the left, illustration of the transmitted data. To the right, illustration of the trend for the accumulated evidence $\beta_{\ell j}(t)$.

DeGroot Model with Untrustworthy Agents

- Consider now DeGroot opinion model where some agents are malicious (could be faulty), i.e., agent i may report false values $x_i(t)$ over time
- Untrustworthy agent may intent to manipulate the consensus value or to prevent consensus from happening
- For modeling purpose, we let $\mathcal{L}$ denote the set of trustworthy agents (legitimate) and we let $\mathcal{M}$ be the set of malicious agents
- *We do not assume that we know their number $\mathcal{M}$ nor their identities*
- Assume that the legitimate agents are suspicious of the presence of untrustworthy neighbors so they update cautiously using the trusted neighbors: for all $\ell \in \mathcal{L}$,

$$x_\ell(t+1) = \sum_{j \in N_\ell(t)} w_{\ell j}(t) x_j(t) + w_{\ell\ell}(t) x_\ell(t)$$

  where the trust values satisfy $w_{\ell j}(t) > 0$, $j \in N_\ell(t) \cup \{\ell\}$, and they sum to 1.
- The agent may assign trust value $w_{\ell j}(t) = 0$ to its neighbor $j \in N_\ell$
- The updates of untrustworthy agent values are not modeled (they are not unknown)

- Assuming that
  - All initial opinions are bounded by some commonly known value η, i.e., $|x_i(t)| \leq \eta$ for all i. Obviously, to avoid detection, no malicious agent will ever report a value that violates this condition
  - Each malicious agent sends the same information to all of its neighbors
  - The legitimate agents are connected in the underlying communication graph
  - The legitimate agents spend time T_0 learning i.e., collecting the trust observations but not updating prior to $T_0 \geq 0$ (parameter). The consensus updates start at T_0
  - For all legitimate agents $\ell \in \mathcal{L}$ and all $t \geq 0$,

$$\mathbb{E}[\alpha_{\ell j}(t)] - \frac{1}{2} = E_{\mathcal{L}} > 0 \qquad \text{for all } j \in N_i \cap \mathcal{L},$$

$$\mathbb{E}[\alpha_{\ell j}(t)] - \frac{1}{2} = E_{\mathcal{M}} < 0 \qquad \text{for all } j \in N_i \cap \mathcal{M}.$$

- Basic questions to answer
  - **Convergence:** Will the protocol lead to consensus among legitimate agents?
  - **Deviation:** How far is the achieved consensus value from the (nominal) consensus value that would have been achieved in the absence of untrustworthy agents?
  - **Convergence Rate:** What is the rate of convergence to the achieved consensus value?

- We have answered all three questions. The main properties we have shown are
  - Consensus is achieved almost surely
    - Almost surely, there exists some (random) finite time $T_f > 0$ such that $W_{\mathcal{L}}(t) = \bar{W}$ for all $t \geq T_f$, where $\bar{W}$ is a stochastic matrix
    - $\lim_{t\to\infty} \bar{W}^t = \mathbf{1}v'$ for a stochastic vector $v > 0$
    - $\lim_{t\to\infty} x_{\mathcal{L}}(t) = z(0)\mathbf{1}$ with $z(0)$ being in the convex hull of both legitimate and malicious agents initial values $x_i(0), i = 1, \ldots, m$
  - For the random finite time we have, for all $k \geq 1$,

$$\mathrm{Prob}[T_f = k] \leq \min\{1, q(k-1)\},$$

where

$$q(k) = D_{\mathcal{L}} \exp(-2kE_{\mathcal{L}}^2) + D_{\mathcal{M}} \exp(-2kE_{\mathcal{M}}^2),$$

with $D_{\mathcal{L}} = \sum_{i\in\mathcal{L}} |N_i \cap \mathcal{L}|$ and $D_{\mathcal{M}} = \sum_{i\in\mathcal{L}} |N_i \cap \mathcal{M}|$.

- We also have

$$\text{Prob}[T_f > k-1] \leq \min\{1, p(k-1)\},$$

where

$$p(k) = D_{\mathcal{L}} \frac{\exp(-2kE_{\mathcal{L}}^2)}{1-\exp(-2E_{\mathcal{L}}^2)} + D_{\mathcal{M}} \frac{\exp(-2kE_{\mathcal{M}}^2)}{1-\exp(-2E_{\mathcal{M}}^2)},$$

- We characterize the deviation of the achieved consensus from the consensus that the legitimate agents would have achieved in the absence of untrustworthy agents
- We show that the convergence rate to the achieved consensus is geometric with a high probability

Deviation from True Consensus Value

- The achieved consensus value $z(0)\mathbf{1}$, with $z(0)$ in the convex hull of $x_i(0), i = 1, \ldots, m$, is due to the opinion dynamic

$$x_{\mathcal{L}}(t+1) = W_{\mathcal{L}}(t)x_{\mathcal{L}}(t) + W_{\mathcal{M}}(t)x_{\mathcal{M}}(t)$$

- The true consensus is the value that the legitimate agents would have reached if they had known their malicious neighbors, i.e, resulting from the opinion dynamic:

$$x_{\mathcal{L}}(t+1) = \bar{W}x_{\mathcal{L}}(t) + \mathbf{0}x_{\mathcal{M}}(t)$$

- The true consensus value is $\langle v, x_{\mathcal{L}}(0)\rangle$ where $v > 0$ is a stochastic vector such that $v'\bar{W} = v'$
- We have shown that for every legitimate agent ℓ and any $\delta > 0$:

$$\mathrm{Prob}\{\limsup_{t\to\infty} |x_\ell(t) - \langle v, x_{\mathcal{L}}(0)\rangle| \leq \Delta(\delta)\} \geq 1 - \delta$$

  where $\Delta(\delta)$ is proportional to the bound η, inversely proportional to δ, depends on the number of malicious agents $|\mathcal{M}|$ (linearly), decreases exponentially with the time T_0, and with the expected values $E_{\mathcal{L}}$ and $E_{\mathcal{M}}$ squared[§]

- We also have a high probability geometric convergence of $x_\ell(t)$ to the consensus value.

[§]M. Yemini, AN, A. Goldsmith, S. Gil "Characterizing Trust and Resilience in Distributed Consensus for Cyberphysical Systems," IEEE Transactions on Robotics (T-RO) 38 (1) 71–91, 2022, https://arxiv.org/abs/2103.05464

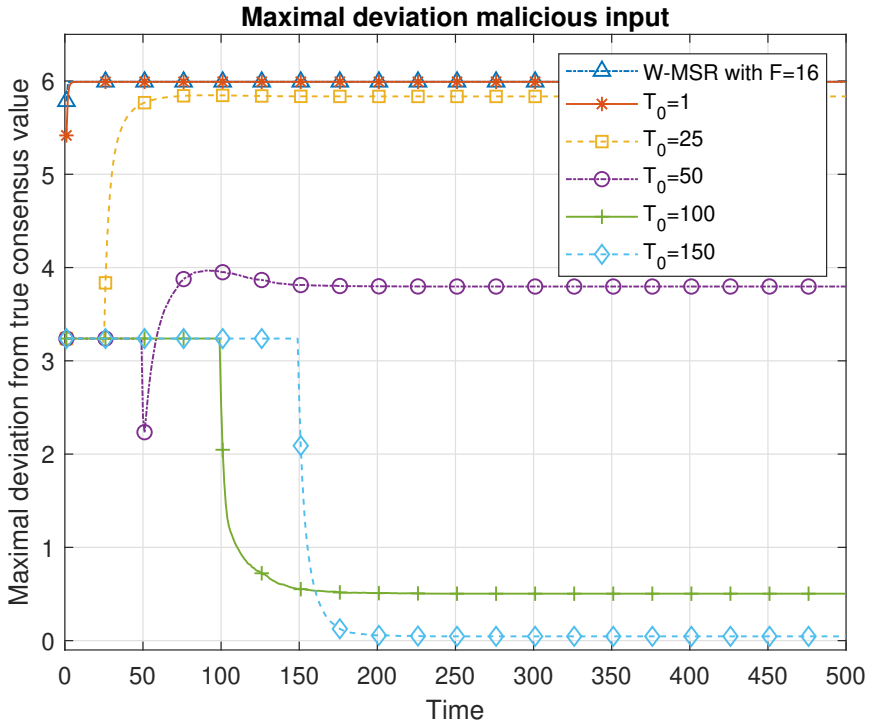

Figure 2: While the W-MSR algorithm (LeBlanc et al. 2013) cannot succeed in achieving true consensus in a graph with 15 legitimate and 30 malicious agents (200% malicious agents), our approach can.

The plot illustrates the maximal deviation achieved by the malicious agents in a system consisting of 45 agents, of which 30 are malicious. It compares our resilient consensus approach with the Weighted Mean Subsequence Reduced (W-MSR) algorithm of LeBlanc

et al. 2013 for the attack where malicious agents inject the input that maximizes the deviation from the true consensus value for the legitimate agents. The time T_0 is the initial time that legitimate agents spend learning the trustworthiness of their neighbors, while from T_0 onward, they start consensus and continue learning. The figure shows that the W-MSR algorithm never succeeds in achieving the true consensus, while our approach can achieve it even though the number of malicious agents is twice the number of the legitimate agents. (Here, $\alpha_{ij}(t)$ is a uniform over an interval, simulation results are averaged over 500 runs; the graph is a ring with extra random connections)

- Key takeaway: The legitimate agents can reach consensus without any restrictions on the number of the malicious agents.

- The number of malicious agents can affect the consensus point; deviation from the nominal consensus can be upper bounded with a high probability.

Distributed Optimization: "Untrustworthy" Agents

- The aim of the legitimate agents is to minimize distributively the sum of their (convex) objective functions over a (convex and closed) constraint set $X \subseteq \mathbb{R}^d$, i.e.,

$$x^\star_{\mathcal{L}} \in \operatorname*{argmin}_{x \in X} \sum_{i \in \mathcal{L}} f_i(x)$$

- A modified distributed gradient projection method takes the following form: for every legitimate agent $i \in \mathcal{L}$,

$$c_i(t) = w_{ii}(t) x_i(t) + \sum_{j \in \mathcal{N}_i} w_{ij}(t) x_j(t),$$

$$x_i(t+1) = \Pi_X \left[c_i(t) - \gamma(t) \nabla f_i(c_i(t)) \right],$$

  where $\Pi_X(\cdot)$ is the Euclidean projection on the set X, and $\gamma(t) \geq 0$ is a stepsize that is common to all agents $i \in \mathcal{L}$ at each time t

- The graph is static but the weights are dynamic
- The set N_i is composed of both legitimate and malicious neighbors of agent $i \in \mathcal{L}$, while the weights $w_{ii}(t)$ and $w_{ij}(t), j \in N_i$ are nonnegative and sum to 1
- Malicious agents' dynamic is not modeled

Assumptions

- The assumptions on the trust observations are the same
- We assume that $X \subset \mathbb{R}^d$ is closed, convex, and bounded, i.e., there is $\eta > 0$ such that

$$\|x\| \leq \eta \quad \text{for all } x \in X$$

- The η value is assumed to be known, and its role is to bound the malicious agents' inputs away from infinity.
- The functions f_i for legitimate agents $i \in \mathcal{L}$ are strongly convex and have Lipschitz continuous gradients
- The stepsize $\gamma(t)$ is diminishing with

$$\sum_t \gamma(t) = \infty, \qquad \sum_t \gamma^2(t) < \infty$$

- The subgraph induced by the legitimate agents is connected.

Under the strong convexity assumption, we have that

- The optimization problem $\min_{x \in X} \sum_{i \in \mathcal{L}} f_i(x)$ has a unique solution $x^*_{\mathcal{L}} \in X$.

Trust Values & Algorithm

- Using the robust consensus directly - not going to work since the resulting weights are not doubly stochastic which would lead to solving a biased problem with the objective of the form $\sum_{i\in\mathcal{L}} v_\ell f_\ell$, with v being the left eigenvector associated with the eigenvalue 1 of the limiting stochastic matrix $\bar{W}$.
- Construction of doubly stochastic (random) matrices in a long run: each legitimate agent $i \in \mathcal{L}$ uses the sum $\beta_{ij}(t)$ of stochastic observations $\alpha_{ij}(k), k = 1, \ldots, t$ and defines

$$N_i(t) = \{j \in N_i \mid \beta_{ij}(t) \geq 0\}$$

- Agents send $d_i(t) = |N_i(t)| + 1$ together with $x_i(t)$ in each round to their neighbors
- Each legitimate agent $i \in \mathcal{L}$ defines the weights:

$$w_{ij}(t) = \begin{cases} \frac{1}{2\cdot\max\{d_i(t),d_j(t)\}} & \text{if } j \in N_i(t), \\ 0 & \text{if } j \notin N_i(t) \cup \{i\}, \\ 1 - \sum_{\ell\in N_i(t)} w_{i\ell}(t) & \text{if } j = i. \end{cases}$$

- Algorithm assumes the form: for legitimate agent i

$$c_i(t) = w_{ii}(t)x_i(t) + \sum_{j \in N_i(t)} w_{ij}(t)x_j(t),$$

$$x_i(t+1) = \Pi_X \left[c_i(t) - \gamma(t) \nabla f_i(c_i(t)) \right],$$

  initiated with an arbitrary $x_i(0) \in X$.

Results[¶]:

- For such weights, there is a finite (random) time T_f such that the (random) matrix $[w_{ij}(t),\ i \in \mathcal{L},\ j \in N_i \cap \mathcal{L}]$ is doubly stochastic and static for $t \geq T_f$
- Consequently, the iterates $\{x_i(t)\}$, $i \in \mathcal{L}$, converge almost surely to the solution $x^*_{\mathcal{L}}$
- Also, they converge in the r-th mean for every $r \geq 1$:

$$\lim_{t \to \infty} \mathbf{E} \left[\|x_i(t) - x^\star_{\mathcal{L}}\|^r \right] = 0 \quad \text{for all } r \geq 1.$$

- The convergence rate of the 2-mean is of the order of $1/t$

[¶]M. Yemini, AN, A. J. Goldsmith, S. Gil, "Resilient Distributed Optimization for Multi-Agent Cyberphysical Systems" https://arxiv.org/abs/2212.02459; to appear in TAC

Convergence Rate in Mean-Square

Under our assumptions, we use the stepsize $\gamma(t) = \frac{2}{\mu(1-\epsilon)(t+2)}$, where $\epsilon \in (0,1)$ is a parameter. Then, for every $T_0 \geq 0$ and $T > T_0$, we have

$$\frac{1}{|\mathcal{L}|}\sum_{i\in\mathcal{L}} \mathbb{E}[\|x_i(T) - x^*_{\mathcal{L}}\|^2] \leq \min\left\{4\eta^2, \frac{4C_e(T-T_0) + C_{\mathcal{M}}(T_0)}{\mu(1-\epsilon)(T-T_0)^2}\right\},$$

where

- $C_e(T-T_0)$ captures the error due to decentralization; not depending on the malicious agent number nor their inputs; it grows linearly with T

- $C_{\mathcal{M}}(T_0)$ is a function that decreases exponentially with T_0 and captures the aggregate influence of malicious agents

Thus, the convergence rate is of the order $O(1/T)$

Simulation Illustration

- We simulated the same scenario of 15 legitimate and 30 malicious agents
- With the same graph structure as for the resilient consensus
- Each legitimate agent has a cost function of the form:

$$f_i(x) = \frac{1}{2}(\langle a_i, x\rangle - b_i)^2 + \frac{\lambda}{2}\|x\|^2, \qquad \lambda = 0.1$$

- The optimal point is determined for the unconstrained problem, then the constraint set is chosen to contain the optimal point
- The constraint set $X = \{x \in \mathbb{R}^d \mid \|x\| \leq \eta\}$, with $\eta = 30$ and $d = 5$.
- The stepsize is set to $\gamma(t) = \frac{2}{0.9(t+2)}$ (corresponds to $\epsilon = 0.1$).
- Simulations run for 100 realizations.

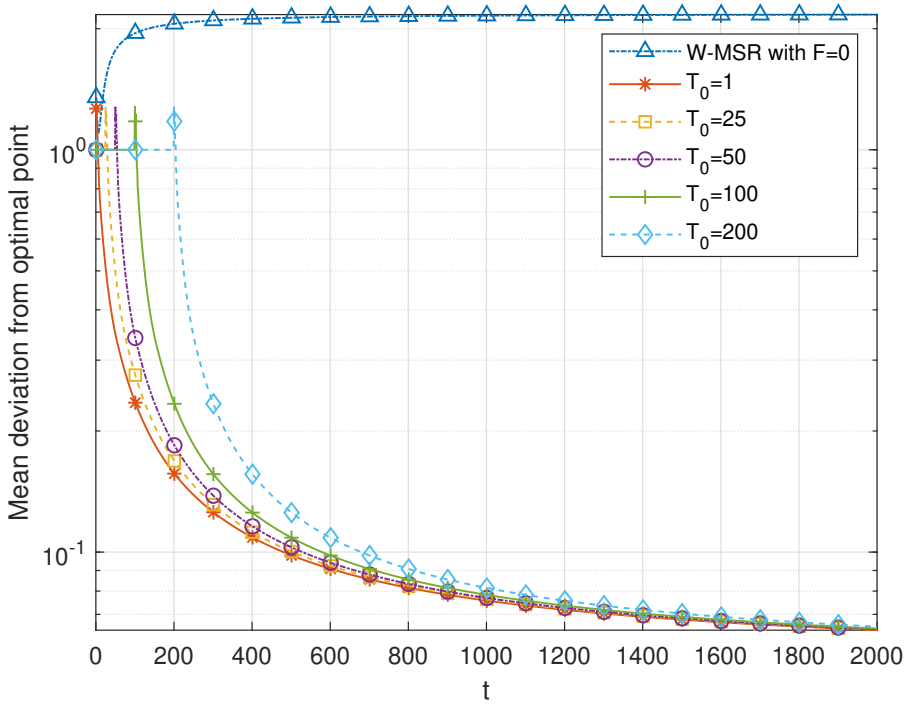

Figure 3: **The W-MSR-based algorithm does not converge to the optimal point in a graph with 15 legitimate and 30 malicious agents, our algorithm does.**

The W-MSR method[‖] not really specially tailored to solve the problem.

Key Observation: Despite the presence of the malicious agents, the method converges almost surely to the solution. The convergence rate is affected as seen from the theory.

[‖] S. Sundaram and B. Gharesifard, "Distributed optimization under adversarial nodes," IEEE Trans. Automat. Contr., vol. 64, no. 3, pp. 1063–1076, 2019.

Conclusion

- Using side information (trust observations), agents can learn trustworthy neighbors
- Leads to robust consensus and optimization in a static undirected graph
- We have ongoing work for the case of learning trustworthy neighbors in a static directed graph:
  O. E. Akgun, A. K. Dayi, S. Gil, and AN, "Learning Trust Over Directed Graphs in Multiagent Systems," L4DC 2023 Conference

- Additional Aspects:
  - Trust observations from physical channels[**]
  - Optimization over a directed static graph[††] - challenging

- Open directions:
  - Learning model - what if observations $\{\alpha_{ij}(t)\}$ are not independent over time?
  - What if the set of malicious agents is not static? Started some work on that[‡‡] but far from the finish
  - What if the underlying communication graph is time-varying?

[**]S. Gil, M. Yemini, A. Chorti, AN, H. V. Poor, A. J. Goldsmith "How Physicality Enables Trust: A New Era of Trust-Centered Cyberphysical Systems," on arxiv

[††]A. K. Dayı, O. E. Akgün, M. Yemini, S. Gil, AN "Fast Distributed Optimization over Directed Graphs under Malicious Attacks using Trust," on arxiv

[‡‡]S. Aydın, O.E. Akgün, S. Gil, and A. Nedić, "Multi-Agent Resilient Consensus under Intermittent Failures and Attack" CDC 2024

THANK YOU!