Exploiting Asynchrony in Multi-agent Consensus to Change the Agreement Point*

Robotics Track

Sasanka Nagavalli
Carnegie Mellon University
Pittsburgh, Pennsylvania
snagaval@andrew.cmu.edu

Ramitha Sundar
Carnegie Mellon University
Pittsburgh, Pennsylvania
rsundar@andrew.cmu.edu

Katia Sycara
Carnegie Mellon University
Pittsburgh, Pennsylvania
katia@cs.cmu.edu

ABSTRACT

Reaching agreement through consensus is fundamental to the operation of distributed systems such as sensor networks, social networks or multi-robot networks. Consensus requires agents in the system to reach an agreement over a variable of interest only through local interactions. In real systems, limitations in resources available to the individual agents and delays in communication typically result in asynchronous discrete time control models for consensus. In this paper, we model the problem where an adversary (or a friend) can exploit asynchrony in updates in a group of agents that use the same control law. By modifying the update frequency of a subset of the agents, the adversary (or friend) can change the final value that the system agrees on.

KEYWORDS

robot communication and teamwork; networked robot/sensor systems, distributed robotics; swarms and collective behavior

ACM Reference Format:
Sasanka Nagavalli, Ramitha Sundar, and Katia Sycara. 2018. Exploiting Asynchrony in Multi-agent Consensus to Change the Agreement Point. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Multi-agent consensus [12], a process by which multiple agents reach agreement on the value of a variable using only local information, is essential to the operation of many distributed systems including sensor networks [10], social networks [1] and multi-robot systems such as robotic swarms [4]. Global consensus may be achieved at a particular time instant or it may be achieved asymptotically as differences between the values assigned by each agent to the variable of interest decrease over time. Mathematically, the consensus problem has been studied in a wide variety of settings and can be characterized along multiple dimensions including but not limited to (a) continuous-time or discrete-time updates [5], (b) synchronous or asynchronous updates [3], (c) fixed or time-varying network topology [9, 15], (d) adversarial agents or cooperating agents [13], (e) homogeneous or heterogeneous agents [11].

In this paper, we study a multi-agent team that performs discrete-time averaging consensus, comprised of agents that are unaware of the presence of agents compromised by an adversary (e.g. robotic swarm with mole agents [14] or with hacked members). The agents on our team have a fixed interaction topology (i.e. each agent has the same set of neighbors over time), but they are heterogeneous in the sense that each agent executes the discrete-time consensus update (i.e. sets its current state to the average of its own and its neighbors' states) with a possibly different frequency. In addition, there is a different delay before the first update for each agent on the team. These delays enable us to model real world features such as initialization delays, timers or deliberate input delays. It has been shown in the literature that delaying an input to a robotic swarm can improve the performance of the swarm [7] — a phenomenon known as Neglect Benevolence — and in some cases it is only possible to meet a deadline by applying a sufficient delay [6, 8].

Asynchrony is present in most real-world multi-robot systems and presents many challenges [2, 3], such as that the agreement point (if the system even converges) may not be the average of the initial values. In this paper, we investigate how asynchrony in a multi-robot system that periodically applies standard consensus updates can actually be exploited to change the agreement point of the team. We show that by *only changing the update periods or initial update delays* of a subset of agents on the team, we can influence the agents to move towards a desired agreement point of our choosing. Studying the effects of changing update periods and delays is essential to understanding potential vulnerabilities in distributed robotic consensus-based systems to (a) enable mitigation strategies to thwart adversarial influence, or (b) enable strategies for beneficial effect (e.g. introduce delay to improve performance without changing the update rules or network topology).

2 PROBLEM FORMULATION

2.1 System Model

We consider a team of n agents with unique identifiers $\mathcal{I} = \{1, \ldots, n\}$ interacting with each other. The initial state of agent i is given by $\mathbf{x}_i(0) \in \mathbb{R}^m$ and joint initial state of the team is given by $\mathbf{x}(0)$. Each agent in the team has a fixed set of neighbors. The graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ captures the neighbor information, where the nodes $\mathcal{V} = \{v_1, \ldots, v_n\}$ represent agents and each edge $(v_i, v_j) \in \mathcal{E} \subseteq (\mathcal{V} \times \mathcal{V})$ represents that agent i can access the state information of agent j and we assume $(v_i, v_i) \in \mathcal{E}$. The adjacency matrix for this graph is given by $\mathbf{A} \in \mathbb{R}^{n \times n}$ and contains non-negative binary elements such that $A_{ij} = 1$ if $(v_i, v_j) \in \mathcal{E}$ (otherwise, $A_{ij} = 0$). The

*This research has been sponsored in part by AFOSR Grant FA9550-15-1-0442, ARI Award W911NF-16-1-0514 and an NSERC PGS D scholarship.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

set of neighbours of node v_i is $\mathcal{N}_i = \left\{v_j \in \mathcal{V} \mid \left(v_i, v_j\right) \in \mathcal{E}\right\}$. Note that $\forall i : v_i \in \mathcal{N}_i$.

Each agent i on our team periodically updates of its own state to the average its own state and the states of neighbor agents $\mathcal{N}_i$. In contrast to many consensus systems, our agents executes the same update rule, but with a possibly different period $T_i \in \mathbb{Z}$ and initial delay $\tau_i \in \mathbb{Z}$. The set of agents that execute an asynchronous discrete-time consensus update at time t is given as follows.

$$U\left(t, T_1, \ldots, T_n, \tau_1, \ldots, \tau_n\right) = \left\{j \in \mathcal{I} \mid t - \tau_j \equiv 0 \mod T_j\right\} \tag{1}$$

Due to limited space, we will drop the explicit dependence of this set on the update periods and delays and write it as $U(t)$. Now the state evolution of each agent can be described using the following recurrence relation where $i, j \in \mathcal{I}$ and $t \geq 0$ is the time.

$$\mathbf{x}_i(t+1) = \begin{cases} \frac{1}{\sum_{j=1}^{n} A_{ij}} \sum_{j=1}^{n} A_{ij}\mathbf{x}_j(t) & \text{if } i \in U(t),\ t \geq \tau_i \\ \mathbf{x}_i(t) & \text{otherwise} \end{cases} \tag{2}$$

2.2 Problem Statement

Imagine that we can influence a subset of agents on the team given by $C \subseteq \mathcal{I}$. However, we cannot change the update rule, the initial states or the set of neighbors for each agent. We can only change the update periods $\{T_i \mid \forall i \in C\}$ and the initial update delays $\{\tau_i \mid \forall i \in C\}$. As shown in Figure 1, changing these periods and delays can significantly change the final agreement point to which the team of agents converge. Consider the situation where we would like the final agreement point of these agents performing consensus to be as close as possible to a desired value $\mathbf{x}_g \in \mathbb{R}^m$. Let τ_{upper} and T_{upper} represent upper bounds on initial delay and period beyond which the influenced agents may be detected as adversarial. Our problem may then be written formally as follows.

$$\begin{aligned} &\underset{\forall i \in C: T_i, \tau_i}{\arg\min} && \sum_{k=1}^{n} \|\mathbf{x}_g - \lim_{t \to \infty} \mathbf{x}_k(t)\|_2^2 \\ &\text{subject to} && \\ &\forall k : \mathbf{x}_k(t+1) && = \begin{cases} \frac{1}{\sum_{j=1}^{n} A_{kj}} \sum_{j=1}^{n} A_{kj}\mathbf{x}_j(t) & \text{if } k \in U(t),\ t \geq \tau_k \\ \mathbf{x}_i(t) & \text{otherwise} \end{cases} \\ &\forall i \in C : && 0 \leq \tau_i \leq \tau_{\text{upper}} \\ &\forall i \in C : && 1 \leq T_i \leq T_{\text{upper}} \end{aligned}$$

2.3 Illustrative Example

Figure 1 presents an example that emphasizes the difficulty of this problem. The interaction topology is shown in Figure 1a. If agent states evolve according to the consensus update rule using their default update periods and initial delays, their states evolve as in Figure 1b. However, changing the update period and initial delay for even one agent (Agent 1) can change the final agreement point (Figure 1c). Figure 1d demonstrates that the choice of initial delay can change the minimal achievable distance to the adversary's desired goal $\mathbf{x}_g$. In addition, it shows that sometimes (but not always) delays (e.g. $\tau_1 = 12$ or $\tau_1 = 18$) are beneficial (the distance of the agreement point to the adversary goal is 0), so the system exhibits Neglect Benevolence. Figure 1e shows that choosing a different agent to influence can have dramatically different effects. It is also shown (Figure 1d) that the optimal delay time is not unique and that these optimal initial delays correspond to different update periods. Additionally and non-intuitively, this example shows that the best choice of agent to influence is not necessarily the one whose initial state begins closest to the value for the desired goal (i.e. agent 3), nor the one whose initial state begins furthest from the desired

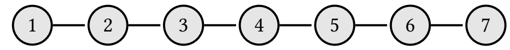

(a) Interaction topology for an example team of seven agents.

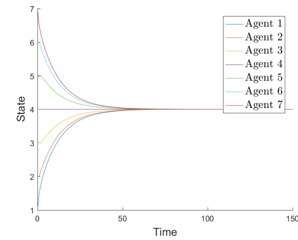

(b) State evolution of multi-agent team with default update periods. The agents converge to an agreement point 4.

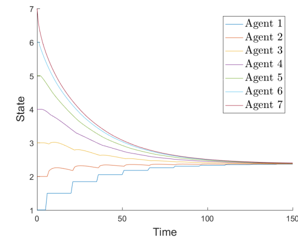

(c) State evolution of multi-agent team when $\tau_1 = 5$ and $T_1 = 15$. The agents converge to an agreement point 2.3751.

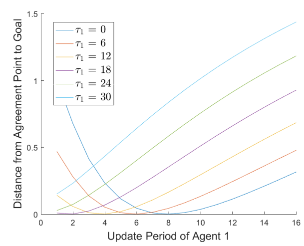

(d) Assume only agent 1 can be influenced. The choice of delay τ_1 changes the agreement point and the minimum achievable distance to the desired goal.

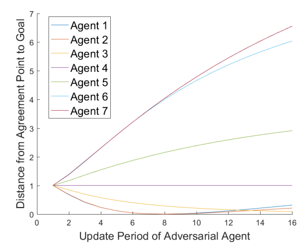

(e) Assume only one of the agents can be influenced. The choice of adversarial agent changes the minimum achievable distance to the desired goal.

Figure 1: A team of seven agents has the interaction topology given by the graph shown in (a). The joint initial state of the team is $\mathbf{x}(0) = [1, 2, 3, 4, 5, 6, 7]^\top$. The default update periods are $\forall i : T_i = 1$ and initial delays are $\forall i : \tau_i = 0$. The adversary's desired goal is $\mathbf{x}_g = 3$.

goal (i.e. agent 7). Instead, changing the update period of agent 1 or agent 2 seems to be the most effective way to minimize the distance between the team's agreement point and the desired goal. Finally, introducing an initial delay reduces the required change in the update period of the influenced adversarial agent. This very simple example shows that the problem is nonlinear and very challenging since it involves interplay of initial delay times, update periods and choice of influenced agents. As we have discovered in further experiments (not shown in the figure), the variance in initial values of the agents also influences the solution.

3 CONCLUSION

In this paper, we considered a team of agents that updated their states using the discrete-time consensus protocol, but with different update periods and delays prior to first update. Given a subset of agents on the team that an adversary could influence, we studied the problem of changing only their update periods and delays to change to the agreement point reached by the team. An adversary could use this technique to bring the team's agreement point close to the adversary's desired goal. In current work, we are developing algorithms to optimally solve this problem.

REFERENCES

[1] Daron Acemoglu and Asuman Ozdaglar. 2011. Opinion dynamics and learning in social networks. *Dynamic Games and Applications* 1, 1 (2011), 3–49.

[2] Yin Chen, Roberto Tron, Andreas Terzis, and Rene Vidal. 2010. Corrective consensus: Converging to the exact average. In *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 1221–1228.

[3] Lei Fang and Panos J Antsaklis. 2005. Information consensus of asynchronous discrete-time multi-agent systems. In *American Control Conference, 2005. Proceedings of the 2005*. IEEE, 1883–1888.

[4] Andreas Kolling, Phillip Walker, Nilanjan Chakraborty, Katia Sycara, and Michael Lewis. 2016. Human interaction with robot swarms: A survey. *IEEE Transactions on Human-Machine Systems* 46, 1 (2016), 9–26.

[5] Peng Lin and Yingmin Jia. 2009. Consensus of second-order discrete-time multi-agent systems with nonuniform time-delays and dynamically changing topologies. *Automatica* 45, 9 (2009), 2154–2158.

[6] Sasanka Nagavalli, Nilanjan Chakraborty, and Katia Sycara. 2017. On Time-Optimal Behavior Scheduling of Robotic Swarms for Achieving Multiple Goals. In *2017 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE.

[7] Sasanka Nagavalli, Shih-Yi Chien, Michael Lewis, Nilanjan Chakraborty, and Katia Sycara. 2015. Bounds of Neglect Benevolence in Input Timing for Human Interaction with Robotic Swarms. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 197–204.

[8] Sasanka Nagavalli, Lingzhi Luo, Nilanjan Chakraborty, and Katia Sycara. 2014. Neglect benevolence in human control of robotic swarms. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 6047–6053.

[9] Reza Olfati-Saber and Richard M Murray. 2004. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on automatic control* 49, 9 (2004), 1520–1533.

[10] Reza Olfati-Saber and Jeff S Shamma. 2005. Consensus filters for sensor networks and distributed sensor fusion. In *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC'05. 44th IEEE Conference on*. IEEE, 6698–6703.

[11] Amanda Prorok, M Ani Hsieh, and Vijay Kumar. 2016. Formalizing the impact of diversity on performance in a heterogeneous swarm of robots. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 5364–5371.

[12] Wei Ren, Randal W Beard, and Ella M Atkins. 2005. A survey of consensus problems in multi-agent coordination. In *American Control Conference, 2005. Proceedings of the 2005*. IEEE, 1859–1864.

[13] David Saldana, Amanda Prorok, Shreyas Sundaram, Mario FM Campos, and Vijay Kumar. 2017. Resilient consensus for time-varying networks of dynamic agents. In *American Control Conference (ACC), 2017*. IEEE, 252–258.

[14] Navyata Sanghvi, Sasanka Nagavalli, and Katia Sycara. 2017. Exploiting Robotic Swarm Characteristics for Adversarial Subversion in Coverage Tasks. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 511–519.

[15] Feng Xiao and Long Wang. 2006. State consensus for multi-agent systems with switching topologies and time-varying delays. *Internat. J. Control* 79, 10 (2006), 1277–1284.