

Dynamic max-consensus with local self-tuning

D. Deplano* M. Franceschelli* A. Giua*

* *Department of Electrical and Electronic Engineering (DIEE),
University of Cagliari, Italy.*

Emails: {diego.deplano,mauro.franceschelli,giua}@unica.it

Abstract: This work describes a novel control protocol for multi-agent systems to solve the dynamic max-consensus problem. In this problem, each agent has access to an external time-varying scalar signal and has the objective to estimate and track the maximum among all these signals by exploiting only local communications. The main strength of the proposed protocol is that it is able to self-tune its internal parameters in order to achieve an arbitrary small steady-state error without significantly affecting the convergence time. We employ the proposed protocol in the context of distributed graph parameter estimations, such as size, diameter, and radius, and provide simulations in the scenario of open multi-agent systems.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Dynamic consensus, max-consensus, distributed estimation, open multi-agent systems.

1. INTRODUCTION

In the context of multi-agent systems, there has been a significant interest in the design of distributed algorithms to solve the so-called consensus or agreement problem.

Problem of interest and motivation. The consensus problem consists in the design of local interaction rule driving the agents to agree upon a common value of interest. Historically, such common value has been considered to be a function of the initial state of the agents. More recently a variation of the problem that goes by the name of *dynamic consensus problem* has been considered, where the agents are required to converge to a state value which is a function of local time-varying reference signals fed to the agents. The current literature has paid major attention to the average value, for which Kia et al. (2019) have provided an insightful tutorial paper, but also other quantities have received some attention, such as the median value (Sanai Dashti et al., 2019; Vasiljevic et al., 2020), and the max/min value (Deplano et al., 2021b).

In a recent paper (Deplano et al., 2021b) we considered the previously unexplored dynamic max/min-consensus problem, which is also the focus of this paper. We propose an improved protocol for dynamic max/min-consensus that allows the designer to decouple convergence rate and steady-state error: this provides a way to jointly increase convergence rate and decrease steady-state error. Applications of max/min-consensus protocols are various and diversified, including monitoring and optimization (Iutzeler et al., 2012); distributed synchronization, such as time-synchronization (Dengchang et al., 2013) and target tracking (Petitti et al., 2011); network parameter estimation,

such as cardinality (Lucchese et al., 2015), diameter and radius (Garin et al., 2012; Oliva et al., 2016; Deplano et al., 2021a), as well as highest/lowest node degree (Borsche and Attia, 2010).

Related literature. The standard max-consensus protocols can be traced back to the works of Tahbaz-Salehi and Jadbabaie (2006) and Cortés (2008), in continuous-time and discrete-time frameworks, respectively. Since then, different approaches and settings were considered, such as switching topologies (Nejad et al., 2010), second-order dynamics (Zhang and Li, 2018), asynchronous, delayed and noisy communications (Agrawal et al., 2019; Muniraju et al., 2019), gossip-based or randomized approaches (Iutzeler et al., 2012; Franceschelli and Gasparri, 2019), as well as open multi-agent systems (Abdelrahim et al., 2017), anonymous and resilient networks (Wang et al., 2018; Muniraju et al., 2019; Shang, 2020). The dynamic version of the max-consensus problem has recently been studied by us in (Deplano et al., 2021b), wherein two protocols are presented: one achieves bounded tracking and steady-state error, while the other one achieves zero steady-state error but requires a high memory burden.

Main contributions. In this paper, we provide a self-tuning version of the protocol we have previously proposed in (Deplano et al., 2021b) to solve the dynamic max-consensus problem. The main novel features are:

- Tracking and steady-state errors are decoupled by design;
- The protocol achieves bounded steady-state error that can be made arbitrarily small;
- The protocol achieves bounded tracking error that can be traded-off for improved convergence time;
- The memory burden does not increase with the dimension of the network.

* This work was partially supported by the Fondazione Banco di Sardegna with the grant "Formal Methods and Technologies for the Future of Energy Systems", cup F72F20000350007.

D. Deplano is the corresponding author.

As a second contribution, we employ the proposed protocol in the context of open multi-agent systems in order to estimate and track some important time-varying parameters of the network, such as the number of agents, the radius, and the diameter. More precisely, we equip the algorithms proposed in (Deplano et al., 2021a) and (Deplano et al., 2021b) with the protocol proposed in this paper achieving the following improvements:

- The size-estimation technique employed in (Deplano et al., 2021b) is highly sensitive to the steady-state error of the dynamic max-consensus protocol, a problem that is completely solved by employing the protocol proposed in this paper.
- The radius/diameter estimation technique employed in (Deplano et al., 2021a) does make use of a static version of the max-consensus protocol, which requires a centralized step. The implementation of the novel protocol allows to avoid such a centralized step and enables the employment of the algorithm in the case of open multi-agent systems.

Structure of the paper. After introducing the notation used throughout the paper along with some theoretical preliminaries, Section 2 first formalizes the dynamic max/min-consensus problem and the main working assumptions, then it presents our self-tuning dynamic max-consensus (STDMC) protocol along with its theoretical characterization. Section 3 provides numerical simulations validating the theoretical characterization and employ the STDMC Protocol in the context of graph parameter estimation for open and time-varying network of agents. Concluding remarks are given in Section 4.

2. SELF-TUNING DYNAMIC MAX-CONSENSUS (STDMC) PROTOCOL

2.1 Notation and Preliminaries

We denote by \mathbb{R} and \mathbb{N} the sets of real numbers and positive integer numbers, respectively. Maximum and minimum of an n -elements vector $u = [u_1, \dots, u_n]^T$, are denoted by

$$\bar{u} = \max_{i=1, \dots, n} u_i, \quad \underline{u} = \min_{i=1, \dots, n} u_i. \quad (1)$$

We consider networks modeled by undirected graphs $\mathcal{G} = (V, E)$, where $V = \{1, \dots, n\}$ with $n \in \mathbb{N}$ is the set of *nodes*, and $E \subseteq V \times V$ is the set of *edges* connecting the nodes. The state and the input of the i -th agent are scalars and are denoted by $x_i \in \mathbb{R}$ and $u_i \in \mathbb{R}$, respectively.

A path between two nodes $i, j \in V$ in a graph is a sequence of consecutive edges $\pi_{ij} = (i, p), (p, q), \dots, (r, s), (s, j)$ where each successive edge shares a node with its predecessor. An undirected graph \mathcal{G} is said to be *connected* if there exists a path π_{ij} between any pair of nodes $i, j \in V$. The *diameter* $\delta_{\mathcal{G}}$ of a graph \mathcal{G} is defined as the length (number of edges) of the longest shortest path between any pair of nodes in the graph.

Nodes i and j are *neighbors* if there exists an edge $(i, j) \in E$, which represents a point-to-point communication channel between nodes i and j . The set of neighbors of the i -th node is denoted by $\mathcal{N}_i = \{j \in V : (i, j) \in E\}$. For sake of simplicity, we consider graphs without self-loops, i.e., $i \notin \mathcal{N}_i$, and define $\mathcal{N}_i^{\circ} = \mathcal{N}_i \cup \{i\}$.

2.2 Problem statement and working assumptions

Consider a network of n agents modeled as discrete-time dynamical systems with scalar state $x_i \in \mathbb{R}$ for $i = 1, \dots, n$. The i -th agent has access to a scalar time-varying reference signal $u_i \in \mathbb{R}$ and interacts with other agents according to an undirected graph $\mathcal{G} = (V, E)$ and a local interaction rule

$$x_i(k) = f_i(u_i(k), x_j(k-1) : j \in \mathcal{N}_i^{\circ}). \quad (2)$$

The *dynamic max-consensus problem* consists in the design of proper local interaction rules $f_i(\cdot)$ for estimating and tracking the maximum $\bar{u}(k) \in \mathbb{R}$ among the time-varying reference signals. The performance can be expressed in terms of the convergence time and the tracking error

$$e(k) = \max_{i \in V} |x_i(k) - \bar{u}(k)|. \quad (3)$$

Our only assumption concerns the boundedness of the reference signals' variation, which is deemed a reasonable assumption in the dynamic consensus literature. We define the variation of the reference signals in one step as

$$\Delta u_i(k) = u_i(k) - u_i(k-1), \quad \forall i \in V. \quad (4)$$

and, in a similar way, we define the variation of the maximum among the reference signals as

$$\overline{\Delta u}(k) = \bar{u}(k) - \bar{u}(k-1). \quad (5)$$

Assumption 1. The absolute variation of the reference signals in one step is bounded by a constant $\Pi \geq 0$, i.e.,

$$|\Delta u_i(k)| \leq \Pi, \quad \forall i \in V, \quad \forall k \geq 0. \quad (6)$$

If the reference signals are sampled versions of continuous-time signals, then by increasing the sampling frequency their variation in one iteration is reduced. Thus, for any signal with bounded variation there exists a sampling frequency such that Assumption 1 is also satisfied. We refer the interested reader to Remark 1 at the end of Section 2.3 for a discussion about the case when the bound on the input variation is not known a-priori, but it can be estimated in a distributed way by the agents.

2.3 Proposed self-tuning protocol

When the reference signals are assumed to be constant over time, i.e., $u(k) = u(0)$ for all $k \in \mathbb{N}$, the problem can be recast as a standard max-consensus problem by

$$x_i(k) = \max_{j \in \mathcal{N}_i^{\circ}} \{x_j(k-1)\}, \quad x_i(0) = u_i(0), \quad (7)$$

which has been proved to converge in finite-time and with zero error (Tahbaz-Salehi and Jadbabaie, 2006). Instead, when the reference signals are assumed to be time-varying, the strategy we have previously proposed (Deplano et al., 2021b) modifies the above interaction rule into

$$x_i(k) = \max_{j \in \mathcal{N}_i^{\circ}} \{x_j(k-1) - \alpha, u_i(k)\}. \quad (8)$$

where $\alpha > 0$ is a design parameter. Such protocol has been proved to converge in finite-time and with bounded error, both depending on the parameter α . More precisely, the

parameter α trades-off estimation errors for convergence time, with larger values of α resulting into faster convergence of the protocol but also higher estimation errors, and lower values of α resulting into smaller estimation errors but also slower convergence of the protocol.

The strategy proposed in this paper is that of equipping the update rule in eq. (8) with local and self-tuning parameters $\alpha_i(k)$ for $i \in V$, as follows

$$x_i(k) = \max_{j \in \mathcal{N}_i} \{x_j(k-1) - \alpha_j(k-1), u_i(k)\}, \quad (9)$$

Intuitively, one can expect that the parameters α_i should take large values when the maximum input varies and the tracking task becomes the priority, while they should take arbitrarily small values when the maximum input remains constant and the estimation accuracy becomes more important. This is correct, and we further clarify this qualitative reasoning by discussing two different cases:

- The maximum input \bar{u} is higher than all states x_i .
In this case, the agent i with the maximum input $u_i = \bar{u}$ updates its state to \bar{u} , regardless of the value of α_i . Thus, the parameter α_i can be an arbitrarily small value. Note that in this case, the state of the agent i increases (or remains the same) after the update.
- The maximum input \bar{u} is lower than all states x_i .
In this case, each agent i update its state to the maximum among $x_j - \alpha_j$ with $j \in \mathcal{N}_i$, which depends on the parameters α_j . Thus, the parameters α_j must be sufficiently large in order to guarantee the tracking of the maximum input. Note that in this case, the state of the agents eventually decreases if the maximum input continues decreasing. This process stops when the input remains constant (or increases) and the first case holds instead.

As a consequence of the above discussion, we propose to make the parameters α_i switch between two design parameters $\alpha^{\text{MAX}} \geq \alpha^{\text{MIN}} > 0$ according to

$$\alpha_i(k) = \begin{cases} \alpha^{\text{MAX}} & \text{if } x_i(k) < x_i(k-1) \\ \alpha^{\text{MIN}} & \text{otherwise} \end{cases}. \quad (10)$$

Due to the above update rule, if the state of agent i is decreasing, then its local parameter α_i is set to a high value α^{MAX} in order to speed up the convergence toward the maximum input that has a smaller value. Contrarily, if the state of agent i is increasing or constant, then its local parameter α_i is set to a small value α^{MIN} in order to improve the estimation accuracy. We describe the steps required to implement the interaction rule (9) with the local self-tuning (10) in the STDMC Protocol, shown in the column on the right, while the next theorem characterizes its convergence time and tracking error.

Remark 1. The agents can make sure to employ the same design parameters α^{MAX} , α^{MIN} and deal with unknown upper bounds Π on the inputs' variation $\Delta u_i(k)$ by locally tuning their local parameters α_i^{MAX} , α_i^{MIN} according to

$$\begin{aligned} \alpha_i^{\text{MAX}}(k) &= \max_{j \in \mathcal{N}_i} \{\alpha_j^{\text{MAX}}(k-1), \theta \cdot \Delta u_i(k)\}, \quad \theta > 1, \\ \alpha_i^{\text{MIN}}(k) &= \min_{j \in \mathcal{N}_i} \{\alpha_j^{\text{MIN}}(k-1)\} \end{aligned} \quad (11)$$

In such case, the characterization of the STDMC Protocol is provided by Corollary 1.

STDMC Protocol :

SELF-TUNING DYNAMIC MAX-CONSENSUS

(Input): Tuning parameters $\alpha^{\text{MAX}} \geq \alpha^{\text{MIN}} > 0$.

(Initialization): $x_i(0) \in \mathbb{R}$ for $i \in V$;
 $\alpha_i(0) \in \{\alpha^{\text{MIN}}, \alpha^{\text{MAX}}\}$ for $i \in V$

(Execution): for $k = 1, 2, 3, \dots$ each node i does

Gather $x_j(k-1)$ and $\alpha_j(k-1)$ from each neighbor $j \in \mathcal{N}_i(k-1)$
Update the current state according to $x_i(k) = \max_{j \in \mathcal{N}_i} \{x_j(k-1) - \alpha_j(k-1), u_i(k)\}$
Update the current parameter according to $\alpha_i(k) = \begin{cases} \alpha^{\text{MAX}} & \text{if } x_i(k) < x_i(k-1) \\ \alpha^{\text{MIN}} & \text{otherwise} \end{cases}$

Theorem 1. Consider a multi-agent system executing the STDMC Protocol under Assumption 1 and let $k = 0$ be the initial time. If the graph \mathcal{G} is connected and if

$$\alpha^{\text{MAX}} > \Pi, \quad (12)$$

then there exists a convergence time $T_c \geq 0$ such that the tracking error is bounded for $k \geq T_c$ by

$$e(k) \leq (\alpha^{\text{MAX}} + \Pi)\delta_{\mathcal{G}}, \quad (13)$$

where $\delta_{\mathcal{G}}$ is the diameter of graph \mathcal{G} , and it holds

$$T_c \leq \delta_{\mathcal{G}} + \max \left\{ \frac{\bar{x}(\delta_{\mathcal{G}}) - \bar{u}(\delta_{\mathcal{G}})}{\alpha^{\text{MAX}} - \Pi}, 0 \right\}. \quad (14)$$

Proof sketch of Theorem 1. It can be shown¹ that there exists a convergence time $T_c \geq 0$ such that the maximum and minimum among the agents' state for $k \geq T_c$ are bounded by the following

$$\bar{x}(k) \leq \bar{u}(k) + (\Pi - \alpha^{\text{MIN}})\delta_{\mathcal{G}} \quad (15)$$

$$\underline{x}(k) \geq \bar{u}(k) - (\Pi + \alpha^{\text{MAX}})\delta_{\mathcal{G}}. \quad (16)$$

The bound (15) comes from the consideration that if $\bar{u}(k) \geq \bar{x}(k)$, then due to the update rule in eq. (9) it must hold $\bar{x}(k+1) = \bar{u}(k+1)$. otherwise, if the $\bar{u}(k) < \bar{x}(k)$ then their distance at subsequent times depends on their maximum variations. In the worst case, the maximum state decreases by α^{MIN} for $\delta_{\mathcal{G}}$ steps and the maximum input decreases by Π , leading to the bound in eq. (15). Indeed, after $\delta_{\mathcal{G}}$ steps, all the parameters α_i are updated to α^{MAX} according to the self-tuning in eq. (10), and thus the distance becomes smaller after $\delta_{\mathcal{G}}$ steps. The bound (16) comes from the consideration that the maximum distance between the minimum state \underline{x} and the maximum input \bar{u} occurs when during the tracking all α_i are set to α^{MAX} and the input suddenly reverses behavior and start increasing with maximum variation Π . Thus, their maximum distance is proportional to $\Pi + \alpha^{\text{MAX}}$, but after at most $\delta_{\mathcal{G}}$ steps the α_i are set to α^{MIN} according to the self-tuning in eq. (10), and thus the distance becomes smaller after $\delta_{\mathcal{G}}$ steps.

Note that the convergence time T_c is upper bounded by the diameter of the network, which is the largest time needed by the network to complete the cascade update. Indeed, $\delta_{\mathcal{G}}$ is the length of the longest shortest path between any pair of agents. Now, at time $k = \delta_{\mathcal{G}}$ two cases may occur:

¹ Formal proofs of the bounds in eqs. (15)-(16)-(17) are given in Appendix A, which is available at (Deplano et al., 2022).

- $\bar{x}(\delta_G) \leq \bar{u}(\delta_G)$. In this case, both bounds in eqs. (15)-(16) hold, and thus $T_c \leq \delta_G$.
- $\bar{x}(\delta_G) > \bar{u}(\delta_G)$. In this case, the maximum state \bar{x} decreases with rate α^{MAX} and, in the worst case, the maximum input decreases with rate Π . Due to eq. (12), then there exists a time T' such that $\bar{x}(\delta_G + T') \leq \bar{u}(\delta_G + T')$ and $T_c \leq \delta_G + T'$, given by

$$T' = \max \left\{ \frac{\bar{x}(\delta_G) - \bar{u}(\delta_G)}{\alpha^{\text{MAX}} - \Pi}, 0 \right\}. \quad (17)$$

From the above discussion one can verify the upper bound to the convergence time given in eq. (14). Now, exploiting the upper bound in eq. (15) and the lower bound in eq. (16), the bound on the tracking error can be derived for $k \geq \delta_G + T'$ as follows

$$e(k) \leq \max\{(\Pi - \alpha^{\text{MIN}})\delta_G, (\Pi + \alpha^{\text{MAX}})\delta_G\} \leq (\Pi + \alpha^{\text{MAX}})\delta_G$$

completing the proof. \square

Corollary 1. Consider a multi-agent system executing the STDMC Protocol under Assumption 1. Assume that the maximum input variation Π is not known and that α^{MAX} is locally updated according to eq. (11). If the graph \mathcal{G} is connected, then there exists a convergence time $T_c \geq 0$ such that the tracking error is bounded for $k \geq T_c$ by

$$e(k) \leq (\theta + 1)\Pi\delta_G, \quad (18)$$

where δ_G is the diameter of graph \mathcal{G} , and it holds

$$T_c \leq k_0 + \delta_G + \max \left\{ \frac{\bar{x}(k_0 + \delta_G) - \bar{u}(k_0 + \delta_G)}{\theta}, 0 \right\}, \quad (19)$$

where k_0 is the time at which the variation of one of the inputs is maximum, i.e., $\exists i \in V$ such that $\Delta u_i(k_0) = \Pi$.

The characterization provided in Theorem 1 reveals that the convergence time is inversely proportional to α^{MAX} , while the tracking error is directly proportional to α^{MAX} . Therefore, α^{MAX} trades off convergence time for tracking error, which are not influenced by the choice of α^{MIN} . On the other hand, α^{MIN} does affect the steady-state error when the maximum among the inputs does not change for a sufficiently large amount of time, as formalized next.

Theorem 2. Consider a generic time $k_0 \geq T_c$. The estimation error for Theorem 1, in the case all reference signals remain constant for $k \geq k_0$, satisfies the next condition

$$e(k) \leq \alpha^{\text{MIN}}\delta_G, \quad k \geq k_0 + 2\delta_G. \quad (20)$$

Proof. By Theorem 1, for any $k \geq T_c$ it holds $\bar{x}(k) \leq \bar{u}(k) + (\Pi - \alpha^{\text{MIN}})\delta_G$, $k \geq T_c$. If this upper bound holds strictly at k , i.e., $\bar{x}(k) = \bar{u}(k) + (\Pi - \alpha^{\text{MIN}})\delta_G$, then:

- All $\alpha_i(k)$ are exactly equal to α^{MAX} ;
- $\bar{u}(k)$ is decreasing, i.e., $\bar{u}(k) < \bar{u}(k - \delta_G)$.

In the worst-case scenario, this is the case at $k = k_0$, since then the inputs remain constant and equal to $u(k_0)$. Therefore, the maximum input starts decreasing at each time step ℓ by a factor α^{MAX} , until the maximum input is reached, i.e., $\bar{x}(k_0 + \ell) \leq \bar{u}(k_0)$, indeed,

$$\begin{aligned} \bar{x}(k_0 + \ell) &= \bar{x}(k_0) - \ell\alpha^{\text{MAX}} \\ &\leq \bar{u}(k_0) + (\Pi - \alpha^{\text{MIN}})\delta_G - \ell\alpha^{\text{MAX}} \leq \bar{u}(k_0) \end{aligned}$$

and since $(\Pi - \alpha^{\text{MIN}}) \leq \alpha^{\text{MAX}}$, it follows that the smallest value satisfying the above inequality is $\ell = \delta_G$.

Now, once δ_G steps have elapsed, the parameter of the agent with the maximum input is updated to α^{MIN} due to the self-tuning rule in eq. (10) for $k \geq k_0 + \delta_G$. Due to the local interaction rule in eq. (9), its 1-hop neighbors update their state to either $\bar{x}(k) - \alpha^*$ with $\alpha^* \in \{\alpha^{\text{MIN}}, \alpha^{\text{MAX}}\}$ or to their own input u_i . In both cases, according to the update rule in eq. (10), at subsequent instant of times their local parameter is updated to the value of α^{MIN} . The cascade effect updates all parameters α_i with $i \in V$ to α^{MIN} in at most δ_G steps. We conclude that

$$\bar{u}(k_0) \geq \bar{x}(k) > \underline{x}(k) = \bar{u}(k) - \alpha^{\text{MIN}}\delta_G, \quad k \geq k_0 + 2\delta_G,$$

from which the statement of the theorem follows. \square

Remark 2. A pragmatic choice for the parameter α^{MIN} is $\alpha^{\text{MIN}} \ll \Pi$, which reduces the steady-state error, even if it is not required for the convergence of the STDMC Protocol and for the validity of Theorems 1-2. However, the intuitive choice $\alpha^{\text{MAX}} \gg \Pi$ greatly reduces the convergence time but it may involve overshoots in the tracking error, see Example 3 in Section 3.

3. NUMERICAL SIMULATIONS

We test the results for the STDMC Protocol by considering the worst-case scenario of a network with line topology, which maximizes the number of steps needed to make the information flow through the network, that is $\delta_G = n - 1$ steps, thus maximizing the bounds on both the tracking and the steady state error errors.

3.1 Example 1: convergence time and steady-state error

We show in Fig. 1 the state evolution of a network with $n = 10$ agents in a line configuration, with diameter $\delta_G = 9$, executing the STDMC Protocol when the maximum input variation Π is known (red curve) and unknown

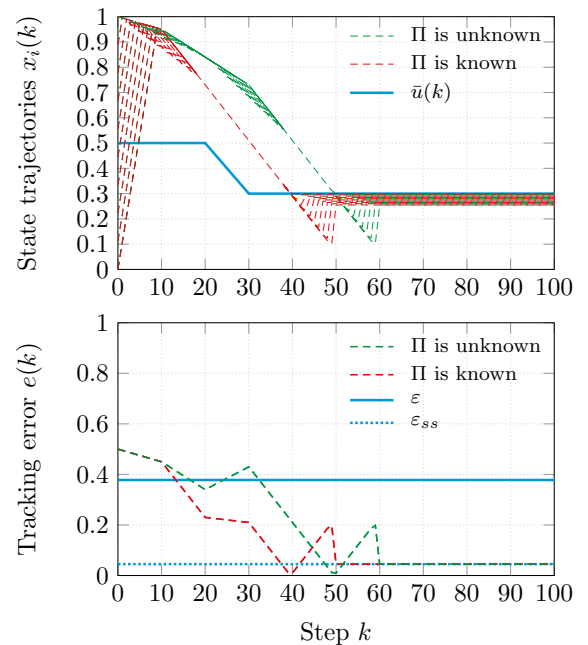


Fig. 1. *Example 1:* Evolution of a network with line topology running the STDMC Protocol; ε and ε_{ss} are the bound on the tracking and steady-state error defined in eq. (21) and provided in Theorems 1-2

(green curve). The agents are uniformly initialized within $[0, 1]$, the reference signals remain constant and equal to 0, except for the 6-th input that is initialized at $u_6(0) = 0.5$ and decreases with rate $\Pi = 0.02$ for $k \in [20, 30)$. When the bound Π is known, the parameters of the STDMC Protocol are chosen accordingly to Theorem 1 and are given by

$$\alpha^{\text{MAX}} = 0.022 > \Pi, \quad \alpha^{\text{MIN}} = 0.005,$$

and Fig. 1 shows the state evolution (red curve). This example allows to verify the validity of the bound ε on the tracking error and the bound ε_{ss} on the steady-state error given in Theorems 1-2, i.e.,

$$\varepsilon = (\alpha^{\text{MAX}} + \Pi)\delta_G = 0.378, \quad \varepsilon_{ss} = \alpha^{\text{MIN}}\delta_G = 0.045. \quad (21)$$

When the bound Π is not known, the agents may also tune the value of $\alpha^{\text{MAX}}(k)$ according to eq. (11) given in Remark 1 with $\theta = 1.1$. We show in Fig. 1 the state evolution (green curve) when the initial guess of the bound is wrongly made according to

$$\alpha_i^{\text{MAX}}(0) = 0.011 < \Pi, \quad \forall i \in V.$$

The 6-th agent updates its parameter α_6^{MAX} to $\theta\Pi\delta_G = 0.378$ at step $k_0 = 20$ when it observes that its input varies with rate Π and, accordingly to Corollary 1, the bound on the tracking error $\varepsilon = (1 + \theta)\Pi = 0.378$ remains unchanged, but the convergence time increases.

3.2 Example 2: tracking error and comparison with non self-tuning dynamic max-consensus

We compare in Fig. 2 the state evolution of a network with $n = 10$ agents in a line configuration executing the STDMC Protocol with self-tuning (red curve) and without self-tuning (green curve). The agents are uniformly initialized in the interval $[0, 1.5]$ and all reference signals remain constant and equal to -1 , except for the 6-th agent. The input of the agent 6 is initialized at $u_6(0) = 0.5$ and varies according to




Fig. 2. Example 2: Evolution of the network in Example 2 running the STDMC Protocol; ε is the bound on the tracking error.

$$u_6(k+1) = \begin{cases} u_6(k) - \Pi & \text{if } k \in [100, 150) \\ u_6(k) + \Pi & \text{if } k \in [150, 200) \\ u_6(k) & \text{otherwise} \end{cases}.$$

where $\Pi = 0.02$ is the maximum absolute variation as in Assumption 1. In Fig. 2, the red curves represent the estimation result provided by the STDMC Protocol when the self-tuning is designed, accordingly to Theorem 1, i.e.,

$$\alpha^{\text{MAX}} = 0.022 > \Pi, \quad \alpha^{\text{MIN}} = 10^{-12}.$$

Instead, in Fig. 2, the green curves represent the estimation result provided by the STDMC Protocol when the self-tuning is disabled, which amounts to the DMC Protocol previously proposed by us in (Deplano et al., 2021b) with constant parameter $\alpha = 0.022$. It is straightforward to notice how the self-tuning logic improves the accuracy of the estimation, without significantly affecting the convergence time. By looking at Fig. 2, one can also validate the characterization of the STDMC Protocol provided in Theorems 1-2. In particular, it can be verified that the tracking error is always bounded by

$$\varepsilon = (\Pi + \alpha^{\text{MAX}})\delta_G = 0.378,$$

and, more importantly, that the steady state error is almost nullified by the choice $\alpha^{\text{MIN}} = 10^{-12}$.

3.3 Example 3: graph parameters estimation in open multi-agent systems

In the third simulation shown in Fig. 3, we employ the proposed protocol as a subroutine of the algorithms presented in our previous works (Deplano et al., 2021a,b) for the estimation and track of some important graph parameters in open multi-agent systems, wherein the agents may join and leave as time goes by. For the convenience of the

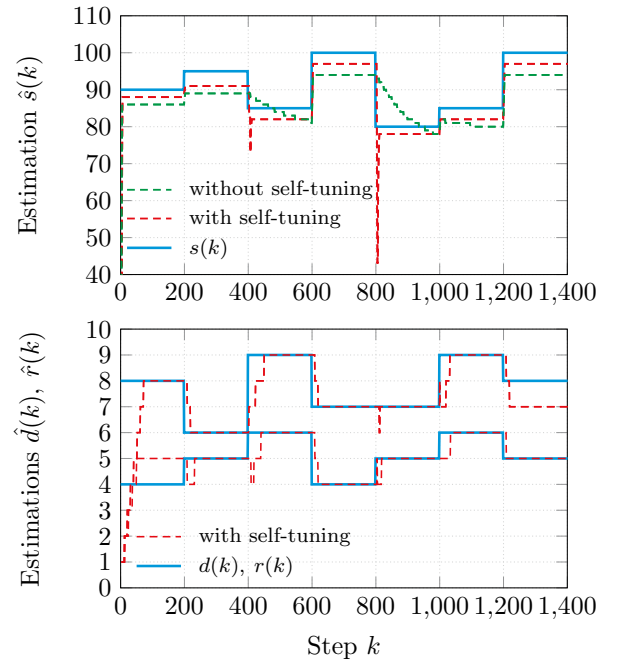


Fig. 3. Example 3: Estimation and tracking of the size (top) and radius/diameter (bottom) in an open multi-agent system running the protocols proposed in (Deplano et al., 2021a,b) implementing the STDMC Protocol as a subroutine.

readers, we detail both these algorithms in the unified SDR Protocol presented in Appendix B, which is available at (Deplano et al., 2022).

In particular, the algorithms enable to track the following time-varying parameters:

- Size $s(k)$: the number of agents within the network;
- Radius $r(k)$: the length of the minimum distance between any pair of agents in the network;
- Diameter $d(k)$: the length of the maximum distance between any pair of agents in the network.

We consider a network of $n \leq 100$ agents that are allowed to join and leave the network every $\Upsilon = 200$ time steps, as well as establishing or closing communications with other agents. This leads to a network with a time-varying number of nodes, as well as time-varying diameter and radius. We denote by $\hat{s}(k)$, $\hat{d}(k)$, $\hat{r}(k)$ their estimations. The parameters are set as shown next

$$\alpha^{\text{MAX}} = 10^{-1}, \quad \alpha^{\text{MIN}} = 10^{-12}$$

Fig. 3 (top) shows the estimation $\hat{s}(k)$ (red dashed curve) of the network's size $s(k)$ (blue solid curve), showing how the algorithm is capable of tracking changes of the time-varying size, with fast convergence rate ruled by parameter α^{MAX} and high accuracy ruled by α^{MIN} .

We also show the estimation when the non self-tuning dynamic max-consensus (DMC) Protocol proposed in (Deplano et al., 2021b) is employed (green dashed curve), which amounts to select $\alpha^{\text{MAX}} = \alpha^{\text{MIN}} = \alpha^*$. In order to provide a fair comparison, we select the good trade-off $\alpha^* = 10^{-4}$.

The comparison reveals that the employment of the STDMC Protocol allows achieving both higher convergence rate and accuracy. Fig. 3 (bottom) shows the estimations $\hat{d}(k)$, $\hat{r}(k)$ (red dashed curve) of the network's diameter and radius $d(k)$, $r(k)$ (blue solid curve), respectively. Since the diameter is always greater or equal to the radius, the curves on top refer to the diameter estimation while the curves below refer to the radius estimation. Note that we cannot provide a comparison with the original algorithm proposed in (Deplano et al., 2021a), since it cannot be employed in an open network setting without the self-tuning logic proposed in this paper.

4. CONCLUSIONS

We have proposed the self-tuning dynamic max-consensus (STDMC) Protocol, which enables the agents to track the time-varying maximum value of a set of reference signals given as inputs to the agents. As the name suggests, it is capable of self-tune some internal parameters in order to minimize both tracking and steady-state errors, which are decoupled by design. Moreover, the tracking error can be traded-off for convergence time, without affecting the steady-state error. We further provided simulations when the STDMC Protocol is employed as a subroutine of two state-of-art algorithms to track size, the diameter, and the radius of the network in open and time-varying multi-agent systems. Future works will investigate the theoretical characterization of the latter algorithms with the implementation of the STDMC Protocol.

REFERENCES

- Abdelrahim, M., Hendrickx, J., and Heemels, W. (2017). Max-consensus in open multi-agent systems with gossip interactions. In *56th IEEE Conference on Decision and Control (CDC)*, 4753–4758.
- Agrawal, N., Frey, M., and Stańczak, S. (2019). A scalable max-consensus protocol for noisy ultra-dense networks. In *2019 IEEE 20th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 1–5.
- Borsche, T. and Attia, S.A. (2010). On leader election in multi-agent control systems. In *22th Chinese Control and Decision Conference (CCDC)*, 102–107.
- Cortés, J. (2008). Distributed algorithms for reaching consensus on general functions. *Automatica*, 44(3), 726–737.
- Dengchang, Z., Zhulin, A., and Yongjun, X. (2013). Time synchronization in wireless sensor networks using max and average consensus protocol. *International Journal of Distributed Sensor Networks*, 9(3), 192128–1/10.
- Deplano, D., Franceschelli, M., and Giua, A. (2021a). Distributed tracking of graph parameters in time-varying anonymous networks. In *60th IEEE Conference on Decision and Control (CDC)*, 6258–6263.
- Deplano, D., Franceschelli, M., and Giua, A. (2021b). Dynamic min and max consensus and size estimation of anonymous multi-agent networks. *IEEE Transactions on Automatic Control*.
- Deplano, D., Franceschelli, M., and Giua, A. (2022). Dynamic max-consensus with local self-tuning. URL <https://unica.it/unica/protected/361711/0/def/ref/DAT258110/>.
- Franceschelli, M. and Gasparri, A. (2019). Multi-stage discrete time and randomized dynamic average consensus. *Automatica*, 99, 69–81.
- Garin, F., Varagnolo, D., and Johansson, K.H. (2012). Distributed estimation of diameter, radius and eccentricities in anonymous networks. In *IFAC Proceedings Volumes*, volume 45, 13–18. Elsevier.
- Iutzeler, F., Ciblat, P., and Jakubowicz, J. (2012). Analysis of max-consensus algorithms in wireless channels. *IEEE Transactions on Signal Processing*, 60(11), 6103–6107.
- Kia, S., Van Scoy, B., Cortes, J., Freeman, R., Lynch, K., and Martinez, S. (2019). Tutorial on dynamic average consensus: The problem, its applications, and the algorithms. *IEEE Control Systems*, 39(3), 40–72.
- Lucchesse, R., Varagnolo, D., Delvenne, J., and Hendrickx, J. (2015). Network cardinality estimation using max consensus: The case of bernoulli trials. In *54th IEEE Conference on Decision and Control*, 895–901.
- Muniraju, G., Tepedelenlioglu, C., and Spanias, A. (2019). Analysis and design of robust max consensus for wireless sensor networks. *IEEE Transactions on Signal and Information Processing over Networks*, 5(4), 779–791.
- Nejad, B.M., Attia, S.A., and Raisch, J. (2010). Max-consensus in a max-plus algebraic setting: The case of switching communication topologies. In *IFAC Proceedings Volumes*, volume 43, 173–180.
- Oliva, G., Setola, R., and Hadjicostis, C.N. (2016). Distributed finite-time calculation of node eccentricities, graph radius and graph diameter. *Systems & Control Letters*, 92, 20–27.
- Petitti, A., Di Paola, D., Rizzo, A., and Cicerelli, G. (2011). Consensus-based distributed estimation for target tracking in heterogeneous sensor networks. In *50th IEEE Conference on Decision and Control (CDC) and European Control Conference (ECC)*, 6648–6653.
- Sanai Dashti, Z.A., Seatzu, C., and Franceschelli, M. (2019). Dynamic consensus on the median value in open multi-agent systems. In *58th IEEE Conference on Decision and Control (CDC)*, 3691–3697.
- Shang, Y. (2020). Resilient consensus in multi-agent systems with state constraints. *Automatica*, 122, 109288.
- Tahbaz-Salehi, A. and Jadbabaie, A. (2006). A one-parameter family of distributed consensus algorithms with boundary: From shortest paths to mean hitting times. In *45th IEEE Conference on Decision and Control (CDC)*, 4664–4669.
- Vasiljevic, G., Petrovic, T., Arbanas, B., and Bogdan, S. (2020). Dynamic median consensus for marine multi-robot systems using acoustic communication. *IEEE Robotics and Automation Letters*, 5(4), 5299–5306.
- Wang, A., Mu, N., and Liao, X. (2018). Min-max consensus algorithm for multi-agent systems subject to privacy-preserving problem. In *International Conference on Neural Information Processing*, 132–142.
- Zhang, Y. and Li, S. (2018). Second-order min-consensus on switching topology. *Automatica*, 96, 293–297.