Resilient Leader-Follower Consensus with Time-Varying Leaders in Discrete-Time Systems

James Usevitch and Dimitra Panagou

Abstract— **The problem of consensus in the presence of adversarially behaving agents has been studied extensively in the literature. The proposed algorithms typically guarantee that the consensus value lies within the convex hull of the initial states of the normally-behaving agents. In leader-follower consensus problems however, the objective for normally behaving agents is to track a time-varying reference state that may take on values outside of this convex hull. In this paper we present a method for agents with discrete-time dynamics to resiliently track a set of leaders' common time-varying reference state despite a bounded subset of the leaders and followers behaving adversarially. The efficacy of our results are demonstrated through simulations.**

I. Introduction

Guaranteeing the resilience of multi-agent systems against adversarial misbehavior and misinformation is a critically needed property in modern autonomous systems. The *resilient consensus problem* has been the focus of much research for the past few decades. In this problem, normally behaving agents in a multi-agent network seek to come to agreement on one or more state variables in the presence of adversarially behaving agents whose identity is unknown. Within the last decade, several algorithms based upon the *Mean-Subsequence-Reduced* family of algorithms [1] have become a popular means to guarantee consensus of the normally behaving agents when the number of adversaries is bounded and the network communication structure satisfies certain *robustness* properties. These discrete-time algorithms, which include the W-MSR, DP-MSR, SW-MSR, and QW-MSR algorithms [2]–[5], guarantee that the final consensus value of the normal agents is within the convex hull of the normal agents' initial states.

A related problem in prior control literature is the *leader-follower consensus problem*, where the objective is for normally behaving agents to come to agreement on the (possibly time-varying) state of a leader or set of leaders [6]–[8]. Prior work in this area typically assumes that there are no adversarially misbehaving agents; i.e. all leaders and followers behave normally. An interesting question to consider is whether the property of resilience can be extended to leader-follower consensus, i.e. whether follower agents can track the leader agents' states while rejecting the influence of adversarial agents whose identity is unknown. One aspect that prevents prior resilient consensus results from being extended to this case is that the (possible time-varying) leaders' states may not lie within the convex hull of the initial states of normally-behaving agents.

James Usevitch is with the Department of Aerospace Engineering, University of Michigan, Ann Arbor; `usevitch@umich.edu`. Dimitra Panagou is with the Department of Aerospace Engineering, University of Michigan, Ann Arbor; `dpanagou@umich.edu`. The authors would like to acknowledge the support of the Automotive Research Center (ARC) in accordance with Cooperative Agreement W56HZV-14-2-0001 U.S. Army TARDEC in Warren, MI, and of the Award No W911NF-17-1-0526.

Recent work most closely related to the resilient leader-follower consensus problem includes [9]–[12]. In [9], the problem of resilient distributed estimation is considered where agents apply a discrete-time resilient consensus algorithm to reduce the estimation error of scalar parameters of interest. The assumption is that certain agents have a precise knowledge of their own parameters. These "reliable agents" drive the errors of the remaining normal agents to the static reference value of zero in the presence of misbehaving agents. In [10], [13], the problem of distributed resilient estimation in the presence of misbehaving nodes is treated under a more general LTI model. The authors show conditions under which information about the decoupled modes of the system is resiliently transmitted from a group of source nodes to other nodes that cannot observe those modes. Their results guarantee exponential convergence to the reference modes of the system. In our prior work [12], we considered the case of leader-follower consensus to arbitrary *static* reference values using the W-MSR algorithm [2]. The resilient leader-follower consensus problem is closely related to the *secure broadcast* problem [14]–[16], where a dealer agent seeks to broadcast a *static message* to an entire network in the presence of misbehaving agents.

In this paper, we address the problem of resilient leader-follower consensus with *time-varying leaders* in the discrete-time domain. We make the following specific contributions:

- We introduce the Multi-Source Resilient Propagation Algorithm (MS-RPA). Under this algorithm, exact tracking of a set of time-varying leaders can be achieved by normally-behaving follower agents in a finite time, despite the presence of a bounded number of arbitrarily misbehaving follower *and leader* agents.
- We demonstrate conditions under which agents applying the MS-RPA algorithm achieve exact tracking of the leaders when agents are subject to input bounds.

Notation and relevant resilient consensus terms from prior literature are reviewed in Section II. The problem formulation is given in Section III. Our main results are outlined in Section IV. Simulations demonstrating our results are shown in Section V, and a brief conclusion is given in Section VI.

II. Notation and Preliminaries

The set of real numbers and integers are denoted $\mathbb{R}$ and $\mathbb{Z}$, respectively. The set of nonnegative reals and integers are denoted $\mathbb{R}_+$ and $\mathbb{Z}_+$, respectively. The cardinality of a set S is denoted $|S|$. The set union, intersection, and set difference operations of two sets S_1 and S_2 are denoted by $S_1 \cup S_2$, $S_1 \cap S_2$, and $S_1 \backslash S_2$ respectively. We denote $\bigcup_{i=1}^{n} S_i = S_1 \cup S_2 \cup \ldots \cup S_n$. We also denote the scalar ball of length $r \in \mathbb{R}$ at $x \in \mathbb{R}$ as $B(x, r) = \{z \in \mathbb{R} : |x - z| \leq r\}$. A digraph is denoted as $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ is the set of vertices or agents, and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of edges. An edge from i to j, $i, j \in \mathcal{V}$, denoted as $(i, j) \in \mathcal{E}$, represents the ability of i to send information to j. Note that for digraphs $(i, j) \in \mathcal{E} \not\Rightarrow (j, i) \in \mathcal{E}$. The set of in-neighbors of agent i is denoted $\mathcal{V}_i = \{j \in \mathcal{V} : (j, i) \in \mathcal{E})\}$. The set of out-neighbors of each agent i is denoted $\mathcal{V}_i^{out} = \{k \in \mathcal{V} : (i, k) \in \mathcal{E}\}$.

A. Resilient Consensus Preliminaries

We briefly review several definitions associated with the resilient consensus literature that will be used in this paper.

Definition 1 ([2]). *Let $F \in \mathbb{Z}_+$. A set $S \subset \mathcal{V}$ is* F-total *if it contains at most F nodes; i.e. $|S| \leq F$.*

Definition 2 ([2]). *Let $F \in \mathbb{Z}_+$. A set $S \subset \mathcal{V}$ is* F-local *if $\forall t \geq t_0$, $t \in \mathbb{Z}$, $|S \cap \mathcal{V}_i(t)| \leq F$ $\forall i \in \mathcal{V} \backslash S$.*

In words, a set S is F-local if it contains at most F nodes in the neighborhood of each node outside of S for all $t \geq t_0$.

The notion of strong r-robustness with respect to (w.r.t.) a subset S is defined as follows:

Definition 3 (Strong r-robustness w.r.t. S [10]). *Let $r \in \mathbb{Z}_+$, $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ be a digraph, and $S \subset \mathcal{V}$ be a nonempty subset. $\mathcal{D}$ is strongly r-robust w.r.t. S if for any nonempty $C \subseteq \mathcal{V} \backslash S$, C is r-reachable.*

Remark 1. *Given a subset $S \subset \mathcal{V}$, it can be verified in polynomial time whether $\mathcal{D}$ is strongly robust w.r.t. S [11].*

III. Problem Formulation

We consider a discrete-time network of n agents with the set of agents denoted $\mathcal{V}$. The communication links between the agents are modeled by the digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$. Without loss of generality we assume $t_0 = 0$, with $t \in \mathbb{Z}_+$. Each agent $i \in \mathcal{V}$ has state $x_i(t) \in \mathbb{R}$. In addition, each agent is designated to be either a leader agent or a follower agent.

Definition 4. *The set of leader agents is denoted $\mathcal{L} \subset \mathcal{V}$. The set of follower agents is denoted $S_f = \mathcal{V} \backslash \mathcal{L}$.*

Assumption 1. *The sets $\mathcal{L}$ and S_f are time-invariant and satisfy $\mathcal{L} \cup S_f = \mathcal{V}$ and $\mathcal{L} \cap S_f = \emptyset$.*

To ensure that agents' states $x_i(t)$ change at synchronous times, both leaders and followers only update their states $x_i(t)$ at time instances t in the set $\mathcal{T} = \{t \in \mathbb{Z}_+ : t \bmod \eta = 0\}$, where $\eta \in \mathbb{Z}_+$, $\eta > 0$ is a pre-specified parameter known to both leaders and followers. Each normally behaving leader agent l has the discrete-time dynamics

$$x_l(t) = r_L(\lfloor t/\eta \rfloor) \tag{1}$$

where $r_L : \mathbb{Z} \to \mathbb{R}$ is a time-varying reference function. The function $r_L(\cdot)$ is assumed to be only known to the leaders. In particular, at every t each leader knows the value of $r_L(\lfloor t/\eta \rfloor + 1)$, i.e. the future value the reference function will take on at the next state update time step.

Each normally behaving follower agent j has dynamics

$$x_j(t) = x_j(t-1) + \alpha(t-1) u_j(t-1) \tag{2}$$

where

$$\alpha(t) = \begin{cases} 1 & \text{if } t \bmod \eta = 0 \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

The control input functions $u_j : \mathbb{R} \to \mathbb{R}$ will be defined later. However, a bound of the form $|u_i(t)| \leq u_M$, $u_M \in \mathbb{R}_+$, $u_M > 0$ is imposed on the follower agents' inputs for all $j \in S_f$. The objective is for the follower agents' states $x_j(t)$ to achieve consensus with the time-varying signal $r_L(\cdot)$. To achieve this, each follower agent i has an internal estimate $r_i \in \mathbb{R}$, which represents its estimate of the reference function $r_L(\cdot)$. At each time t each leader $l \in \mathcal{L}$ is able to broadcast the value $r_L(\lfloor t/\eta \rfloor + 1)$ to its out-neighbors. Note that $r_L(\lfloor t/\eta \rfloor + 1)$ is the value which the reference function will be at the *next* state update time step. In addition, each follower agent i is able to receive the value r_j from each of its in-neighbors $j \in \mathcal{V}_i$ and is able to broadcast r_i to its out-neighbors.

However, instead of assuming that all agents apply nominally specified control laws, this paper considers the presence of *misbehaving agents*.

Definition 5. *An agent $k \in \mathcal{V}$ is* misbehaving *if at least one of the following conditions holds:*

1) *There exists a $t \geq 0$ where agent k does not update its state according to* (1) *and also does not update its state according to* (2).
2) *There exists a $t \geq 0$ where k communicates some value $\hat{r}_k \neq r_L(t)$ to at least one of its out-neighbors.*
3) *There exists a $t \geq 0$ where k communicates different values to different out-neighbors; i.e. k communicates r_{k_1} to out-neighbor i_1 and r_{k_2} to out-neighbor i_2 such that $r_{k_1} \neq r_{k_2}$.*

The set of misbehaving agents is denoted $\mathcal{A} \subset \mathcal{V}$.

Intuitively, misbehaving agents are agents which update their states arbitrarily or communicate arbitrary false information to their out-neighbors. By Definition 5, the set of misbehaving agents $\mathcal{A}$ includes faulty agents, malicious agents, and Byzantine agents [2]. Both followers *and leaders* are vulnerable to adversarial attacks and faults, and therefore the sets $\mathcal{A} \cap \mathcal{L}$ and $\mathcal{A} \cap S_f$ may possibly be nonempty. The following notation will be used:

Definition 6. *The set of agents which are* not *misbehaving are denoted $\mathcal{N} = \mathcal{V} \backslash \mathcal{A}$. Agents in $\mathcal{N}$ are referred to as* normally behaving agents.

Definition 7. *The set of normally behaving leaders is denoted as $\mathcal{L}^{\mathcal{N}} = \mathcal{L} \backslash \mathcal{A}$. The set of normally behaving followers is denoted $S_f^{\mathcal{N}} = S_f \backslash \mathcal{A}$.*

Definition 8. *The set of misbehaving leaders is denoted as $\mathcal{L}^{\mathcal{A}} = \mathcal{L} \cap \mathcal{A}$. The set of misbehaving followers is denoted as $S_f^{\mathcal{A}} = S_f \cap \mathcal{A}$.*

The purpose of this paper is to determine conditions under which normally behaving follower agents resiliently achieve consensus with the time-varying reference state of the set of normally-behaving leader agents in the presence of a nonempty set of misbehaving agents. Before giving the problem statement, we define the error function $e(t) = \max_{i \in S_f^{\mathcal{N}}, l \in \mathcal{L}^{\mathcal{N}}} |x_i(t) - x_l(t)|$.

Problem 1. *Determine conditions under which $e(t)$ is non-increasing for all $t \geq 0$, and there exists a finite $T \in \mathbb{Z}_+$ such that $e(t) = 0$ for all $t \geq 0 + T$, in the presence of a nonempty adversarial set $\mathcal{A}$.*

IV. Main Results

By Definition 5, misbehaving agents may either update their state arbitrarily, or communicate misinformation to their out-neighbors. This brings up several challenges for normal followers seeking to track the leaders' reference state:

1) Agents in $S_f^{\mathcal{N}}$ may have no knowledge about the set $\mathcal{L}$. It is not assumed that leader agents are identifiable to any follower agent.
2) The set of misbehaving leaders $\mathcal{L}^{\mathcal{A}}$ may possibly be nonempty, implying some normally behaving agents may receive misinformation from the leader set $\mathcal{L}$.
3) There may exist agents $i \in S_f^{\mathcal{N}}$ that are not directly connected to any leaders, i.e. $\mathcal{V}_i \cap \mathcal{L} = \emptyset$.

To address these challenges, we introduce the *Multi-Source Resilient Propagation Algorithm* with parameter F, which is described in Algorithm 1.

Remark 2. *The MS-RPA algorithm is based on the Certified Propagation Algorithm (CPA) in [14] (see also [17]), but contains several significant theoretical differences. First, the CPA algorithm considers a single source agent invulnerable to adversarial attacks propagating a reference value to the network, whereas the MS-RPA considers multiple sources that each are* vulnerable *to attacks (i.e. the set $\mathcal{L}$). Second, the CPA algorithm considers the broadcast of one static message to the network, whereas the MS-RPA considers a time-varying reference signal. Third, the CPA algorithm does not take any state update input bounds into account, whereas the MS-RPA algorithm explicitly considers input bounds.*

For brevity, let $\tau = \lfloor t/\eta \rfloor$. Our first Lemma demonstrates conditions such that on every time interval $[\tau\eta, (\tau+1)\eta - 1]$, the MS-RPA algorithm guarantees that each normal follower receives the reference value $r_L(\tau+1)$ before the next state update timestep $(\tau+1)\eta$.

Algorithm 1 MS-RPA WITH PARAMETER F:

Initialization:

1) Each leader $l \in \mathcal{L}$ begins with $x_l(0) = r_L(0)$
2) Each follower $i \in S_f$ begins with $x_i(0) \in \mathbb{R}$ and $u_i(0) = 0$.

At each time step $t \geq 0$, $t \in \mathbb{Z}_+$:

(For brevity, define $\tau = \lfloor t/\eta \rfloor$, where $(\tau+1)\eta$ is the next state update time step and $\tau\eta$ is the most recent state update time step.)

1) Each leader agent $l \in \mathcal{L}$ broadcasts the value $r_L(\tau+1)$ to its out-neighbors.
2) If $t \geq \tau\eta + 1$, each follower agent i which received some value from at least $F+1$ in-neighbors at the previous time $t-1$, with the value denoted $c \in \mathbb{R}$, sets
$$\begin{aligned} r_i(t) &= c \\ u_i(t) &= \frac{(r_i(t) - x_i(t))u_M}{\max(u_M, |r_i(t) - x_i(t)|)} \end{aligned} \tag{4}$$
where $u_M \in \mathbb{R}$, $u_M > 0$. Denote this set of follower agents as $\mathcal{C}(t)$.
3) If $t \leq (\tau+1)\eta - 1$, each agent $i \in \mathcal{C}$ broadcasts $r_i(t)$ to its out-neighbors.
4) Each agent $i \notin \mathcal{C}(t)$ sets $r_i(t) = r_i(t-1)$ and $u_i(t) = u_i(t-1)$, and does not broadcast any value.

Lemma 1. *Let $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ be a nonempty, nontrivial, simple digraph with S_f nonempty, and let $F \in \mathbb{Z}_+$. Suppose that $\mathcal{A}$ is an F-local set, $\mathcal{D}$ is strongly $(2F+1)$-robust w.r.t. the set $\mathcal{L}$, and all normally behaving agents apply the MS-RPA algorithm with parameter F. If $\eta > |S_f|$, then $\forall \tau \in \mathbb{Z}_+$ all normal agents receive $r_L(\tau+1)$ and set $r_i(t) = r_L(\tau+1)$ before the next state update time $(\tau+1)\eta$. More precisely, the reference $r_i(t)$ for each normally-behaving follower $i \in S_f^{\mathcal{N}}$ satisfies $r_i(t) = r_L(\tau+1)$ at time $t = (\tau+1)\eta - 1$.*

Proof. Consider the system at time step $t = \tau\eta$ for some $\tau \in \mathbb{Z}_+$. At this time step, each leader agent $l \in \mathcal{L}^{\mathcal{N}}$ broadcasts the value $r_L(\lfloor t/\eta \rfloor + 1) = r_L(\tau+1)$ to its out-neighbors. By the definition of strong $2F+1$-robustness, $|\mathcal{L}| \geq 2F+1$.

Consider the set $C_1 = S_f^{\mathcal{N}} = (\mathcal{V} \backslash \mathcal{L}) \backslash \mathcal{A}$. Note that $C_1 \subset \mathcal{V} \backslash \mathcal{L}$. By the definition of strong $(2F+1)$-robustness, there exists a nonempty subset $S_1 = \{i_1 \in C_1 : |\mathcal{V}_{i_1} \backslash C_1| \geq 2F+1\}$, with $(\mathcal{V}_{i_1} \backslash C_1) \subseteq (\mathcal{L} \cup \mathcal{A})$. As per Algorithm 1, the only agents which broadcast a value at time $\tau\eta$ will be either agents in $\mathcal{L}$ or agents in $\mathcal{A}$. This implies that agents in S_1 receive messages only from agents in the set $\mathcal{L} \cup \mathcal{A}$. Since $\mathcal{A}$ is F-local, it is impossible for any $i_1 \in S_1$ to receive any malicious value from more than F misbehaving in-neighbors. Therefore since $|V_{i_1} \cap (\mathcal{L} \cup \mathcal{A})| \geq 2F+1$ and $\mathcal{A}$ is F-local, each i_1 receives the value $r_L(\tau+1)$ from $F+1$ behaving leaders. Note that all i_1 will continue receiving $r_L(\tau+1)$ from at least $F+1$ behaving leaders for all times in the interval $[\tau\eta+1, (\tau+1)\eta - 1]$.

Let $t_1 = \tau\eta + 1$. At time t_1, each agent $i_1 \in S_1$ sets $r_i(t) = r_L(\tau+1)$ and broadcasts $r_L(\tau+1)$ to all its out-neighbors. Consider the set $C_2 = C_1 \backslash S_1$, which satisfies $C_2 \subseteq \mathcal{V} \backslash \mathcal{L}$. If C_2 is nonempty, by the definition of strong $2F+1$-robustness there exists a nonempty subset $S_2 = \{i_2 \in C_2 : |\mathcal{V}_{i_2} \backslash C_2| \geq 2F+1\}$, with $(\mathcal{V}_{i_2} \backslash C_2) \subseteq (S_1 \cup \mathcal{L} \cup \mathcal{A})$. As per Algorithm 1, the only agents which broadcast values at time t_1 will be agents in the set $S_1 \cup \mathcal{L} \cup \mathcal{A}$. Since $\mathcal{A}$ is F-

local, it is impossible for any $i_2 \in S_2$ to receive a malicious signal from more than F misbehaving in-neighbors. Instead, each agent $i_2 \in S_2$ receives the value $r_L(\tau+1)$ from at least $F+1$ behaving in-neighbors in the set $S_1 \cup \mathcal{L}^{\mathcal{N}}$. In addition, note that all agents in S_2 will continue receiving the value $r_L(\tau+1)$ from at least $F+1$ agents in $(S_1 \cup \mathcal{L})$ for all times in the interval $[t_1+1, (\tau+1)\eta - 1]$.

This process can be iteratively repeated: at each time step $t_p = \tau\eta + p$, $0 \le p < \eta$, each agent $i_p \in S_p$ will have received the value $r_L(\tau+1)$ from at least $F+1$ behaving in-neighbors in the previous timestep and therefore broadcast $r_L(\tau+1)$ to all its out-neighbors. The set C_{p+1} is iteratively defined as $C_{p+1} = C_p \backslash S_p$. By the definition of strong $(2F+1)$-robustness, if C_{p+1} is nonempty there exists a nonempty $S_{p+1} = \{i_{p+1} \in C_{p+1} : |\mathcal{V}_{i_{p+1}} \backslash C_{p+1}| \ge 2F+1\}$, with $\mathcal{V}_{i_{p+1}} \subseteq \left(\left(\bigcup_{j=1}^{p} S_j \right) \cup \mathcal{L} \cup \mathcal{A} \right)$. As per Algorithm 1, the only agents which broadcast values at time t_p will be agents in the set $\left(\left(\bigcup_{j=1}^{p} S_j \right) \cup \mathcal{L} \cup \mathcal{A} \right)$. Since $\mathcal{A}$ is F-local, it is impossible for any $i_{p+1} \in S_{p+1}$ to receive a misbehaving signal from more than F misbehaving in-neighbors. Instead, each agent $i_{p+1} \in S_{p+1}$ receives the value of $r_L(\tau+1)$ from at least $F+1$ behaving in-neighbors in $\left(\left(\bigcup_{j=1}^{p} S_j \right) \cup \mathcal{L} \right)$ and will continue receiving $r_L(\tau+1)$ for all times in the interval $[t_p+1, (\tau+1)\eta - 1]$.

Observe that C_p being nonempty implies $C_{p+1} < C_p$, since S_p is nonempty. Therefore there exists a finite $p' \in \mathbb{Z}_+$ such that $C_{p'}$ is empty. Since $C_p \subseteq S_f^{\mathcal{N}}$ for all $0 < p \le p'$, this implies that $p' \le |S_f^{\mathcal{N}}| \le |S_f| < \eta$. Let $t_{p'} = \tau\eta + p'$, and observe that $t_{p'} < (\tau+1)\eta$. An empty $C_{p'}$ implies that $\bigcup_{k=1}^{p'-1} S_k = S_f^{\mathcal{N}}$, and that all $i \in S_f^{\mathcal{N}}$ have $r_i(t_{p'}) = r_L(\tau+1)$. In addition, by prior arguments all $i \in S_f^{\mathcal{N}}$ and $l \in \mathcal{L}^{\mathcal{N}}$ will broadcast the value $r_L(\tau+1)$ for all timesteps $t \in [t_{p'}, \tau\eta - 1)$. Since $\mathcal{A}$ is F-local, no agent $i \in S_f^{\mathcal{N}}$ will set $r_i(t)$ to any other value than $r_L(\tau+1)$ on time steps $t \in [t_{p'}, (\tau+1)\eta - 1]$. This implies that $r_i((\tau+1)\eta) - 1) = r_L(\tau+1)$ for all normal agents $i \in S_f^{\mathcal{N}}$, which concludes the proof. □

In essence, Lemma 1 states that all agents are able to receive the leaders' reference value $r_L(\tau+1)$ before the next state update time $(\tau+1)\eta$. We next demonstrate conditions under which all normally-behaving followers are able to achieve consensus to the time-varying reference $r_L(\cdot)$ with the input bound u_M. We define the following functions:

$$m(t) = \min_{i \in S_f^{\mathcal{N}}} (x_i(t), x_l(t)), \quad M(t) = \max_{i \in S_f^{\mathcal{N}}} (x_i(t), x_l(t)).$$
$$V(t) = M(t) - m(t) \tag{5}$$

The following Theorem will use the Lyapunov candidate $V(t)$, which satisfies $V(t) = 0$ if and only if $x_i(t) = x_l(t)$ for all $i \in S_f^{\mathcal{N}}$, $l \in \mathcal{L}$.

Theorem 1. *Let $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ be a nonempty, nontrivial, simple digraph with S_f nonempty. Let $F \in \mathbb{Z}_+$, $u_M \in \mathbb{R}_+$. Let $V(t)$ be defined as in* (5)*. Define the error function $e(t) = \max_{i \in S_f^{\mathcal{N}}} |x_i(t) - x_l(t)|$, $l \in \mathcal{L}^{\mathcal{N}}$. Suppose that $\mathcal{A}$ is an F-local set, $\mathcal{D}$ is strongly $(2F+1)$-robust w.r.t. the set $\mathcal{L}$, and all normally behaving agents apply the MS-RPA algorithm with parameter F. Suppose further that the input of follower agents is constrained by $|u_i(t)| \le u_M$ $\forall t \ge 0$, and the reference signal satisfies $|r_L(\tau+1) - r_L(\tau)| \le u_M - \epsilon$ $\forall \tau \ge 0$, for some $\epsilon > 0$, $\epsilon \in \mathbb{R}$ and with $\tau \in \mathbb{Z}_+$. If $\eta > |S_f|$, then $e(t)$ is nonincreasing for $t \in [, \infty)$ and there exists a $T \in \mathbb{Z}_+$ such that $e(t) = 0$ for all $t \ge T$, $\forall i \in S_f^{\mathcal{N}}$, $\forall l \in \mathcal{L}^{\mathcal{N}}$. Furthermore, the value of T can be found as*

$$T = \left\lceil \frac{V(0)}{\epsilon} \right\rceil + 1$$

Proof. Since all behaving leaders states satisfy $|x_{l_1}(t) - x_{l_2}(t)| = 0$ $\forall t \ge \forall l_1, l_2 \in \mathcal{L}^{\mathcal{N}}$ as per the MS-RPA algorithm and (1), we will simply write $x_l(t)$ for brevity. Notice that Lemma 1 and $\eta > |S_f|$ imply that all normal agents receive and set $r_i(t) = r_L(\tau+1)$ at or before the time $t = (\tau+1)\eta - 1$. In addition, the input bound u_M guarantees that $x_i((\tau+1)\eta)$ lies within the (one dimensional) ball $B(x_i(\tau\eta)$ for all normal agents $i \in S_f^{\mathcal{N}}$ and for all $\tau \in \mathbb{Z}_+$.

The main idea of the proof is to show that $V(t)$ decreases by at least ϵ at every state update time since every change in leader's states is bounded by $u_M - \epsilon$ while every change in followers' states is only bounded by u_M. Once $V(\tau\eta) \le \epsilon$ for $t = \tau\eta$ for some $\tau \in \mathbb{Z}_+$, we will show that $V((\tau+q)\eta) = 0$ for all future state update times with integer $q \ge 1$. We begin by demonstrating that all of the following three statements hold true $\forall \tau \in \mathbb{Z}_+$:

1) $r_L(\tau+1) \in B(m(\tau\eta), u_M) \cap B(M(\tau\eta), u_M) \implies V((\tau+1)\eta) = 0$
2) $V(\tau\eta) \le \epsilon \implies V((\tau+1)\eta) = 0$
3) $V(\tau\eta) > \epsilon$ and $r_L(\tau+1) \notin B(m(\tau\eta), u_M) \cap B(M(\tau\eta), u_M) \implies V((\tau+1)\eta) - V(\tau\eta) \le -\epsilon$

In words, these conditions can be interpreted as follows:

1) If the reference value $r_L(\tau+1)$ is within distance u_M from all agents with maximum and all agents with minimum state values, all normal agents' states at $t = \tau\eta$ will be within distance u_M from $r_L(\tau+1)$. They are therefore able to change their states to $r_L(\tau+1)$ at the next state update time $t = \tau\eta$, which implies $V((\tau+1)\eta) = 0$.
2) If $V(\tau\eta) \le \epsilon$, then all normal agents will be able to change their state values to $r_L(\tau+1)$ at the next state update time, implying $V((\tau+1)\eta) = 0$.
3) If neither of the hypotheses for 1) nor 2) are satisfied, then $V((\tau+1)\eta)$ is guaranteed to be at least ϵ less than $V(\tau\eta)$.

Proof of 1): Suppose $r_L(\tau+1) \in B(m(\tau\eta), u_M) \cap B(M(\tau\eta), u_M)$. In words, this implies that $r_L(\tau+1)$ is within a distance u_M from all normal agents with either minimum or maximum state values. We show that this implies $r_L(\tau+1)$ is within a distance u_M of *all* normal agents' states. First, $r_L(\tau+1) \in B(m(\tau\eta), u_M) \cap B(M(\tau\eta), u_M)$ implies that $|r_L(\tau+1) - m(\tau\eta)| \le u_M$ and $|r_L(\tau+1) - M(\tau\eta)| \le u_M$, which implies $M(\tau\eta) - u_M \le r_L(\tau+1) \le m(\tau\eta) + u_M$.

We have

$$r_L(\tau{+}1) \geq M(\tau\eta){-}u_M \geq x_i(\tau\eta){-}u_M,$$
$$r_L(\tau{+}1) \leq m(\tau\eta){+}u_M \leq x_i(\tau\eta){+}u_M,$$

for all $i \in S_f^{\mathcal{N}}$, which implies $|x_i(\tau\eta) - r_L(\tau+1)| \leq u_M$ for all $i \in S_f^{\mathcal{N}}$. Next, since every normal agent i is within a distance u_M of $r_L(\tau+1)$ at $t = \tau\eta$, each i can apply a control action such that $x_i(\tau\eta) = r_L(\tau+1)$ at $t = \tau\eta$. To see this, by Lemma 1 each agent $i \in S_f^{\mathcal{N}}$ receives and accepts the value $r_L(\tau+1)$ at or before time step $t = (\tau+1)\eta - 1$. Each agent i therefore selects $u_i((\tau+1)\eta - 1) = \frac{(r_L(\tau+1)-x_i((\tau+1)\eta-1))u_M}{\max(u_M,|r_L(\tau+1)-x_i((\tau+1)\eta-1)|)} = r_L(\tau+1) - x_i((\tau+1)\eta-1) = r_L(\tau+1) - x_i(\tau\eta)$, implying $x_i((\tau+1)\eta) = x_i(\tau\eta) + (r_L(\tau+1) - x_i(\tau\eta)) = r_L(\tau+1)$. Using (1) we therefore have $|x_i((\tau+1)\eta) - x_l((\tau+1)\eta)| = 0$ for all $i \in S_f^{\mathcal{N}}$ at time $(\tau+1)\eta$, implying $V((\tau+1)\eta) = 0$.

Proof of 2): We show that $V(\tau\eta) \leq \epsilon$ implies that condition 1) holds, and therefore by prior arguments we have $V((\tau+1)\eta) = 0$. First, $V(\tau\eta) \leq \epsilon$ implies $M(\tau\eta) - m(\tau\eta) \leq \epsilon$. By the definitions of $m(t)$ and $M(t)$ we have $m(\tau\eta) \leq x_l(\tau\eta) = r_L(\tau) \leq M(\tau\eta)$. Using the bound $|r_L(\tau+1) - r_L(\tau)| \leq u_M - \epsilon$ from the Theorem statement yields $m(\tau\eta){-}u_M{+}\epsilon \leq r_L(\tau) \leq M(\tau\eta){+}u_M{-}\epsilon$. However, since $V(\tau\eta) \leq \epsilon$ we therefore have $M(\tau\eta) - m(\tau\eta) \leq \epsilon$, which after rearranging and subtracting u_M from both sides yields $M(\tau\eta) - u_M \leq m(\tau\eta) - u_M + \epsilon$. Similarly, we have $m(\tau\eta) \geq M(\tau\eta) - \epsilon$ implying $m(\tau\eta) + u_M \geq M(\tau\eta) + u_M - \epsilon$. Combining these relationships with prior equations yields $M(\tau\eta) - u_M \leq m(\tau\eta) - u_M + \epsilon \leq r_L(\tau+1) \leq M(\tau\eta){+}u_M{-}\epsilon \leq m(\tau\eta){+}u_M$. This implies that $r_L(\tau{+}1)$ must therefore be within a distance u_M of both $m(\tau\eta)$ and $M(\tau\eta)$; i.e. $r_L(\tau+1) \in B(m(\tau\eta), u_M) \cap B(M(\tau\eta), u_M)$. This precisely matches the conditions of condition 1) implying by the arguments of 1) that $V((\tau+1)\eta) = 0$.

x*Proof of 3):* $V(\tau\eta) > \epsilon$ and $r_L(\tau{+}1) \notin B(m(\tau\eta), u_M) \cap B(M(\tau\eta), u_M)$ imply that either $r_L(\tau{+}1) \notin B(m(\tau\eta), u_M)$ or $r_L(\tau+1) \notin B(M(\tau\eta), u_M)$. Without loss of generality, consider the case where $r_L(\tau+1) \notin B(M(\tau\eta), u_M)$. Since $m(\tau\eta) \leq x_l(\tau\eta) = r_L(\tau)$ by definition, and since $|r_L(\tau) - r_L(\tau{+}1)| \leq u_M{-}\epsilon$, we therefore have $m(\tau\eta){-}r_L(\tau{+}1) \leq u_M{-}\epsilon$, which after rearranging yields $r_L(\tau{+}1) \geq m(\tau\eta) - (u_M - \epsilon)$. Next, we show that $m((\tau+1)\eta) - m(\tau\eta) \leq (u_M{-}\epsilon)$. Observe that if $r_L(\tau{+}1) \geq m(\tau\eta)$ we have $m((\tau+1)\eta) \geq m(\tau\eta)$, since any normal agent i with $x_i(\tau\eta) \leq r_L(\tau+1)$ will apply a control action towards $r_L(\tau+1)$. Therefore, $m(\tau\eta) - m((\tau+1)\eta) \leq 0$. On the other hand, if $r_L(\tau+1) < m(\tau\eta)$ we have $m((\tau+1)\eta) \geq m(\tau\eta) - (u_M - \epsilon)$. This follows since $|r_L(\tau+1) - r_L(\tau-1)| \leq u_M - \epsilon$ and therefore $m((\tau+1)\eta)$ is lower bounded by the case when $r_L(\tau+1) < m(\tau\eta)$. These arguments imply $m((\tau+1)\eta) - m(\tau\eta) \leq (u_M - \epsilon)$. By Lemma 1 all agents $i \in S_f^{\mathcal{N}}$ will receive and accept the value $r_L(\tau+1)$ at some time $\tau\eta + 1 \leq t \leq (\tau+1)\eta - 1$. This implies that any agent i such that $x_i(\tau\eta) = M(\tau\eta)$ will select the control action $u_i((\tau{+}1)\eta{-}1) = \frac{(r_L(\tau+1)-x_i((\tau+1)\eta-1))u_M}{\max(u_M,|r_L(\tau+1)-x_i((\tau+1)\eta-1)|)} = u_M \mathrm{sgn}(r_L(\tau+1) - x_i((\tau+1)\eta-1)) = -u_M$. This follows since $r_L(\tau)$ is less than or equal to $M(\tau\eta)$, $r_L(\tau+1) \notin B(M(\tau\eta), u_M)$, and $|r_L(\tau{+}1){-}r_L(\tau)| \leq u_M{-}\epsilon$. All i with state values equal to $M(\tau\eta)$ will therefore apply maximum control effort in the direction of $r_L(\tau+1)$. This implies $M((\tau+1)\eta) - M(\tau\eta) = -u_M$.

Using these arguments, the difference between $V((\tau{+}1)\eta)$ and $V(\tau\eta)$ is

$$\begin{aligned} V((\tau{+}1)\eta){-}V(\tau\eta) &= M((\tau{+}1)\eta){-}M(\tau\eta) \\ &\quad +m(\tau\eta){-}m((\tau{+}1)\eta), \\ &\leq -u_M{+}u_M{-}\epsilon \ \leq \ -\epsilon. \end{aligned}$$

Similar arguments hold if $x_l((\tau+1)\eta) \notin B(m(\tau\eta), u_M)$, which yields 3).

From 1), 2) and 3) we can infer that if $V(\tau\eta) > \epsilon$ and $r_L(\tau+1)\eta) \notin B(m(\tau\eta), u_M) \cap B(M(\tau\eta), u_M)$, then $V((\tau+1)\eta) - V(\tau\eta) \leq -\epsilon$. This implies that the value of $V(\cdot)$ decreases until $V((\tau+k')\eta) \leq \epsilon$ or $r_L(\tau+1+k') \in B(m((\tau+k')\eta), u_M) \cap B(M((\tau+k')\eta), u_M)$ for some finite $k' \geq 1$, either of which conditions imply that $V((\tau+1+k')\eta) = 0$ by conditions 1) and 2). In addition, note that for any time $t = \tau^*\eta$ such that $V(\tau^*\eta) = 0$ we have $V(\tau^*\eta) \leq \epsilon$, implying by 2) that $V(t) = 0$ for all $t \geq \tau^*$. Since k' is finite, we can conclude that there exists a $T \in \mathbb{Z}_+$ such that $|x_i(t){-}x_l(t)| = 0\ \forall i \in S_f^{\mathcal{N}}, \forall t \geq T$, which implies $e(t) = 0\ \forall t \geq T$. In addition, statements 1), 2), and 3) imply that $(V((\tau+1)\eta) - V(\tau\eta) \leq 0\ \forall \tau \in \mathbb{Z}_+$, which along with (1), (2), and (3) implies that $e(t)$ is nonincreasing $\forall t \geq 0$.

An exact value of T can be found by using statement 3) which implies that the slowest rate of decrease of $V(\cdot)$ is $V((\tau{+}1)\eta){-}V(\tau\eta) \leq -\epsilon$. Letting $k' = (\lceil V(0)/\epsilon \rceil {+} 1)$, it is straighforward to show that after $k'\eta$ time steps, $V(k'\eta) = 0$. Therefore T can be taken to be $T = k' = (\lceil V(0)/\epsilon \rceil + 1)$. $\square$

V. Simulations

This section presents simulations of the MS-RPA algorithm conducted on an undirected k-circulant digraph [12], [18] denoted $\mathcal{D}_1$ with $n = 14$ agents and parameter $k = 5$. The set of leaders is $\mathcal{L} = \{1, \ldots, 5\}$, and $S_f = \{6, \ldots, 14\}$. Our results in [12] demonstrate conditions under which k-circulant undirected and directed graphs are strongly r-robust with respect to $\mathcal{L}$. Briefly, if a k-circulant undirected or directed graph $\mathcal{D}$ contains a set of consecutive agents by index P_L such that $|P_L| \leq k$ and $|P_L \cap \mathcal{L}| \geq r$, $\mathcal{D}$ is strongly r-robust w.r.t. $\mathcal{L}$. Letting $P_L = \{1, \ldots, 5\}$, we clearly have $|P_L| \leq k = 5$ and $|P_L \cap \mathcal{L}| \geq 5$, implying $\mathcal{D}_1$ is 5-robust. The parameter F is set to $F = 2$, implying that $\mathcal{D}_1$ is strongly $(2F+1)$-robust w.r.t. $\mathcal{L}$. In the first simulation, the follower agents' inputs are bounded by $u_M = 10.1$. All agents begin with randomly initial states on the interval $[-25, 25]$. The initial time $= 0$, the reference function is $r_L(\tau+1) = 10\sin(\tau/\pi)$, and the communication rate is defined by $\eta = 10$. Note that $\eta > |S_f| = 9$. One leader agent and one follower agent are misbehaving, with $A = \{4, 11\}$, by updating their state according to an arbitrary function *and* sending random values to their out-neighbors at each time step; the misbehavior is *malicious* [2] because each

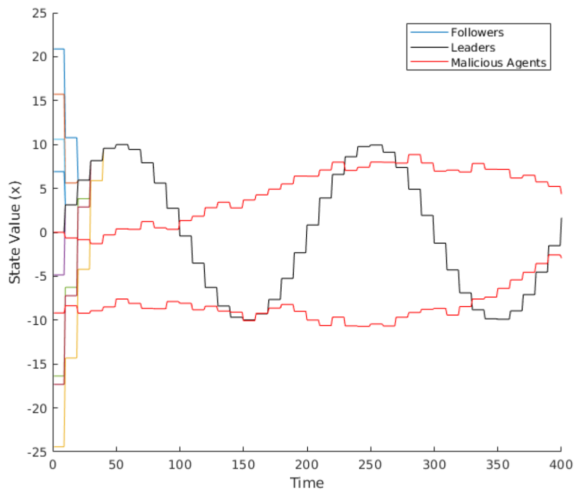

Fig. 1. The MS-RPA algorithm with input bounds. When the inputs are bounded, exact convergence to the time-varying leaders' states is still guaranteed according to Theorem 1. In this simulation, $\eta = 10$.

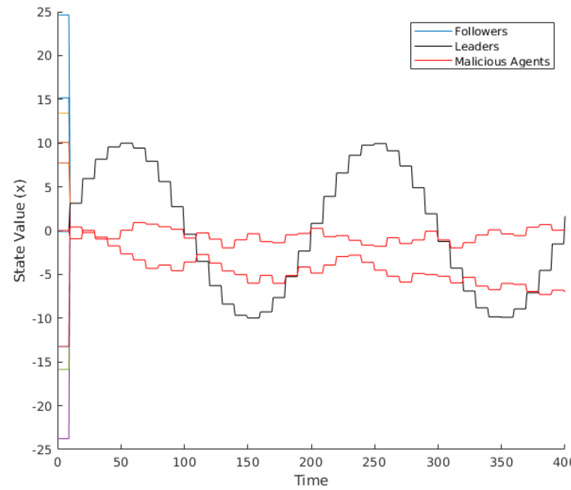

Fig. 2. The MS-RPA algorithm with no input bounds. Leader agents' states are black lines, misbehaving agent's states are red lines, and follower agents' states are the multi-colored lines. When the inputs are unbounded, exact convergence to the leaders' trajectory occurs for all $t \geq \eta$, where $\eta = 10$.

misbehaving agent sends the same misinformation to each of its out-neighbors. We emphasize that *the normally behaving agents have no knowledge about which agents are misbehaving*. All other normally behaving agents apply the MS-RPA algorithm (Algorithm 1). Note that $|r_L(\tau+1)-r_L(\tau)| < u_M$ for all $\tau \in Z_+$, and therefore convergence is guaranteed after a finite number of time steps T as per Theorem 1. This is shown in Figure 1, where the followers converge exactly to the leaders after a finite number of time steps. In the second simulation, all parameters are the same as in the first simulation except 1) agents' inputs are *unbounded* and 2) the set of misbehaving agents is $\mathcal{A} = \{1, 5\}$ (two leader agents). Again, normally behaving agents have no knowledge about which agents are misbehaving. For this particular case with unbounded inputs, the normally behaving follower agents achieve consensus to the leader agents at time step η, where $t_0 = 0$, and successfully track the leaders' time-varying state *exactly* for all remaining time.

VI. Conclusion

In this paper we presented conditions for agents with discrete-time dynamics to resiliently track a time-varying reference signal propagated by a set of leader agents in the presence of adversarial agents. Future work includes extending these results to time-varying graphs with asynchronous communication.

References

[1] R. M. Kieckhafer and M. H. Azadmanesh, "Reaching approximate agreement with mixed-mode faults," *IEEE Transactions on Parallel and Distributed Systems*, vol. 5, no. 1, pp. 53–63, 1994.

[2] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 4, pp. 766–781, 2013.

[3] S. M. Dibaji and H. Ishii, "Resilient consensus of second-order agent networks: Asynchronous update rules with delays," *Automatica*, vol. 81, pp. 123–132, 2017.

[4] D. Saldana, A. Prorok, S. Sundaram, M. F. Campos, and V. Kumar, "Resilient consensus for time-varying networks of dynamic agents," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 252–258.

[5] S. M. Dibaji, H. Ishii, and R. Tempo, "Resilient randomized quantized consensus," *IEEE Transactions on Automatic Control*, vol. 63, no. 8, pp. 2508–2522, 2018.

[6] D. V. Dimarogonas, P. Tsiotras, and K. J. Kyriakopoulos, "Leader–follower cooperative attitude control of multiple rigid bodies," *Systems & Control Letters*, vol. 58, no. 6, pp. 429–435, 2009.

[7] W. Ren, "Multi-vehicle consensus with a time-varying reference state," *Systems and Control Letters*, vol. 56, no. 7-8, pp. 474–483, 2007.

[8] ——, "Consensus tracking under directed interaction topologies: Algorithms and experiments," in *American Control Conference, 2008*. IEEE, 2008, pp. 742–747.

[9] H. J. LeBlanc and F. Hassan, "Resilient distributed parameter estimation in heterogeneous time-varying networks," in *Proceedings of the 3rd international conference on High confidence networked systems*. ACM, 2014, pp. 19–28.

[10] A. Mitra and S. Sundaram, "Secure distributed observers for a class of linear time invariant systems in the presence of Byzantine adversaries," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*. IEEE, 2016, pp. 2709–2714.

[11] ——, "Resilient distributed state estimation for LTI systems," *arXiv preprint arXiv:1802.09651*, 2018.

[12] J. Usevitch and D. Panagou, "Resilient leader-follower consensus to arbitrary reference values," in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 1292–1298.

[13] A. Mitra and S. Sundaram, "Secure distributed state estimation of an lti system over time-varying networks and analog erasure channels," in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 6578–6583.

[14] C.-Y. Koo, V. Bhandari, J. Katz, and N. H. Vaidya, "Reliable broadcast in radio networks: The bounded collision case," in *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*. ACM, 2006, pp. 258–264.

[15] C. Litsas, A. Pagourtzis, and D. Sakavalas, "A graph parameter that matches the resilience of the certified propagation algorithm," in *International Conference on Ad-Hoc Networks and Wireless*. Springer, 2013, pp. 269–280.

[16] S. Bonomi, G. Farina, and S. Tixeuil, "Reliable broadcast in dynamic networks with locally bounded byzantine failures," in *International Symposium on Stabilizing, Safety, and Security of Distributed Systems*. Springer, 2018, pp. 170–185.

[17] Haotian Zhang and S. Sundaram, "Robustness of information diffusion algorithms to locally bounded adversaries," in *2012 American Control Conference (ACC)*, Montreal, Canada, 2012, pp. 5855–5861.

[18] J. Usevitch and D. Panagou, "r-robustness and (r, s)-robustness of circulant graphs," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 4416–4421.