Resilient Asymptotic Consensus in Asynchronous Robust Networks

Heath J. LeBlanc and Xenofon Koutsoukos

***Abstract*— In this paper, we study the problem of reaching consensus asymptotically in the presence of adversary nodes whenever the network is asynchronous under a local broadcast model of communication. The type of adversary considered is omniscient and may collude with other adversaries to achieve the goal of disrupting consensus among the normal nodes. The main limitation on the behavior of the adversary nodes is that whenever the adversary nodes communicate with neighbors, they must broadcast their messages so that all neighbors receive the same information. The asynchronous consensus algorithm studied here uses local strategies to ensure resilience against the adversary nodes. The class of topologies studied are those that are *robust*. Network robustness formalizes a notion of redundancy of direct information exchange between subsets of nodes in the network, and is an important property for analyzing the behavior of resilient distributed algorithms that use only local information.**

I. INTRODUCTION

Consensus problems have a rich history in distributed computing [1] and communication [2]. More recently, consensus has become an active area of control research [3], [4]. This is because reaching agreement is a fundamental task in distributed and multi-agent systems, and arises in diverse applications such as agent flocking [5], [6], synchronized path following [7], distributed estimation [8], and load balancing for parallel processors [9]. A major concern in large-scale distributed systems is whether the group objectives can be achieved in the presence of uncertainties such as communication delays, data loss, or node failures. While researchers in control have studied consensus algorithms that have been shown to be robust to communication delays [10], data loss [11], and quantization [12], the robustness of such algorithms to node failures has been shown to be lacking [13].

Of course, consensus algorithms that are robust to node failures have been studied in distributed computing [1], [14], communication networks [15], and mobile robotics [16]–[18]. In these works, the faulty nodes may be characterized by *fault models* and *scope of fault* assumptions. Two common fault models are the *crash fault* [16], [17] and the *Byzantine fault* [16], [19]–[21]. Whenever a node suffers a crash fault, it simply stops somewhere in its execution (this is also referred to as a stopping failure [1]). A Byzantine faulty node, on the other hand, may behave in an arbitrary manner. Therefore, worst case executions must be considered. For the scope of faults, it is typically assumed that at most F out of n nodes fail. We refer to this as the F-total model. On the other hand, a local bound on the number of faulty nodes has been considered in fault-tolerant broadcasting [22]–[24] and consensus [24], [25], where it is assumed that at most F of any normal node's neighbors fail. We refer to this as the F-local model.

Another important concern with respect to networked systems is the issue of malicious attacks and security breaches [26]. Attacks on the network may include jamming [27], denial-of-service [28], false data injection [29], replay [30], or deception [31]. In jamming and denial-of-service attacks, the attacker reduces (or entirely eliminates) the availability of data from the communication network. False data injection and deception attacks affect the integrity of the data. Likewise, replay attacks inject incorrect and outdated information by repeating previously transmitted information.

In addition to direct attacks on the communication network, another important type of attack – especially relevant to the cooperative control of multi-agent networks – occurs when a subset of nodes are compromised and behave as adversaries. In the context of consensus, the adversary nodes participate with the goal of disrupting consensus or leading the consensus process to an invalid value. By analogy with node failures, the adversary nodes may be characterized by threat models and scope of threat assumptions. For such a scenario, the Byzantine model is a suitable threat model. However, depending on the communication realization, the full generality of the Byzantine model may not be necessary. In general, Byzantine nodes may simultaneously send different information to different neighbors in the network [1]. However, if the nodes broadcast their information to neighbors, such as in wireless broadcast, then duplicity (of this type) is not possible. We refer to the Byzantine node under the local broadcast model as a *malicious node* [24], [25], [32]–[34].

One approach to overcoming adversary nodes is to try to identify the adversary nodes so that their influence can be removed [32], [35]. While identifying misbehaving nodes is clearly an interesting and important problem, detection and identification techniques require each node to have information of the network topology beyond its local neighborhood [32], [35]. This requirement of *nonlocal information* may not be suitable for large-scale networks. Furthermore, the detection algorithms are computationally expensive.

A second approach is to design computationally efficient

This work is supported in part by the National Science Foundation (CNS-1035655, CCF-0820088), the U.S. Army Research Office (ARO W911NF-10-1-0005), and Lockheed Martin.

Heath J. LeBlanc is with the Department of Electrical & Computer Engineering and Computer Science, Ohio Northern University, Ada, OH, USA h-leblanc@onu.edu

Xenofon Koutsoukos is with the Department of Electrical Engineering and Computer Science, Vanderbilt University, Nashville, TN, USA xenofon.koutsoukos@vanderbilt.edu

algorithms that filter the information received from neighbors to ensure resilience [24], [25], [33], [36]. A class of such algorithms has been introduced for the study of Byzantine approximate agreement [20], and has been extended to a family of algorithms, called the *Mean-Subsequence-Reduced (MSR)* algorithms [37]. MSR algorithms are iterative consensus algorithms that are designed under the assumption that at most F nodes fail (or are compromised). The main idea is for each normal node to eliminate from consideration the largest and smallest F values (i.e., the extreme values) in its neighborhood. The values from a subset of the remaining nodes are then averaged to determine the value for the next round. MSR algorithms have been used extensively to achieve fault tolerant and resilient consensus (e.g., in clock synchronization [38] and robot gathering [16]–[18]). However, the network topological condition for characterizing convergence has long been an open problem.

Recently, it has been shown that traditional graph theoretic metrics (such as connectivity) are inadequate for characterizing the conditions under which MSR algorithms achieve resilient consensus [24], [33]. Because of the removal of extreme values in MSR algorithms, a property that encapsulates the notion of sufficient *local* redundancy of incoming information is needed. This idea is captured by *network robustness* [24], [25] and a similar property studied in [21], [39], [40]. Equipped with these properties, the necessary and sufficient conditions for convergence of a class of MSR algorithms have been given for the Byzantine model under the assumption that at most F nodes are compromised (F-total model) in synchronous [21], [39], [40] and asynchronous networks [40]. For the local broadcast version of the Byzantine model (referred to as the malicious adversary), the tight condition on the network topology has been given only for the case of synchronous networks under the F-total model [25].

This paper formulates the resilient asymptotic consensus (RAC) problem in an asynchronous framework, with a local broadcast model of communication. We characterize, for the first time, the necessary and sufficient condition on the network topology for the existence of an algorithm that achieves RAC in an asynchronous network under the local broadcast model of communication, and in the presence of up to F malicious nodes (F-total model). To show sufficiency, we adapt the Weighted MSR (W-MSR) algorithm of [24], [34] to an asynchronous setting and prove that the algorithm succeeds under the necessary condition. Finally, we provide a sufficient condition for the existence of an algorithm that achieves RAC in the asynchronous network model under the F-local model.

The rest of the paper is organized as follows. Section II defines the networked system model, the threat model, the scope of threat assumptions, and the problem statement. Section III describes how the W-MSR algorithm is adapted to an asynchronous setting. Section IV contains the main results. Section V provides context for how these results relate to the literature. Finally, Section VI summarizes the paper.

II. SYSTEM MODEL AND PROBLEM FORMULATION

Consider a time-invariant network topology modeled by the finite simple directed graph, or just *digraph*, $\mathcal{D} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, ..., n\}$ is the *node set* and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the *directed edge set*. Without loss of generality, the node set is partitioned into a set of N *normal nodes* $\mathcal{N} = \{1, 2, \ldots, N\}$ and a set of M *adversary nodes* $\mathcal{A} = \{N+1, N+2, \ldots, n\}$, with $M = n - N$. The adversary nodes are assumed to be unknown a priori to the normal nodes. Each directed edge $(i, j) \in \mathcal{E}$ models the capability of node i to influence node j. The set of *in-neighbors*, or just *neighbors*, of node i is defined as $\mathcal{N}_i = \{j \in \mathcal{V} \colon (j, i) \in \mathcal{E}\}$ and the (in-)degree of i is denoted $d_i = |\mathcal{N}_i|$. Likewise, the set of *out-neighbors* of node i is defined as $\mathcal{N}_i^{\text{out}} = \{j \in \mathcal{V} \colon (i, j) \in \mathcal{E}\}$. Because each node has access to its own state, we also consider the *inclusive neighbors* of node i, denoted $\mathcal{J}_i = \mathcal{N}_i \cup \{i\}$.

The digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ models the communication of the networked system. The nodes communicate under a *local broadcast model*, in which the out-neighbors of each node i in $\mathcal{D}$ are precisely those nodes capable of receiving messages from i. The communication is assumed to be *reliable* (meaning all transmitted messages are eventually delivered successfully), but the messages may incur different (arbitrary) delays in transmission to different nodes and messages may be received out of order. It is further assumed that the sender of each message is identifiable by the receiver.

For the class of algorithms studied in this work, the execution of each node proceeds in a sequence of rounds $r \in \mathbb{Z}_{\geq 0}$ that consists of *transmit*, *receive*, and *update stages*. A node (or subset of nodes) that is in the process of executing its algorithm (i.e., it is in either the transmit, receive, or update stage of some round $r \in \mathbb{Z}_{\geq 0}$) is said to be *active*. The networked system is *asynchronous*, meaning the nodes do not necessarily execute rounds at the same rate and the nodes are *not* synchronized. Therefore, at any point in real time $t \in \mathbb{R}_{\geq 0}$, the nodes may be in different stages of different rounds. The reference time $t = 0$ is defined as the point in time at which the first subset of normal nodes becomes active. Because there is no synchrony among the nodes, it is possible that at some time $t > 0$, some of the nodes may have not yet become active. To handle this situation, we assume there exists a *dormant stage* of round $r = 0$, in which nodes may accumulate messages from in-neighbors but do not act upon them until becoming active. Note that we place no limitations on the amount of storage available at each node.

In order to keep track of the messages corresponding to a given round, each message is tagged by the round $r \in \mathbb{Z}_{\geq 0}$ in which it is sent. Every node (including adversary nodes) sends at most one message to its out-neighbors in each round, and each normal node sends exactly one message per round. Since the normal nodes follow this protocol, any adversary node that sends multiple messages tagged by a single round would be easily detected as an adversary. However, an adversary may skip rounds without detection (in finite time) because the messages may be received out

of order and with arbitrary delay. Moreover, an adversary may decide at some point in time to stop sending messages altogether.

A. Update Model

Suppose that each node $i \in \mathcal{V}$ maintains a scalar value $x_i(r) \in \mathbb{R}$ called the *state*[1] of node i in round r. In particular, each node begins with the private value $x_i(0)$ (which could represent a measurement, optimization variable, etc.). At the beginning of each round $r \in \mathbb{Z}_{\geq 0}$ (once the node becomes active), each normal node i broadcasts its value to its out-neighbors in the network (transmit stage). The value sent by node j in round $r \in \mathbb{Z}_{\geq 0}$ is denoted $x_j(r)$. Once each normal node i transmits its value, it then waits[2] to receive $d_i^* < d_i$ values (receive stage) from its in-neighbors (d_i^* depends on the adversary model, scope of threat assumptions, and size of $\mathcal{N}_i$). Observe that d_i^* should be small enough so as to avoid deadlock. Once node $i \in \mathcal{N}$ collects d_i^* values from in-neighbors, it updates its value for round $r+1$ according to the prescribed rule

$$x_i(r+1) = f_i\left(x_i(r), \{x_j(r)\}\right), \;\; i \in \mathcal{N}, j \in \mathcal{N}_i^*(r),$$

where $\mathcal{N}_i^*(r)$ is the set of nodes corresponding to the d_i^* values received from node i's in-neighbors in round r. The update rule $f_i(\cdot)$ can be an arbitrary function, and may be different for each node, depending on its role in the network. These functions are designed *a priori* so that the normal nodes reach consensus. However, some of the nodes may not follow the prescribed strategy if they are compromised by an adversary. Such misbehaving nodes threaten the group objective, and it is important to design the $f_i(\cdot)$'s in such a way that the influence of such nodes can be eliminated or reduced without prior knowledge about their identities.

B. Threat Model

The type of adversary considered is referred to as a *malicious* node. Malicious nodes are omniscient adversaries. In particular, they know all other values and the full network topology; they are aware of the update rules $f_i(\cdot)$, $\forall i \in \mathcal{N}$; they are aware of which other nodes are adversaries; and they know the plans of the other adversaries.[3] Although malicious nodes have complete knowledge, their ability to affect other nodes is limited. Specifically, a malicious node $k \in \mathcal{A}$ may choose whether or not to broadcast its value to its out-neighbors in any round $r \in \mathbb{Z}_{\geq 0}$, but if it does, it must send *at most* one value $x_k(r)$ tagged with round r (otherwise, the normal nodes receiving more than one value from node k in round r would know that k is an adversary). Additionally, in any finite time interval, a malicious node must send a finite number of messages. Finally, a malicious node may update its value in an arbitrary fashion. Since it is omniscient, one must assume that the update is one that causes the most disruption to the normal nodes. Note that malicious nodes are Byzantine nodes restricted to the local broadcast model. Byzantine nodes differ in that they are capable of sending different messages to different out-neighbors, which is possible under a point-to-point communication model [40].

[1]Throughout this paper, we use interchangeably the terms "value" and "state" of a node.

[2]If node i is one of the last normal nodes to become active, it is possible that it may have already received at least d_i^* messages from in-neighbors for multiple rounds. In such a case, there is no need for node i to wait.

[3]One may take the viewpoint that a centralized omniscient adversary informs and directs the behavior of the malicious nodes.

C. Scope of Threats

Having defined the threat model, it is necessary to define the *number* of adversary nodes. While there are various stochastic models that could be used to formalize the scope of threats, we use a deterministic approach and consider an upper bound on the number of compromised nodes in the network (F-total), or in the neighborhood of a normal node (F-local).

Definition 1 (F-total model): A set $\mathcal{S} \subset \mathcal{V}$ is **F-total** (with $F \in \mathbb{Z}_{\geq 0}$) if it contains at most F nodes in the network, i.e., $|\mathcal{S}| \leq F$. The **F-total model** refers to the case when the set of adversaries $\mathcal{A}$ is an F-total set. □

Definition 2 (F-local model): A set $\mathcal{S} \subset \mathcal{V}$ is **F-local** (with $F \in \mathbb{Z}_{\geq 0}$) if it contains at most F nodes in the neighborhood of the other nodes, i.e., $|\mathcal{N}_i \bigcap \mathcal{S}| \leq F$, $\forall i \in \mathcal{V} \setminus \mathcal{S}$. The **$F$-local model** refers to the case when the set of adversaries $\mathcal{A}$ is an F-local set. □

Note that whenever the set of M adversary nodes $\mathcal{A}$ is an F-total set, we know $M \leq F$. On the other hand if $\mathcal{A}$ is an F-local set, it is possible that $M > F$. Indeed, there is no upper bound on M for F-local set $\mathcal{A}$ since it is feasible that many adversaries may not be neighbors with any normal node. The F-total fault model has been studied in distributed computing [1], [19], [21] and mobile robotics [16]–[18] for both stopping (or crash) failures and Byzantine failures. The F-local fault model has been studied in the context of Byzantine-resilient broadcasting [22]–[24] and consensus in synchronous networks [24], [25], [34].

D. Resilient Asymptotic Consensus

Given the adversary model and scope of threats, we formally define the *resilient asymptotic consensus problem*. Let $M_{\mathcal{N}}(r)$ and $m_{\mathcal{N}}(r)$ denote the *maximum* and *minimum* values of the normal nodes in round r, respectively.

Definition 3 (Resilient Asymptotic Consensus): The normal nodes are said to achieve **resilient asymptotic consensus (RAC)** in the presence of adversary nodes (given a particular adversary model) if

(i) $m_{\mathcal{N}}(r+1) \geq m_{\mathcal{N}}(r)$ and $M_{\mathcal{N}}(r+1) \leq M_{\mathcal{N}}(r)$, for every round $r \in \mathbb{Z}_{\geq 0}$, and

(ii) $\lim_{r\to\infty} M_{\mathcal{N}}(r) - m_{\mathcal{N}}(r) = 0$,

for any choice of initial values $x_i(0)$ for $i \in \mathcal{V}$. □

The RAC problem consists of two conditions. The first (i) is a *validity* or *safety* condition. If it is satisfied, then the states of the normal nodes always remain inside the initial interval $[m_{\mathcal{N}}(0), M_{\mathcal{N}}(0)]$ (safety), and any value selected (i.e., through termination) is guaranteed to lie in this interval (validity). The second (ii) is a *convergence* condition on agreement. Observe that any asynchronous algorithm that achieves resilient asymptotic consensus in the presence of

adversary nodes under either the F-total or F-local model must wait for no more than $d_i^* = d_i - F$ values in its receive stage in order to avoid deadlock.

III. ASYNCHRONOUS W-MSR ALGORITHM

In this section, we modify the Weighted MSR (W-MSR) algorithm with parameter F that has been studied in synchronous networks [24], [25], [34]. The modifications made to accommodate asynchrony are analogous to the modifications made for MSR algorithms [20], [40], and consist of the following two changes: (i) the messages are tagged with the corresponding round index, and (ii) each normal node j waits to receive only $d_j^* = d_j - F$ messages from in-neighbors for a given round r before updating its value.

Asynchronous W-MSR with parameter F:
In each round $r \in \mathbb{Z}_{\geq 0}$, once active, normal node i performs the following steps:

1) *Transmit step*: Send the current value $x_i(r)$ to outgoing neighbors, along with round index r.
2) *Receive step*: Wait to receive exactly $d_i^* = d_i - F$ messages from different nodes tagged by round index r, and break ties arbitrarily. Sort the $d_i - F$ values in ascending order. If there are less than F values strictly larger than its own value, $x_i(r)$, then normal node i removes all values that are strictly larger than its own. Otherwise, it removes precisely the largest F values in the sorted list (breaking ties arbitrarily). Likewise, if there are less than F values strictly smaller than its own value, then node i removes all values that are strictly smaller than its own. Otherwise, it removes precisely the smallest F values.
3) *Update step*: Let $\mathcal{R}_i(r)$ denote the set of nodes whose values are removed or disregarded by normal node i in step 2 of round r. Each normal node i applies the update

$$x_i(r+1) = \sum_{j \in \mathcal{J}_i \setminus \mathcal{R}_i(r)} w_{ij}(r) x_j(r), \qquad (1)$$

where the weights $w_{ij}(r)$ satisfy the following conditions for all rounds $r \in \mathbb{Z}_{\geq 0}$ and for some $0 < \alpha < 1$.

- $w_{ij}(r) = 0$ whenever $j \notin \mathcal{J}_i$ or $j \in \mathcal{R}_i(r)$;
- $w_{ij}(r) \geq \alpha$, $\forall j \in \mathcal{J}_i \setminus \mathcal{R}_i(r), i \in \mathcal{N}$;
- $\sum_{j=1}^{n} w_{ij}(r) = 1$, $\forall i \in \mathcal{N}$.

Together, these conditions imply that the updated value is a convex combination of values in $\mathcal{J}_i \setminus \mathcal{R}_i(r)$ with a uniform lower bound on the weights given by α.

The rest of this paper is concerned with determining necessary and sufficient conditions on the network topology for the normal nodes using Asynchronous W-MSR with parameter F to achieve resilient asymptotic consensus.

IV. RESILIENT CONSENSUS ANALYSIS

We begin with the following result showing that W-MSR with parameter F always satisfies the validity condition for resilient asymptotic consensus (RAC) under the F-total and F-local models. We then provide the definition of network robustness used in the analysis. Recall that $M_{\mathcal{N}}(r)$ and $m_{\mathcal{N}}(r)$ are the maximum and minimum values of the *normal* nodes in round r, respectively.

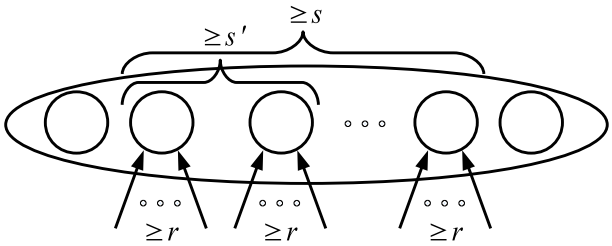

Fig. 1. Illustration of an (r, s)-edge reachable set of nodes.

Lemma 1: Suppose each normal node updates its value according to the Asynchronous W-MSR algorithm with parameter F under the F-total or F-local model. Then, for each normal node $i \in \mathcal{N}$, $x_i(r+1) \in [m_{\mathcal{N}}(r), M_{\mathcal{N}}(r)]$, regardless of the network topology. From this we conclude $m_{\mathcal{N}}(r+1) \geq m_{\mathcal{N}}(r)$ and $M_{\mathcal{N}}(r+1) \leq M_{\mathcal{N}}(r)$.

Proof: Suppose that one value, say $x_j(r)$, used in the update (1) satisfies $x_j(r) > M_{\mathcal{N}}(r)$. Then, by definition of $M_{\mathcal{N}}(r)$, j must be an adversary and $x_j(r) > x_i(r)$. Since i uses $x_j(r)$ in round r, there must be at least F more nodes in the neighborhood of i with values at least as large as $x_j(r)$. Hence, these nodes must also be adversaries, which contradicts the assumption that at most F in-neighbors of i are adversary nodes. Thus, $x_j(r) \leq M_{\mathcal{N}}(r)$. Similarly, we can show that $x_j(r) \geq m_{\mathcal{N}}(r)$. The result follows since $x_i(r+1)$ in (1) is a convex combination of values in $[m_{\mathcal{N}}(r), M_{\mathcal{N}}(r)]$. ■

A. Network Robustness

Before stating our main results, we require the following definitions.

Definition 4 ((r, s)-edge reachable set): Given a nontrivial digraph $\mathcal{D}$ and a nonempty subset of nodes $\mathcal{S}$, we say that $\mathcal{S}$ is an (r, s)-**edge reachable set** if there are at least s nodes in $\mathcal{S}$ with at least r in-neighbors outside of $\mathcal{S}$, where $r, s \in \mathbb{Z}_{\geq 0}$; i.e., given $\mathcal{X}_{\mathcal{S}}^r = \{i \in \mathcal{S} \colon |\mathcal{N}_i \setminus \mathcal{S}| \geq r\}$, then $|\mathcal{X}_{\mathcal{S}}^r| \geq s$. □

A general illustration of an (r, s)-edge reachable set of nodes is shown in Figure 1. The parameter s in the definition of (r, s)-edge reachability quantifies a lower bound on the number of nodes in the set with at least r in-neighbors outside $\mathcal{S}$. Observe that, in general, a set is (r, s')-edge reachable, for $s' \leq s$, if it is (r, s)-edge reachable. At one extreme, whenever there are no nodes in $\mathcal{S}$ with at least r in-neighbors outside of $\mathcal{S}$, then $\mathcal{S}$ is only $(r, 0)$-edge reachable. At the other extreme, $\mathcal{S}$ can be at most $(r, |\mathcal{S}|)$-edge reachable. Edge reachability is used to define the global property of robustness.

Definition 5 ((r, s)-robustness): A nonempty, nontrivial digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ on n nodes ($n \geq 2$) is (r, s)-**robust**, for nonnegative integers $r \in \mathbb{Z}_{\geq 0}$, $1 \leq s \leq n$, if for every pair of nonempty, disjoint subsets $\mathcal{S}_1$ and $\mathcal{S}_2$ of $\mathcal{V}$ at least one

of the following holds (recall $\mathcal{X}^r_{\mathcal{S}_k} = \{i \in \mathcal{S}_k \colon |\mathcal{N}_i \backslash \mathcal{S}_k| \geq r\}$ for $k \in \{1, 2\}$):

(i) $|\mathcal{X}^r_{\mathcal{S}_1}| = |\mathcal{S}_1|$;
(ii) $|\mathcal{X}^r_{\mathcal{S}_2}| = |\mathcal{S}_2|$;
(iii) $|\mathcal{X}^r_{\mathcal{S}_1}| + |\mathcal{X}^r_{\mathcal{S}_2}| \geq s$.

By convention, if $\mathcal{D}$ is empty or trivial ($n \leq 1$), then $\mathcal{D}$ is (0,1)-robust. If $\mathcal{D}$ is trivial, $\mathcal{D}$ is also (1,1)-robust.[4] □

Note that an $(r, 1)$-edge reachable set is abbreviated as r-edge reachable, and an $(r, 1)$-robust digraph is abbreviated as r-robust.

B. Necessary Condition for F-Total Malicious Model

The following is the main contribution of the paper and provides, for the first time, a *necessary and sufficient* condition for there to exist an algorithm that can achieve RAC in asynchronous networks with a local broadcast communication model under the F-total malicious model. First, we prove necessity. Then we show sufficiency by demonstrating that Asynchronous W-MSR achieves RAC consensus under this condition.

Theorem 1: If an asynchronous algorithm achieves resilient asymptotic consensus under the F-total or F-local malicious model in a nontrivial ($n \geq 2$) time-invariant asynchronous network under the local broadcast model, then the network is $(2F + 1, F + 1)$-robust.

Proof: Suppose there exists an asynchronous algorithm that achieves resilient asymptotic consensus in a nontrivial network that is not $(2F + 1, F + 1)$-robust. Then, there are nonempty, disjoint $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ such that none of the conditions $(i) - (iii)$ in Definition 5 hold (with $r = 2F + 1$ and $s = F + 1$). Suppose the initial value of each node in $\mathcal{S}_1$ is a and each node in $\mathcal{S}_2$ is b, with $a < b$. Let all other nodes have initial values taken from the interval $[a, b]$. Since $|\mathcal{X}^{2F+1}_{\mathcal{S}_1}| + |\mathcal{X}^{2F+1}_{\mathcal{S}_2}| \leq F$, suppose all nodes in $\mathcal{X}^{2F+1}_{\mathcal{S}_1}$ and $\mathcal{X}^{2F+1}_{\mathcal{S}_2}$ are malicious (which is allowed under both the F-total and F-local models) and keep their values constant for all rounds. With this assignment of adversaries, there is still at least one normal node in both $\mathcal{S}_1$ and $\mathcal{S}_2$ since $|\mathcal{X}^{2F+1}_{\mathcal{S}_1}| < |\mathcal{S}_1|$ and $|\mathcal{X}^{2F+1}_{\mathcal{S}_2}| < |\mathcal{S}_2|$, respectively.

Fix any normal node $i \in \mathcal{S}_1$ (and therefore, $i \in \mathcal{S}_1 \setminus \mathcal{X}^{2F+1}_{\mathcal{S}_1}$), and note that $|\mathcal{N}_i \backslash \mathcal{S}_1| \leq 2F$. Suppose the delays for messages from $q_i = \min\{F, |\mathcal{N}_i \setminus \mathcal{S}_1|\}$ nodes in $\mathcal{N}_i \setminus \mathcal{S}_1$ are arbitrarily large compared to all the other incoming messages to node i in round 0 (and the delays are large enough so that node i has become active). Then $\mathcal{N}_i \setminus \mathcal{R}_i(0)$ includes at most $|\mathcal{N}_i \setminus \mathcal{S}_1| - q_i \leq F$ values outside of $\mathcal{S}_1$ (which from the perspective of the update rule could all be adversary values). Other values used by the update rule for node i are from inside $\mathcal{S}_1$ (including node i's own value), so they have value a. Therefore, the update rule must set $x_i(1) = a$ to ensure the validity condition (more specifically, to ensure $M_{\mathcal{N}}(1) \leq M_{\mathcal{N}}(0)$). In a similar manner, one can argue that any normal node $j \in \mathcal{S}_2 \setminus \mathcal{X}^{2F+1}_{\mathcal{S}_2}$ must select $x_j(1) = b$. Finally, since $[m_{\mathcal{N}}(0), M_{\mathcal{N}}(0)] = [a, b]$, any normal node k in $\mathcal{V} \setminus (\mathcal{S}_1 \cup \mathcal{S}_2)$ must set its value $x_k(1) \in [a, b]$ to ensure the validity condition. Therefore, round 1 has the same distribution of values as round 0. By induction, we conclude that for each round $r \in \mathbb{Z}_{\geq 0}$ each node in $\mathcal{S}_1$ has value a, each node in $\mathcal{S}_2$ has value b, and all other nodes have values in $[a, b]$. Therefore, no consensus is achieved, which contradicts the assumption that there exists an asynchronous algorithm that achieves resilient asymptotic consensus in a network that is not $(2F + 1, F + 1)$-robust. ■

[4]The trivial graph is defined to be both (0,1)-robust and (1,1)-robust for consistency with properties of robust networks for $n > 1$.

C. Sufficient Condition for F-Total Malicious Model

Theorem 2 (Sufficiency): Consider a time-invariant asynchronous network under the local broadcast model. Suppose the communication is described by a digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$, where each normal node uses the Asynchronous W-MSR algorithm with parameter F. Then, under the F-total malicious model, resilient asymptotic consensus is achieved if the network topology is $(2F + 1, F + 1)$-robust.

Proof: Define $\Psi(r) = M_{\mathcal{N}}(r) - m_{\mathcal{N}}(r)$, which is a nonincreasing function of r by Lemma 1. Whenever the normal nodes are in agreement at some round $r_0 \in \mathbb{Z}_{\geq 0}$, then consensus is maintained in future rounds $r \geq r_0$. In the analysis that follows, recall that r is the round index and does not correspond to a common point in real time among the nodes. The difference in real time between when any two nodes actually execute round r may be quite large. The main point is that eventually each node will execute round $r \in \mathbb{Z}_{>0}$ because the network delay is finite and normal nodes only wait for at most $d_i - F$ incoming messages from neighbors. With this in mind, fix $r_0 \geq 0$ and assume $\Psi(r_0) > 0$. For $r \geq r_0$ and $\eta > 0$, define $\mathcal{S}_M(r, r_0, \eta) = \{j \in \mathcal{V} \colon x_j(r) > M_{\mathcal{N}}(r_0) - \eta\}$ and $\mathcal{S}_m(r, r_0, \eta) = \{j \in \mathcal{V} \colon x_j(r) < m_{\mathcal{N}}(r_0) + \eta\}$. Define $\epsilon_0 = \Psi(r_0)/2$ and define $\epsilon_j = \alpha\epsilon_{j-1}$ for $j = 1, 2, \ldots, N - 1$, where $N = \mathcal{N}$. It follows that $\epsilon_j = \alpha^j \epsilon_0 > 0$. By definition, the sets $\mathcal{S}_M(r_0, r_0, \epsilon_0)$ and $\mathcal{S}_m(r_0, r_0, \epsilon_0)$ are nonempty and disjoint. Because $\mathcal{D}$ is $(2F+1, F+1)$-robust and there are at most F malicious nodes in the network (F-total model), it follows that either there exists $i \in \mathcal{S}_M(r_0, r_0, \epsilon_0) \cap \mathcal{N}$ or there exists $i \in \mathcal{S}_m(r_0, r_0, \epsilon_0) \cap \mathcal{N}$, or there exists such i in both, such that i has at least $2F+1$ neighbors outside of its set. Suppose first that $i \in \mathcal{S}_M(r_0, r_0, \epsilon_0) \cap \mathcal{N}$ has at least $2F+1$ neighbors outside its set. Since at most $2F$ of these values will be ignored or removed (up to F ignored due to delays and F removed for being the smallest values in the in-neighborhood of node i), it follows that

$$\begin{aligned} x_i(r_0 + 1) &= \sum_{j \in \mathcal{J}_i \backslash \mathcal{R}_i(r_0)} w_{ij}(r_0) x_j(r_0) \\ &\leq \alpha(M_{\mathcal{N}}(r_0) - \epsilon_0) + (1 - \alpha) M_{\mathcal{N}}(r_0) \\ &\leq M_{\mathcal{N}}(r_0) - \alpha\epsilon_0 = M_{\mathcal{N}}(r_0) - \epsilon_1. \end{aligned}$$

Note that for any normal node not in $\mathcal{S}_M(r_0, r_0, \epsilon_0)$, the above inequality holds as well because any normal node always uses its own value in the update. From this, we conclude

$$|\mathcal{S}_M(r_0 + 1, r_0, \epsilon_1) \cap \mathcal{N}| < |\mathcal{S}_M(r_0, r_0, \epsilon_0) \cap \mathcal{N}|.$$

Similarly, if $i \in \mathcal{S}_m(r_0, r_0, \epsilon_0) \cap \mathcal{N}$ has at least $2F+1$ neighbors outside its set, then

$$\begin{aligned} x_i(r_0+1) &= \sum_{j \in \mathcal{J}_i \setminus \mathcal{R}_i(r_0)} w_{ij}(r_0) x_j(r_0) \\ &\geq \alpha(m_{\mathcal{N}}(r_0) + \epsilon_0) + (1-\alpha) m_{\mathcal{N}}(r_0) \\ &\geq m_{\mathcal{N}}(r_0) + \alpha\epsilon_0 = m_{\mathcal{N}}(r_0) + \epsilon_1. \end{aligned}$$

Similarly as above, this inequality holds for any normal node not in $\mathcal{S}_m(r_0, r_0, \epsilon_0)$. From this, we conclude

$$|\mathcal{S}_m(r_0+1, r_0, \epsilon_1) \cap \mathcal{N}| < |\mathcal{S}_m(r_0, r_0, \epsilon_0) \cap \mathcal{N}|.$$

By repeating this analysis, we can show by induction that as long as $\mathcal{S}_M(r_0+j, r_0, \epsilon_j) \cap \mathcal{N}$ and $\mathcal{S}_m(r_0+j, r_0, \epsilon_j) \cap \mathcal{N}$ are both nonempty, then either

$$|\mathcal{S}_M(r_0+j+1, r_0, \epsilon_{j+1}) \cap \mathcal{N}| < |\mathcal{S}_M(r_0+j, r_0, \epsilon_j) \cap \mathcal{N}|,$$

or

$$|\mathcal{S}_m(r_0+j+1, r_0, \epsilon_{j+1}) \cap \mathcal{N}| < |\mathcal{S}_m(r_0+j, r_0, \epsilon_j) \cap \mathcal{N}|,$$

or both hold. Since

$$|\mathcal{S}_M(r_0, r_0, \epsilon_0) \cap \mathcal{N}| + |\mathcal{S}_m(r_0, r_0, \epsilon_0) \cap \mathcal{N}| \leq |\mathcal{N}| = N,$$

there exists $T < N$ such that one of the sets

$$\begin{aligned} &\mathcal{S}_M(r_0+T, r_0, \epsilon_T) \cap \mathcal{N}, \\ &\mathcal{S}_m(r_0+T, r_0, \epsilon_T) \cap \mathcal{N}, \end{aligned}$$

or both, is empty. It follows that in the former case,

$$M_{\mathcal{N}}(r_0+T) \leq M_{\mathcal{N}}(r_0) - \epsilon_T,$$

and in the latter case,

$$m_{\mathcal{N}}(r_0+T) \geq m_{\mathcal{N}}(r_0) + \epsilon_T.$$

Since

$$\epsilon_0 > \epsilon_1 > \cdots > \epsilon_T \geq \epsilon_{N-1} > 0,$$

we have

$$\begin{aligned} \Psi(r_0+N-1) - \Psi(r_0) &\leq \Psi(r_0+T) - \Psi(r_0) \\ &\leq (M_{\mathcal{N}}(r_0+T) - M_{\mathcal{N}}(r_0)) \\ &\qquad + (m_{\mathcal{N}}(r_0) - m_{\mathcal{N}}(r_0+T)) \\ &\leq -\epsilon_T \\ &\leq -\epsilon_{N-1}. \end{aligned}$$

Therefore,

$$\Psi(r_0+N-1) \leq \Psi(r_0)(1-\alpha^{N-1}/2).$$

Define $c = (1-\alpha^{N-1}/2)$. Since c is not a function of r_0 and r_0 was chosen arbitrarily, it follows that

$$\Psi(r_0 + k(N-1)) \leq c^k \Psi(r_0),$$

for all $k \in \mathbb{Z}_{\geq 0}$. Because $c < 1$, it follows that $\Psi(r) \to 0$ as $r \to \infty$. ■

D. Sufficient Condition for F-Local Malicious Model

Theorem 3 (Sufficient Condition, F-Local): Consider a time-invariant asynchronous network under the local broadcast model. Suppose the communication is described by a digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$, where each normal node uses the Asynchronous W-MSR algorithm with parameter F. Then, under the F-local malicious model, resilient asymptotic consensus is achieved if the network topology is $(3F+1)$-robust.

Proof: In this case, the sets $\mathcal{S}_M$ and $\mathcal{S}_m$ are defined to include only normal nodes. Then, the $(3F+1)$-robust assumption under the F-local model ensures at least one normal value outside of either $\mathcal{S}_M$ or $\mathcal{S}_m$ will be used in the update (at most up to F values outside the set could be from adversaries, F smallest or largest values are removed, and F values are ignored due to time delays, leaving still one normal value from outside that is used). The rest of the analysis is identical to the proof of Theorem 2. ■

V. RELATED WORK

Although the Byzantine approximate agreement problem was posed more than twenty-five years ago [20], the necessary and sufficient topological condition on a time-invariant network for the existence of a successful asymptotic consensus algorithm in the presence of up to F Byzantine nodes has been an open problem (for both synchronous and asynchronous networks) until very recently [21], [39], [40]. Synchronous networks under the F-total Byzantine model are studied in [21], [39], and both synchronous and asynchronous networks are studied in [40]. In [21], Vaidya et al. provide the tight condition required in synchronous directed networks for the existence of a successful algorithm that ensures resilient asymptotic consensus in the presence of up to F Byzantine faulty nodes (F-total model). In order to state the condition, we require the following definition, which provides a common notation for the definitions considered separately in [21] and [40].

Definition 6: For nonempty, disjoint sets of nodes $A, B \subset \mathcal{V}$, $A \overset{r}{\Rightarrow} B$ if and only if there exists a node $v \in B$ that has at least r in-neighbors in A; i.e., $|\mathcal{N}_v \cap A| \geq r$. $A \overset{r}{\not\Rightarrow} B$ if and only if $A \overset{r}{\Rightarrow} B$ is not true. □

Given the relation of Definition 6, the tight condition for the synchronous case may be stated as follows. Fix any quadruple of sets of nodes $\mathcal{F}, L, C, R$ that form a partition[5] of $\mathcal{V}$ such that $0 \leq |\mathcal{F}| \leq F$, $|L| > 0$, and $|R| > 0$. Then at least one of the two following conditions must hold true: (i) $R \cup C \overset{F+1}{\Rightarrow} L$ or (ii) $L \cup C \overset{F+1}{\Rightarrow} R$. Observe that this condition requires sufficient redundancy of directed edges between subsets of normal nodes in the network (the nodes in $\mathcal{F}$ account for the Byzantine nodes). Note that the condition can be restated in terms of robustness; i.e., the subdigraph induced by the normal nodes must be $(F+1)$-robust.

Vaidya et al. present an equivalent condition in [40], which uses the following concepts. The *decomposition digraph*

[5]Here, sets $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_p \subseteq \mathcal{S}$ are said to form a **partition** of set $\mathcal{S}$ if $\cup_{i=1}^{p} \mathcal{S}_i = \mathcal{S}$ and $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for $i \neq j$. Note that in this context, some of the sets in the partition may be empty.

$\mathcal{D}^d = (\mathcal{V}^d, \mathcal{E}^d)$ of $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is constructed from $\mathcal{D}$ by associating a node $v_k \in \mathcal{V}^d$ to each strongly connected component $\mathcal{C}_k$ of $\mathcal{D}$. A directed edge $(i, j) \in \mathcal{E}^d$ exists if and only if there is a node in component $\mathcal{C}_j$ reachable from every node in component $\mathcal{C}_i$. Note that the decomposition digraph is always a directed acyclic graph [41]. A *source component* of $\mathcal{D}$ is a strongly connected component $\mathcal{C}_k$ of $\mathcal{D}$ such that v_k is not reachable from any other node in $\mathcal{D}^d$. Finally, a *reduced digraph* $\mathcal{D}_\mathcal{F} = (\mathcal{V}_\mathcal{F}, \mathcal{E}_\mathcal{F})$ of $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is any subdigraph of $\mathcal{D}$ such that $\mathcal{F} \subset \mathcal{V}$, $\mathcal{V}_\mathcal{F} = \mathcal{V} \setminus \mathcal{F}$, and $\mathcal{E}_\mathcal{F}$ is obtained by first removing all directed edges in $\mathcal{E}$ that are incident with nodes in $\mathcal{F}$ and then removing up to F other incoming edges at each node in $\mathcal{V}_\mathcal{F}$. The alternative condition states that every reduced digraph $\mathcal{D}_\mathcal{F}$ with $|\mathcal{F}| < |\mathcal{V}|$ and $|\mathcal{F}| \leq F$ must contain *exactly* one source component. It is shown in [40] that the unique source component in any such reduced digraph must contain at least $F + 1$ nodes. By associating $\mathcal{F}$ with the set of Byzantine nodes, these results say there must be a set of normal nodes (the source nodes in the reduced digraph) that are capable of disseminating their information resiliently throughout the rest of the network. Moreover, the number of source nodes in any reduced digraph $\mathcal{D}_\mathcal{F}$ must outnumber the Byzantine faulty nodes.

The necessary and sufficient condition for time-invariant asynchronous networks with a point-to-point communication model in the presence of up to F Byzantine nodes is given in [40]. The condition can also be stated using the relation of Definition 6. Fix any quadruple of sets of nodes $\mathcal{F}, L, C, R$ that form a partition of $\mathcal{V}$ such that $0 \leq |\mathcal{F}| \leq F$, $|L| > 0$, and $|R| > 0$. Then at least one of the two following conditions must hold true: (i) $R \cup C \overset{2F+1}{\Rightarrow} L$ or (ii) $L \cup C \overset{2F+1}{\Rightarrow} R$. Note that the condition can be restated in terms of robustness; i.e., the subdigraph induced by the normal nodes must be $(2F + 1)$-robust.

The concept of robust networks is introduced by Zhang and Sundaram in [24], where it is shown to be a useful property in studying the resilience of distributed algorithms (including consensus and broadcast algorithms) in the presence of F-local adversaries. A refined definition (which is the one presented here, with finer granularity through the introduction of parameter s) is given in [25], [34], in order to formulate the necessary and sufficient condition to achieve resilient asymptotic consensus in time-invariant synchronous networks under the F-total malicious model.

Note that robust networks are quite common. In [24], [25], [34], it is shown that the robustness of the seed graph in the well-known preferential attachment model for scale-free networks [42] is maintained throughout the growth of the network provided the number of edges added each iteration is sufficiently large. Moreover, it is shown in [43] that random networks also exhibit robustness properties.

VI. CONCLUSION

This paper provides the necessary and sufficient condition required on the network topology for the existence of a consensus algorithm that achieves resilient asymptotic consensus (RAC) in time-invariant asynchronous networks in the presence of up to F malicious nodes under a local broadcast model of communication (F-total malicious model). A sufficient condition is given for the case when the number of malicious nodes in any normal node's neighborhood is bounded by F (F-local malicious model). The topological conditions are stated in terms of the robustness of the network, which is a novel graph theoretic property that provides a measure on the amount of redundancy of directed edges that exist between subsets of nodes in the network. Together with other recent results [21], [25], [39], [40], the tight conditions on the network topology for successful RAC algorithms have been determined for the F-total Byzantine and malicious models (for both synchronous and asynchronous time-invariant networks). The tight topological condition for the F-local malicious model is still an open problem left for future work.

REFERENCES

[1] N. A. Lynch, *Distributed Algorithms*. San Francisco, California: Morgan Kaufmann Publishers Inc., 1997.

[2] J. Hromkovic, R. Klasing, A. Pelc, P. Ruzicka, and W. Unger, *Dissemination of Information in Communication Networks*. Springer-Verlag, 2005.

[3] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[4] W. Ren, R. W. Beard, and E. M. Atkins, "Information consensus in multivehicle cooperative control," *IEEE Control Systems Magazine*, vol. 27, no. 2, pp. 71–82, April 2007.

[5] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, June 2003.

[6] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: algorithms and theory," *IEEE Transactions on Automatic Control*, vol. 51, no. 3, pp. 401–420, March 2006.

[7] I. A. F. Ihle, M. Arcak, and T. I. Fossen, "Passivity-based designs for synchronized path-following," *Automatica*, vol. 43, no. 9, pp. 1508–1518, 2007.

[8] I. D. Schizas, G. Mateos, and G. B. Giannakis, "Distributed LMS for consensus-based in-network adaptive processing," *IEEE Transactions on Signal Processing*, vol. 57, no. 6, pp. 2365–2382, June 2009.

[9] D. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Englewood Cliffs, NJ: Prentice-Hall, Inc., 1989.

[10] U. Münz, A. Papachristodoulou, and F. Allgöwer, "Delay robustness in consensus problems," *Automatica*, vol. 46, no. 8, pp. 1252–1265, 2010.

[11] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004.

[12] A. Kashyap, T. Başar, and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007.

[13] V. Gupta, C. Langbort, and R. M. Murray, "On the robustness of distributed algorithms," in *IEEE Conference on Decision and Control*, San Diego, California, Dec. 2006, pp. 3473–3478.

[14] M. Pease, R. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *Journal of the ACM*, vol. 27, no. 2, pp. 228–234, 1980.

[15] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, and M. Medard, "Resilient network coding in the presence of Byzantine adversaries," in *26th IEEE International Conference on Computer Communications, INFOCOM*, Anchorage, AL, May 2007, pp. 616–624.

[16] N. Agmon and D. Peleg, "Fault-tolerant gathering algorithms for autonomous mobile robots," *SIAM Journal on Computing*, vol. 36, no. 1, pp. 56–82, July 2006.

[17] X. Défago, M. Gradinariu, S. Messika, and P. Raipin-Parvédy, "Fault-tolerant and self-stabilizing mobile robots gathering," in *Distributed Computing*, ser. Lecture Notes in Computer Science, S. Dolev, Ed. Springer Berlin, Heidelberg, 2006, vol. 4167, pp. 46–60.

[18] Z. Bouzid, M. G. Potop-Butucaru, and S. Tixeuil, "Optimal Byzantine-resilient convergence in uni-dimensional robot networks," *Theoretical Computer Science*, vol. 411, no. 34-36, pp. 3154–3168, July 2010.

[19] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 2, pp. 382–401, 1982.

[20] D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl, "Reaching approximate agreement in the presence of faults," *Journal of the ACM*, vol. 33, no. 3, pp. 499–516, 1986.

[21] N. H. Vaidya, L. Tseng, and G. Liang, "Iterative approximate Byzantine consensus in arbitrary directed graphs," in *Proceedings of the ACM symposium on Principles of distributed computing (PODC)*, Madeira, Portugal, 2012, pp. 365–374.

[22] A. Pelc and D. Peleg, "Broadcasting with locally bounded Byzantine faults," in *Information Processing Letters*, 2005, pp. 109–115.

[23] A. Ichimura and M. Shigeno, "A new parameter for a broadcast algorithm with locally bounded Byzantine faults," *Information Processing Letters*, vol. 110, pp. 514–517, 2010.

[24] H. Zhang and S. Sundaram, "Robustness of information diffusion algorithms to locally bounded adversaries," in *Proceedings of the American Control Conference*, Montréal, Canada, 2012, pp. 5855–5861.

[25] H. J. LeBlanc, H. Zhang, S. Sundaram, and X. Koutsoukos, "Consensus of multi-agent networks in the presence of adversaries using only local information," in *Proceedings of the 1st International Conference on High Confidence Networked Systems (HiCoNS)*, Beijing, China, 2012, pp. 1–10.

[26] A. A. Cárdenas, S. Amin, and S. S. Sastry, "Research challenges for the security of control systems," in *Proceedings of the 3rd conference on Hot topics in security*, San Jose, CA, July 2008, pp. 1–6.

[27] S. Bhattacharya and T. Başar, "Spatial approaches to broadband jamming in heterogeneous mobile networks: a game-theoretic approach," *Autonomous Robots*, vol. 31, no. 4, pp. 367–381, 2011.

[28] S. Amin, A. A. Cárdenas, and S. S. Sastry, "Safe and secure networked control systems under denial-of-service attacks," in *Proceedings of the 12th International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '09, San Francisco, CA, 2009, pp. 31–45.

[29] Y. Mo, E. Garone, A. Casavola, and B. Sinopoli, "False data injection attacks against state estimation in wireless sensor networks," in *IEEE Conference on Decision and Control*, Dec. 2010, pp. 5967–5972.

[30] Y. Mo and B. Sinopoli, "Secure control against replay attacks," in *47th Annual Allerton Conference on Communication, Control, and Computing*, Sept. 2009, pp. 911–918.

[31] M. Zhu and S. Martinez, "Attack-resilient distributed formation control via online adaptation," in *IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, Orlando, FL, Dec. 2011, pp. 6624–6629.

[32] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation via linear iterative strategies in the presence of malicious agents," *IEEE Transactions on Automatic Control*, vol. 56, no. 7, pp. 1495–1508, July 2011.

[33] H. J. LeBlanc and X. D. Koutsoukos, "Low complexity resilient consensus in networked multi-agent systems with adversaries," in *Proceedings of the 15th international conference on Hybrid systems: computation and control*, ser. (HSCC '12), Beijing, China, 2012, pp. 5–14.

[34] H. J. LeBlanc, H. Zhang, X. D. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE Journal on Selected Areas in Communications*, 2012, submitted and under review.

[35] F. Pasqualetti, A. Bicchi, and F. Bullo, "Consensus computation in unreliable networks: A system theoretic approach," *IEEE Transactions on Automatic Control*, vol. 57, no. 1, pp. 90–104, Jan. 2012.

[36] H. J. LeBlanc and X. D. Koutsoukos, "Consensus in networked multi-agent systems with adversaries," in *Proceedings of the 14th international conference on Hybrid systems: computation and control*, ser. (HSCC '11), Chicago, IL, 2011, pp. 281–290.

[37] 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.

[38] Q. Li and D. Rus, "Global clock synchronization in sensor networks," *IEEE Transactions on Computers*, vol. 55, no. 2, pp. 214–226, Feb. 2006.

[39] N. H. Vaidya, "Matrix representation of iterative approximate Byzantine consensus in directed graphs," *CoRR*, vol. abs/1201.1888, 2012.

[40] N. H. Vaidya, L. Tseng, and G. Liang, "Iterative approximate Byzantine consensus in arbitrary directed graphs - Part II: Synchronous and asynchronous systems," *CoRR*, vol. abs/1202.6094, 2012.

[41] S. Dasgupta, C. Papadimitriou, and U. Vazirani, *Algorithms*. McGraw-Hill Higher Education, 2006.

[42] R. Albert and A. L. Barabási, "Statistical mechanics of complex networks," *Rev. Mod. Phys.*, vol. 74, no. 1, pp. 47–97, Jan. 2002.

[43] H. Zhang and S. Sundaram, "Robustness of complex networks: Reaching consensus despite adversaries," *CoRR*, vol. abs/1203.6119, 2012, to appear at the 2012 Conference on Decision and Control.