

Resilient Asymptotic Consensus in Robust Networks

Heath J. LeBlanc, *Member, IEEE*, Haotian Zhang, *Student Member, IEEE*,
Xenofon Koutsoukos, *Senior Member, IEEE*, Shreyas Sundaram, *Member, IEEE*

Abstract

This paper addresses the problem of resilient in-network consensus in the presence of misbehaving nodes. Secure and fault-tolerant consensus algorithms typically assume knowledge of nonlocal information; however, this assumption is not suitable for large-scale dynamic networks. To remedy this, we focus on local strategies that provide resilience to faults and compromised nodes. We design a consensus protocol based on local information that is resilient to worst-case security breaches, assuming the compromised nodes have full knowledge of the network and the intentions of the other nodes. We provide necessary and sufficient conditions for the normal nodes to reach asymptotic consensus despite the influence of the misbehaving nodes under different threat assumptions. We show that traditional metrics such as connectivity are not adequate to characterize the behavior of such algorithms, and develop a novel graph-theoretic property referred to as *network robustness*. Network robustness formalizes the notion of redundancy of direct information exchange between subsets of nodes in the network, and is a fundamental property for analyzing the behavior of certain distributed algorithms that use only local information.

Index Terms

H. LeBlanc is with the Electrical & Computer Engineering and Computer Science Department, Ohio Northern University, Ada, OH, 45810 USA h-leblanc@onu.edu.

H. Zhang and S. Sundaram are with the Department of Electrical and Computer Engineering at the University of Waterloo, Waterloo, Ontario, Canada {h223zhan,ssundara}@uwaterloo.ca.

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

Manuscript received April 9, 2012; revised December 30, 2012. Some of the results in this paper were presented in preliminary form at the First Conference on High-Confidence Networked Systems (HiCoNS 2012) [1] and at the 2012 American Control Conference [2].

I. INTRODUCTION

Engineered systems have undergone a paradigm shift from centralized to distributed, propelled by advances in networking and low-cost, high performance embedded devices. These advances have enabled a transition from end-to-end routing of information in large-scale networked systems to *in-network computation* of aggregate quantities of interest [3]. In-network computing offers certain performance advantages, including reduced latency, smaller communication overhead, and greater robustness to node and link failures.

A fundamental challenge of in-network computation is that the quantities of interest must be calculated using only *local information*, i.e., information obtained by each node through sensor measurements, calculations, or communication only with neighbors in the network. Another important challenge is that large-scale distributed systems have many potential vulnerable points for failures or attacks. To obtain the desired computational results, it is important to design the in-network algorithms to be able to withstand the compromise of a subset of the nodes and still ensure *some notion of correctness* (possibly at a degraded level of performance). We refer to such a networked system as being *resilient* to adversaries. Given the growing threat of malicious attacks in large-scale cyber-physical systems, this is an important and challenging problem [4].

One of the most important objectives in networked systems is to reach consensus on a quantity of interest [5]–[8]. Consensus is fundamental to diverse applications such as data aggregation [9], distributed estimation [10], distributed optimization [11], distributed classification [12], and flocking [13]. Reaching consensus (and more generally, transmitting information) resiliently in the presence of faulty or misbehaving nodes has been studied extensively in distributed computing [5], [14]–[18], communication networks [19], [20], and mobile robotics [21]–[23]. Among other things, it has been shown that given F (worst-case) adversarial nodes, there exists a strategy for these nodes to disrupt consensus if the network connectivity¹ is $2F$ or less. Conversely, if the network connectivity is at least $2F + 1$, then there exist strategies for the

¹The network connectivity is defined as the smaller of the two following values: (i) the size of a minimal vertex cut and (ii) $n - 1$, where n is the number of nodes in the network.

normal nodes to use that ensure consensus is reached (under the local broadcast model of communication) [5], [24], [25]. However, these consensus algorithms either require that normal nodes have at least some nonlocal information (e.g., knowledge of multiple independent paths in the network between themselves and other nodes) or assume that the network is *complete*, i.e., all-to-all communication or sensing [14], [21]–[23], [26]. Moreover, these algorithms tend to be computationally expensive. Therefore, there is a need for resilient consensus algorithms that have *low complexity* and operate *using only local information* (i.e., without knowledge of the network topology and the identities of non-neighboring nodes). A key challenge is to characterize fundamental topological properties that allow the normal nodes to compute an appropriate consensus value, despite the influence of misbehaving nodes.

The faulty or misbehaving nodes can be characterized by threat models and scope of threat assumptions. Examples of fault or threat models include *non-colluding* [25], *malicious* [24]–[26], *Byzantine* [14], [21], [27], [28], or *crashed* [21], [22] nodes. Typically, the scope of the faults or threats is assumed to be bounded by a constant, i.e., at most F out of n nodes fail or are compromised. We refer to this as the *F-total model*. Alternatively, the scope may be local; e.g., at most F neighbors of any normal node fail (*F-local model*), or at most a fraction f of neighbors are compromised (*f-fraction local model*).

A. Previous Work on Consensus With Only Local Information

In [29], the authors introduced the *Approximate Byzantine Consensus* problem, in which the normal nodes are required to achieve *approximate agreement*² (i.e., they should converge to a relatively small convex hull contained in the convex hull defined by their initial values) in the presence of F -total Byzantine faults in finite time. They consider only complete networks (where there is a direct connection between every pair of nodes), and they propose the following algorithm: each node disregards the largest and smallest F values received from its neighbors and updates its state to be the average of a carefully chosen subset of the remaining values. This algorithm was extended to a family of algorithms, named the *Mean-Subsequence-Reduced (MSR)* algorithms, in [30]. Although the research on Approximate Byzantine Consensus for complete

²If the network is synchronous, and if one allows $t \rightarrow \infty$, then approximate agreement is equivalent to asymptotic consensus.

networks is mature, there are few papers that have attempted to analyze this algorithm in more general topologies [31], and even then, only certain special networks have been investigated.

Recently, we have studied resilient algorithms in the presence of misbehaving nodes [32], [33]. In [26], we proposed a continuous-time variation of the MSR algorithms, named the *Adversarial Robust Consensus Protocol (ARC-P)*, to solve asymptotic consensus under the F -total malicious model. The results of [26] were extended to both malicious and Byzantine threat models in networks with constrained information flow and dynamic network topology in [27]. The sufficient conditions studied in [27] are stated in terms of in-degrees and out-degrees of nodes in the network and are shown to be sharp, i.e., if the conditions are relaxed, even minimally, then there are examples in which the relaxed conditions are not sufficient. In [2], we generalized the MSR algorithm as the *Weighted-Mean-Subsequence-Reduced (W-MSR)* algorithm and studied general distributed algorithms with F -local malicious adversaries.

In a recent paper, developed independently of our work, Vaidya *et al.* have characterized tight conditions for resilient consensus using the MSR algorithm when the threat model is Byzantine and the scope is F -total [28]. The network constructions used in [28] are very similar to the robust digraphs presented here. In particular, the networks in [28] also require redundancy of direct information exchange between subsets of nodes in the network.

In contrast to the deterministic approach taken here, gossip algorithms have been studied for in-network computation of aggregate functions such as sums, averages, and quantiles [9]. In such algorithms, each node chooses at random a single neighbor to communicate with in each round. This scheme limits the required computational, communication, and energy resources, and provides some robustness against time-varying topologies and random node and link failures [34]. However, we are not aware of any work that studies the resilience of gossip-based algorithms to malicious attacks.

B. Contributions

In this paper, we show that traditional graph theoretic properties such as connectivity and minimum degree, which have played a vital role in characterizing the resilience of distributed algorithms (see [5], [24]), are not adequate when the nodes make purely local decisions (i.e., without knowing nonlocal aspects of the network topology). Instead, we introduce a novel topological property, referred to as *network robustness*, and show that this concept is *the key*

property for reasoning about the ability of purely local algorithms to succeed. In particular, we provide a comprehensive characterization of the network topologies where algorithms such as W-MSR (which uses only local information and operates in synchronous networks) can succeed despite the presence of a broad class of adversaries. We establish results for both malicious and Byzantine threats, where the scope is F -total, F -local, and f -fraction local, and the network is time-invariant or time-varying. For the case of time-invariant networks, we provide, for the first time, a tight condition for the W-MSR algorithm to succeed under the F -total malicious model. Furthermore, we give tight conditions for F -local and F -total Byzantine threats (the proof for the F -total Byzantine model is different than the proof given in [28], and is stated for the more general W-MSR algorithm and in terms of network robustness). We prove separate necessary and sufficient conditions for the W-MSR algorithm under the F -local malicious, f -fraction local malicious, and f -fraction local Byzantine threat models. While there is a gap between the necessary and sufficient conditions in these latter cases, we show that the sufficient condition for the F -local malicious model is *sharp*, i.e., relaxing the condition leads to examples in which consensus is not achieved. For all threat models, we provide sufficient conditions for the case of time-varying networks.

In addition to the results on resilient asymptotic consensus, we also examine properties of robust digraphs. We demonstrate the connectivity and degree properties of robust digraphs, explore the robustness maintained after edge removal, and describe how to compare the relative robustness of different digraphs. Finally, we provide a method that enables the construction of robust networks and show that the preferential attachment mechanism for generating complex networks is a special case of this method (and therefore produces robust networks).

The rest of the paper is organized as follows. Section II introduces the problem of resilient consensus. Section III presents the W-MSR algorithm. Section IV demonstrates the inadequacy of connectivity as a metric to analyze the behavior of the W-MSR algorithm, and formally introduces the notion of *network robustness*. The main results are given in Section V. A simulation example is presented in Section VI. In Section VII, we discuss properties of network robustness and provide a construction for robust networks. Finally, some conclusions are given in Section VIII.

C. Notation and Graph Terminology

Throughout this paper, we denote the set of integers by \mathbb{Z} and the set of real numbers by \mathbb{R} . The set of integers greater than or equal to some integer $q \in \mathbb{Z}$ is denoted $\mathbb{Z}_{\geq q}$. Given $a \in \mathbb{R}$, the *ceiling* of a , denoted $\lceil a \rceil$, is the smallest integer that is greater than or equal to a . Similarly, the *floor* of a , denoted $\lfloor a \rfloor$, is the largest integer less than or equal to a . The cardinality of a set \mathcal{S} is denoted by $|\mathcal{S}|$. Given sets $\mathcal{S}_1, \mathcal{S}_2$, the reduction of \mathcal{S}_1 by \mathcal{S}_2 is denoted $\mathcal{S}_1 \setminus \mathcal{S}_2 = \{x \in \mathcal{S}_1 : x \notin \mathcal{S}_2\}$.

A finite simple directed graph, or just *digraph*, is denoted $\mathcal{D} = (\mathcal{V}, \mathcal{E})$, in which \mathcal{V} is the *node set* and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the *directed edge set*. With a slight abuse of terminology, we will often refer to the network and the digraph that models the topology of the network synonymously. The *underlying graph* $\mathcal{G}(\mathcal{D})$ is defined by replacing directed edges of \mathcal{D} by undirected ones, resulting in the edge set $\mathcal{E}_{\mathcal{G}}$. A digraph $\mathcal{D}' = (\mathcal{V}', \mathcal{E}')$ is a *subdigraph* of \mathcal{D} , written $\mathcal{D}' \subseteq \mathcal{D}$, if $\mathcal{V}' \subseteq \mathcal{V}$ and $\mathcal{E}' \subseteq \mathcal{E}$. A digraph $\mathcal{D}' = (\mathcal{V}', \mathcal{E}')$ is *isomorphic* to \mathcal{D} if there exists a bijection $\psi: \mathcal{V} \rightarrow \mathcal{V}'$ such that $(i, j) \in \mathcal{E}$ if and only if $(\psi(i), \psi(j)) \in \mathcal{E}'$.

A *path* is a sequence of distinct vertices i_0, i_1, \dots, i_k such that $(i_j, i_{j+1}) \in \mathcal{E}$, $j = 0, 1, \dots, k-1$. We use the notion of path to define different forms of connectedness. We say that \mathcal{D} is *strongly connected* if for every $i, j \in \mathcal{V}$, there exists a path starting at i and ending at j . If the underlying graph is connected, then \mathcal{D} is *weakly connected*. Alternatively, if the underlying graph is disconnected, then \mathcal{D} is *disconnected*. A digraph has a *rooted out-branching* if there exists a node r , the root, such that for each $i \in \mathcal{V}$, there exists a path from r to i .

II. PROBLEM FORMULATION

Consider a time-varying network modeled by the *digraph* $\mathcal{D}[t] = (\mathcal{V}, \mathcal{E}[t])$, where $\mathcal{V} = \{1, \dots, n\}$ is the *node set* and $\mathcal{E}[t] \subset \mathcal{V} \times \mathcal{V}$ is the *directed edge set* at time-step $t \in \mathbb{Z}_{\geq 0}$. The node set is partitioned into a set of *normal nodes* \mathcal{N} and a set of *adversary nodes* \mathcal{A} which is unknown a priori to the normal nodes. Each directed edge $(j, i) \in \mathcal{E}[t]$ models *information flow* and indicates that node i can be influenced by (or receive information from) node j at time-step t . The set of *in-neighbors*, or just *neighbors*, of node i at time-step t is defined as $\mathcal{V}_i[t] = \{j \in \mathcal{V} : (j, i) \in \mathcal{E}[t]\}$ and the (in-)degree of i is denoted $d_i[t] = |\mathcal{V}_i[t]|$. Likewise, the set of *out-neighbors* of node i at time-step t is defined as $\mathcal{V}_i^{\text{out}}[t] = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}[t]\}$. Because each node has access to its own state at time-step t , we also consider the *inclusive neighbors*

of node i , denoted $\mathcal{J}_i[t] = \mathcal{V}_i[t] \cup \{i\}$. Note that time-invariant networks are represented simply by dropping the dependence on t .

A. Update Model

Suppose that each node $i \in \mathcal{N}$ begins with some private value $x_i[0] \in \mathbb{R}$ (which could represent a measurement, optimization variable, vote, etc.). The nodes interact synchronously by conveying their value to (out-)neighbors in the network. Each normal node updates its own value over time according to a prescribed rule, which is modeled as

$$x_i[t + 1] = f_i(\{x_j^i[t]\}), \quad j \in \mathcal{J}_i[t], i \in \mathcal{N}, t \in \mathbb{Z}_{\geq 0},$$

where $x_j^i[t]$ is the value sent from node j to node i at time-step t , and $x_i^i[t] = x_i[t]$. 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 compute some desired function. 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

Definition 1: A node $i \in \mathcal{A}$ is said to be **Byzantine** if it does not send the same value to all of its neighbors at some time-step, or if it applies some other function $f'_i(\cdot)$ at some time-step.

Definition 2: A node $i \in \mathcal{A}$ is said to be **malicious** if it sends $x_i[t]$ to all of its neighbors at each time-step, but applies some other function $f'_i(\cdot)$ at some time-step.

Note that both malicious and Byzantine nodes are allowed to update their states arbitrarily (perhaps colluding with other malicious or Byzantine nodes to do so). The only difference is in their capacity for duplicity. If the network is realized through sensing or broadcast communication, it is natural to assume that the out-neighbors receive the same information, thus motivating the definition of a malicious node. If the network is point-to-point, however, Byzantine behavior is possible. Note that all malicious nodes are Byzantine, but not vice versa. When we do not need to explicitly distinguish between Byzantine and malicious threats, we simply say those nodes are *misbehaving*.

C. Scope of Threats

Having defined the types of misbehavior in the system, it is necessary to define the *number* of misbehaving nodes. While there are various stochastic models that could be used to formalize the scope of threats, we use a deterministic approach and consider upper bounds on the number of compromised nodes either in the network (F -total) or in each node's neighborhood (F -local). To account for varying degrees of different nodes, we also introduce a fault model that considers an upper bound on the *fraction* of compromised nodes in any node's neighborhood.

Definition 3 (F -total set): A set $\mathcal{S} \subset \mathcal{V}$ is **F -total** if it contains at most F nodes in the network, i.e., $|\mathcal{S}| \leq F$, $F \in \mathbb{Z}_{\geq 0}$.

Definition 4 (F -local set): A set $\mathcal{S} \subset \mathcal{V}$ is **F -local** if it contains at most F nodes in the neighborhood of the other nodes for all t , i.e., $|\mathcal{V}_i[t] \cap \mathcal{S}| \leq F$, $\forall i \in \mathcal{V} \setminus \mathcal{S}$, $\forall t \in \mathbb{Z}_{\geq 0}$, $F \in \mathbb{Z}_{\geq 0}$.

Definition 5 (f -fraction local set): A set $\mathcal{S} \subset \mathcal{V}$ is **f -fraction local** if it contains at most a fraction f of nodes in the neighborhood of the other nodes for all t , i.e., $|\mathcal{V}_i[t] \cap \mathcal{S}| \leq f|\mathcal{V}_i[t]|$, $\forall i \in \mathcal{V} \setminus \mathcal{S}$, $\forall t \in \mathbb{Z}_{\geq 0}$, $0 \leq f \leq 1$.

It should be noted that in time-varying network topologies, the local properties defining an F -local set (or an f -fraction local set) must hold at all time instances. These definitions facilitate the following scope of threat models.

Definition 6: A set of adversary nodes is **F -totally bounded**, **F -locally bounded** or **f -fraction locally bounded** if it is an F -total set, F -local set or f -fraction local set, respectively. We refer to these threat scopes as the **F -total**, **F -local**, and **f -fraction local models**, respectively.

F -totally bounded faults have been studied in distributed computing [5], [14], [28] and mobile robotics [21]–[23] for both stopping (or crash) failures and Byzantine failures. The F -locally bounded fault model has been studied in the context of fault-tolerant broadcasting [35], [36]. However, to the best of our knowledge, there are no prior works discussing the f -fraction local model; our investigation of this model is inspired by ideas pertaining to *contagion* in social and economic networks [37], where a node will accept some new information (behavior or technology) if more than a certain fraction of its neighbors has adopted it. However, these previous works do not consider faulty or malicious behavior, and our definition is a natural extension to the existing interpretations.

D. Resilient Asymptotic Consensus

Given the threat model and scope of threats, we formally define resilient asymptotic consensus. Let $M[t]$ and $m[t]$ be the *maximum* and *minimum* values of the *normal* nodes at time-step t , respectively.

Definition 7 (Resilient Asymptotic Consensus): The normal nodes are said to achieve **resilient asymptotic consensus** in the presence of (a) F -totally bounded, (b) F -locally bounded, or (c) f -fraction locally bounded misbehaving (Byzantine or malicious) nodes if

- $\exists L \in \mathbb{R}$ such that $\lim_{t \rightarrow \infty} x_i[t] = L$ for all $i \in \mathcal{N}$, and
- $[m[0], M[0]]$ is an invariant set (i.e., the normal values remain in the interval for all t),

for any choice of initial values. Whenever the scope of threat is understood, we simply say that the normal nodes reach **asymptotic consensus**.

The resilient asymptotic consensus problem has two important conditions. First, the normal nodes must reach asymptotic consensus in the presence of misbehaving nodes given a particular threat model (e.g., malicious) and scope of threat (e.g., F -total). This is a condition on *agreement*. Additionally, it is required that the interval containing the initial values of the normal nodes is an invariant set for the normal nodes; this is a *safety* condition. This condition is important in safety critical processes where the interval $[m[0], M[0]]$ is known to be safe. The agreement and safety conditions, when combined, imply a third condition on *validity*: the consensus quantity that the values of the normal nodes converge to must lie within the range of initial values of the normal nodes.

The validity condition is reasonable in applications where any value in the range of initial values of normal nodes is acceptable to select as the consensus value. For instance, consider a large sensor network where every sensor takes a measurement of its environment, captured as a real number. Suppose that at the time of measurement, all values taken by correct sensors fall within a range $[a, b]$, and that all sensors are required to come to an agreement on a common measurement value. If the range of measurements taken by the normal sensors is relatively small, it will likely be the case that reaching agreement on a value within that range will form a reasonable estimate of the measurements taken by all sensors. However, if a set of malicious nodes is capable of biasing the consensus value outside of this range, the error in the measurements could be arbitrarily large.

More generally, suppose the nodes are trying to distributively minimize $\sum h_i(\theta)$, where each of the h_i 's is a local convex function and θ is the optimization variable. If the initial value of each node i represents the value of θ that minimizes h_i , a convex combination of these initial values will represent an estimate of the optimal θ , within some bounded error. On the other hand, if an adversary is capable of biasing the consensus value arbitrarily, the resulting value of the objective function will also be arbitrarily far away from its minimum value. One can formulate similar motivating examples for the validity condition in other applications as well; for instance, a swarm of robots that are trying to flock should not be pulled in arbitrary directions by a malicious agent in the network.

III. CONSENSUS ALGORITHM

While there are various approaches to facilitate consensus, a class of *linear algorithms* have attracted significant interest in recent years [6], [38], due to their applicability in a variety of contexts. In such strategies, at time t , each node senses or receives information from its neighbors, and changes its value according to

$$x_i[t+1] = \sum_{j \in \mathcal{J}_i[t]} w_{ij}[t] x_j^i[t], \quad (1)$$

where $w_{ij}[t]$ is the weight assigned to node j 's value by node i at time-step t . The above strategy is the so-called *Linear Consensus Protocol (LCP)*.

Different conditions have been reported in the literature to ensure asymptotic consensus is reached [7], [13], [39]–[41]. In discrete time, it is common to assume that there exists a constant $\alpha \in \mathbb{R}$, $0 < \alpha < 1$ such that all of the following conditions hold:

- $w_{ij}[t] = 0$ whenever $j \notin \mathcal{J}_i[t]$, $i \in \mathcal{N}$, $t \in \mathbb{Z}_{\geq 0}$;
- $w_{ij}[t] \geq \alpha$, $\forall j \in \mathcal{J}_i[t]$, $i \in \mathcal{N}$, $t \in \mathbb{Z}_{\geq 0}$;
- $\sum_{j=1}^n w_{ij}[t] = 1$, $\forall i \in \mathcal{N}$, $t \in \mathbb{Z}_{\geq 0}$.

Given these conditions, a necessary and sufficient condition for reaching asymptotic consensus in time-invariant networks is that the digraph has a *rooted out-branching*, also called a *rooted directed spanning tree* [38]. The case of dynamic networks is not quite as straightforward. In this case, under the conditions stated above, a sufficient condition for reaching asymptotic consensus is that there exists a uniformly bounded sequence of contiguous time intervals such that the union of digraphs across each interval has a rooted out-branching [40]. Recently, a more general

condition referred to as the *infinite flow property* has been shown to be both necessary and sufficient for asymptotic consensus for a class of discrete-time stochastic models [42]. Finally, the lower bound on the weights is needed because there are examples of asymptotically vanishing weights in which consensus is not reached [43].

Given a fixed, bidirectional network topology, the selection of the optimal weights in (1) with respect to the speed of the consensus process can be done by solving a semidefinite program (SDP) [39]. However, this SDP is solved at design time with global knowledge of the network topology. A simple suboptimal choice of weights that requires only local information is to let $w_{ij}[t] = 1/(1 + d_i[t])$ for $j \in \mathcal{J}_i[t]$.

One problem with the linear update given in (1) is that it is not resilient to misbehaving nodes. In fact, it was shown in [13], [44] that a single ‘leader’ node can cause all agents to reach consensus on an arbitrary value of its choosing (potentially resulting in a dangerous situation in physical systems) simply by holding its value constant. Thus, by themselves, the dynamics given by (1) do not facilitate resilient asymptotic consensus for any of the fault models. We now describe a simple modification to the update rule, and then provide a comprehensive characterization of network topologies in which resilient asymptotic consensus is reached under such dynamics.

A. Description of W-MSR

At every time-step t , each normal node i obtains the values of other nodes in its neighborhood. At most F of node i 's neighbors may be misbehaving; however, node i is unsure of which neighbors may be compromised. To ensure that node i updates its value in a safe manner, we consider a protocol where each node removes the extreme values with respect to its own value. More specifically:

- 1) At each time-step t , each normal node i obtains the values of its neighbors, and forms a sorted list.
- 2) If there are less than F values strictly larger than its own value, $x_i[t]$, 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) Let $\mathcal{R}_i[t]$ denote the set of nodes whose values were removed by normal node i in step 2 at time-step t . Each normal node i applies the update

$$x_i[t + 1] = \sum_{j \in \mathcal{J}_i[t] \setminus \mathcal{R}_i[t]} w_{ij}[t] x_j^i[t], \quad (2)$$

where the weights $w_{ij}[t]$ satisfy the conditions stated above, but with $\mathcal{J}_i[t]$ replaced by $\mathcal{J}_i[t] \setminus \mathcal{R}_i[t]$.³

To accommodate the f -fraction local model, the parameter F in step 2 above is replaced by $F_i = \lfloor f d_i[t] \rfloor$. As a matter of terminology, we refer to the bound on the number (or fraction) of larger or smaller values that could be thrown away as the *parameter* of the algorithm. Above, the parameter of W-MSR with the F -local and F -total models is F , whereas the parameter with the f -fraction local model is f , and the meaning of the parameter will be clear from the context.

Observe that the set of nodes removed by normal node i , $\mathcal{R}_i[t]$, is possibly time-varying. Hence, even if the underlying network topology is fixed, the W-MSR algorithm effectively induces switching behavior, and can be viewed as the linear update of (1) with a specific rule for state-dependent switching (the rule given in step 2).

The above algorithm is extremely lightweight, and does not require any normal node to have any knowledge of the network topology or of the identities of non-neighbor nodes. Given these highly desirable properties, the question that we answer in this paper is: in what networks will the above algorithm facilitate resilient asymptotic consensus?

B. Use of Related Algorithms in Previous Work

As mentioned in the introduction, the use of similar algorithms that remove extreme values and then form an average from a subset of the remaining values have been studied for decades. In [29], functions that perform this type of operation are referred to as *approximation functions*, and both synchronous and asynchronous algorithms are studied that use these approximation functions in complete networks for resilience to F -total Byzantine faults. These approximation functions are used in the family of *Mean-Subsequence-Reduced (MSR)* algorithms [30]. There are a few key differences between the operations used in the W-MSR algorithm and the MSR algorithm of [30]. First, W-MSR does not always remove the largest and smallest F values as

³In this case, a simple choice for the weights is to let $w_{ij}[t] = 1/(1 + d_i[t] - |\mathcal{R}_i[t]|)$ for $j \in \mathcal{J}_i[t] \setminus \mathcal{R}_i[t]$.

in the MSR algorithm [30]. Instead, only the extreme values that are strictly larger or strictly smaller than the given node’s value are removed. Since the node’s own value may be one of the F extreme values, the MSR algorithm may throw away this useful (correct) information. Second, W-MSR uses all values retained after removing the extreme values. MSR, on the other hand, may select only a subsequence of the remaining values to use in the update. Finally, MSR averages the remaining values instead of allowing for weighted averages as in W-MSR.

MSR algorithms have also been used for Byzantine point convergence of mobile robots in complete networks [23]. Besides Byzantine faults, some works also consider other threat models [30]. However, few papers have addressed the convergence of MSR algorithms in less restrictive (non-complete) networks. Some exceptions include [31], [45], [46]. In [31], the authors studied *local convergence* (convergence of a subset of nodes) in undirected regular graphs⁴; the results are extended to asynchronous networks in [46] and global convergence of a class of undirected regular graphs, named *Partially Fully Connected Networks (PFCN)*, in [45]. More recently, [28] provides necessary and sufficient conditions on the network topology required for a special case of the MSR algorithm (which retains all of the values after removing the extreme ones) to achieve consensus in the presence of F -total Byzantine faults. In the following sections, we will develop a novel topological property and show that this property is essential for studying MSR (and more generally, W-MSR) algorithms in arbitrary networks for the broad class of adversarial models described in Section II.

Finally, it is worth noting the relationship between the W-MSR algorithm and robust consensus algorithms designed to withstand outliers [47], [48]. The problem of robust consensus to outliers does not assume a threat model, such as malicious or Byzantine nodes. Instead, some measurements may be statistical outliers (caused by noise) and the goal is to reach consensus on the measurements in a manner that reduces the error introduced by the outliers. In these works the nodes with outlier measurements are cooperative in the consensus process. Therefore, such techniques are not designed to work in the presence of misbehaving nodes. Furthermore, the W-MSR algorithm will also handle the case where the misbehaving nodes change their initial values, but behave normally otherwise.

⁴A regular graph is a graph where each vertex has the same number of neighbors.

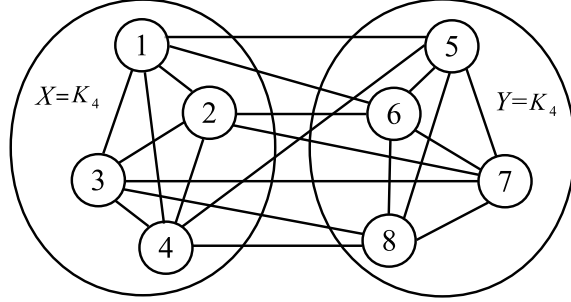


Fig. 1. Example of a 5-connected graph satisfying Prop. 1 with $F = 2$.

IV. ROBUST NETWORK TOPOLOGIES

Traditionally, network connectivity has been the key metric for studying robustness of distributed algorithms because it formalizes the notion of redundant information flow across the network through independent paths. Due to the fact that each independent path may include multiple intermediate nodes, network connectivity is well-suited for studying resilient distributed algorithms that assume such nonlocal information is available (for example, by explicitly relaying information across multiple hops in the network [5], by ‘inverting’ the dynamics on the network to recover the needed information [24], [25], or by resiliently encoding information along multiple paths [20]). However, when the nodes in the network use only local information (as in W-MSR), the following proposition suggests that connectivity is no longer a promising metric.

Proposition 1: For any $n, F \in \mathbb{Z}_{>0}$ with $F \leq \lfloor \frac{n}{2} \rfloor$, there exists a graph with connectivity $\kappa = \lfloor \frac{n}{2} \rfloor + F - 1$ in which W-MSR with parameter F does not ensure asymptotic consensus.

The proof of Proposition 1 can be found in the Appendix. Figure 1 illustrates an example of this kind of graph with $n = 8$, $F = 2$, and $\kappa = 5$. In this graph, there are two cliques (complete subgraphs), $X = K_4$ and $Y = K_4$, where K_n is the complete graph on n nodes. Each node in X has exactly $F = 2$ neighbors in Y , and vice versa. One can see that if the initial values of nodes in X and Y are $a \in \mathbb{R}$ and $b \in \mathbb{R}$, respectively, with $a \neq b$, then asymptotic consensus is not achieved whenever W-MSR is used with parameter F , even in the absence of misbehaving nodes. This is because each node views the values of its F neighbors from the opposing set as extreme, and removes all of these values from its list. The only remaining values for each node are from its own set, and thus no node ever changes its value.

The situation can be even worse in the more general case of digraphs. Examples of digraphs are illustrated in [27] that have minimum out-degree $n - 2$ and the underlying graph is complete, yet W-MSR still cannot guarantee resilient asymptotic consensus. Thus, even a relatively large connectivity (or minimum out-degree) in digraphs is not sufficient to guarantee consensus of the normal nodes, indicating the inadequacy of these traditional metrics to analyze the convergence properties of W-MSR. Taking a closer look at the graph in Fig. 1, we see that the reason for the failure of consensus is that no node has enough neighbors in the opposite set; this causes every node to throw away all useful information from outside of its set, and prevents consensus. What is needed is a metric that formalizes the notion of sufficient redundancy of information flow *directly* to at least one node in a subset. To capture this intuition, we develop a novel graph-theoretic property framed in terms of *reachable sets* and *network robustness* [2].

Definition 8 (r -reachable set): Given a digraph \mathcal{D} and a nonempty subset \mathcal{S} of nodes of \mathcal{D} , we say \mathcal{S} is an **r -reachable set** if $\exists i \in \mathcal{S}$ such that $|\mathcal{V}_i \setminus \mathcal{S}| \geq r$, where $r \in \mathbb{Z}_{\geq 0}$.

Definition 9 (p -fraction reachable set): Given a digraph \mathcal{D} and a nonempty subset \mathcal{S} of nodes of \mathcal{D} , we say \mathcal{S} is a **p -fraction reachable set** if $\exists i \in \mathcal{S}$ such that $|\mathcal{V}_i| > 0$ and $|\mathcal{V}_i \setminus \mathcal{S}| \geq p|\mathcal{V}_i|$, where $0 \leq p \leq 1$. If $|\mathcal{V}_i| = 0$ or $|\mathcal{V}_i \setminus \mathcal{S}| = 0$ for all $i \in \mathcal{S}$, then \mathcal{S} is 0-fraction reachable.

A set \mathcal{S} is r -reachable (or p -fraction reachable) if it contains a node that has at least r (or $\lceil pd_i \rceil$) neighbors outside of \mathcal{S} . The parameter r (or p) quantifies the redundancy of information flow from nodes outside of \mathcal{S} to *some* node inside \mathcal{S} . Intuitively, the r -reachability (or p -fraction reachability) property captures the idea that some node inside the set is influenced by a sufficiently large number of nodes from outside the set. The above reachability property pertains to a given set \mathcal{S} . The following definitions generalize this notion of redundancy to the entire network.

Definition 10 (r -robustness): A nonempty, nontrivial digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ on n nodes ($n \geq 2$) is **r -robust**, with $r \in \mathbb{Z}_{\geq 0}$, if for every pair of nonempty, disjoint subsets of \mathcal{V} , at least one of the subsets is r -reachable. By convention, if \mathcal{D} is empty or trivial ($n \leq 1$), then \mathcal{D} is 0-robust. The trivial graph is also 1-robust.⁵

Definition 11 (p -fraction robustness): A nonempty, nontrivial digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ on n nodes ($n \geq 2$) is **p -fraction robust**, with $0 \leq p \leq 1$, if for every pair of nonempty, disjoint subsets of

⁵The trivial graph is defined to be both 0-robust and 1-robust for consistency with properties shown to hold for larger digraphs in Section VII.

\mathcal{V} , at least one of the subsets is p -fraction reachable. If \mathcal{D} is empty or trivial ($n \leq 1$), then \mathcal{D} is 0-fraction robust.

Note that the notions of robustness and fraction robustness are similar to the concept of *vertex expanders*⁶ [49], [50]. However, the definition of vertex expanders only requires that the *whole set* has sufficient neighbors outside the set; for this reason, even a high expansion ratio may not guarantee that the set contains some node that by itself has enough neighbors outside the set. Thus, the concept of vertex expanders is not applicable to characterize the network topology required to succeed using the W-MSR algorithm.

The reason that pairs of nonempty, disjoint subsets of nodes are considered in the definition of r -robustness can be seen in the example of Fig. 1. If either X or Y were 3-reachable ($r = F + 1 = 3$), then at least one node would be sufficiently influenced by a node outside its set (because each node only removes up to $F = 2$ nodes that have values lower or higher than its own). This would drive it away from the values of its group, and thereby allow it to lead its group to the values of the other set.

However, if there are misbehaving nodes in the network, then the situation becomes more complex. For example, consider the network modeled by the graph in Fig. 2. One can verify that the graph is 3-robust by checking every possible pair of disjoint subsets, and confirming that at least one of them is 3-reachable. Consider the disjoint subsets X and Y shown in the figure, and note that both of them are 3-reachable – nodes 2 and 8 each have three neighbors outside of their respective sets. However, no other nodes in those two sets have more than two neighbors outside their own set, and thus nodes 2 and 8 are the only ones with access to sufficient information outside their own set. Suppose these two nodes 2 and 8 are malicious (or Byzantine) and the initial values of nodes in X and Y are a and b , respectively. Then, by stubbornly maintaining their initial values, nodes 2 and 8 are able to prevent consensus whenever the normal nodes use W-MSR with parameter $F = 2$. One way to remedy this is to require the whole network to be more robust. Another way is to introduce another form of information redundancy by specifying a minimum number of nodes that are sufficiently influenced from outside of their set. In order to capture this intuition, we define the following concept.

⁶A digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is an r vertex expander if for all $\mathcal{S} \subset \mathcal{V}$ of size at most $\lceil \frac{n}{2} \rceil$, the neighborhood $\mathcal{V}(\mathcal{S}) = \{j \in \mathcal{V} \setminus \mathcal{S} : \exists i \in \mathcal{S} \text{ s.t. } (j, i) \in \mathcal{E}\}$ is of size at least $r|\mathcal{S}|$.

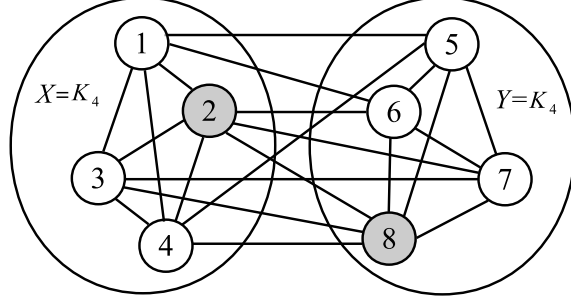


Fig. 2. A 3-robust graph in which sets X and Y are each 3-reachable. Nodes 2 and 8 are malicious (shown in grey).

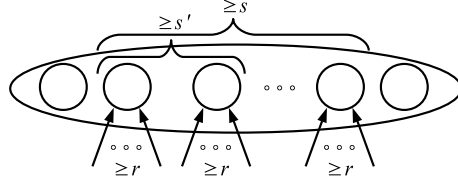


Fig. 3. Illustration of an (r, s) -reachable set of nodes.

Definition 12 ((r, s) -reachable set): Given a digraph \mathcal{D} and a nonempty subset of nodes \mathcal{S} , we say that \mathcal{S} is an (r, s) -**reachable set** if there are at least s nodes in \mathcal{S} , each of which has at least r neighbors outside of \mathcal{S} , where $r, s \in \mathbb{Z}_{\geq 0}$; i.e., given $\mathcal{X}_{\mathcal{S}}^r = \{i \in \mathcal{S} : |\mathcal{V}_i \setminus \mathcal{S}| \geq r\}$, then $|\mathcal{X}_{\mathcal{S}}^r| \geq s$.

An illustration of an (r, s) -reachable set of nodes is shown in Fig. 3. Observe that, in general, a set \mathcal{S} is (r, s') -reachable, for $s' \leq s$, whenever \mathcal{S} is (r, s) -reachable. At one extreme, whenever there are no nodes in \mathcal{S} with at least r neighbors outside of \mathcal{S} , then \mathcal{S} is only $(r, 0)$ -reachable. At the other extreme, \mathcal{S} can be at most $(r, |\mathcal{S}|)$ -reachable. Also note that r -reachability is equivalent to $(r, 1)$ -reachability. Hence, (r, s) -reachability strictly generalizes r -reachability, and better quantifies the number of nodes with redundant information flow from outside of their set. This additional specificity is useful for defining a more general notion of robustness.

Definition 13 ((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 $\mathcal{X}_{\mathcal{S}_k}^r = \{i \in \mathcal{S}_k : |\mathcal{V}_i \setminus \mathcal{S}_k| \geq r\}$ for $k \in \{1, 2\}$):

- (i) $|\mathcal{X}_{\mathcal{S}_1}^r| = |\mathcal{S}_1|$;

$$(ii) \quad |\mathcal{X}_{\mathcal{S}_2}^r| = |\mathcal{S}_2|;$$

$$(iii) \quad |\mathcal{X}_{\mathcal{S}_1}^r| + |\mathcal{X}_{\mathcal{S}_2}^r| \geq s.$$

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.⁷

The definition of (r, s) -robustness aims to capture the idea that “enough” nodes in every pair of nonempty, disjoint sets $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ have at least r neighbors outside of their respective sets. To quantify what is meant by “enough” nodes, it is necessary to take the maximal $s_{r,k}$ for which \mathcal{S}_k is $(r, s_{r,k})$ -reachable for $k \in \{1, 2\}$ (since \mathcal{S}_k is $(r, s'_{r,k})$ -reachable for $s'_{r,k} \leq s_{r,k}$). Since $s_{r,k} = |\mathcal{X}_{\mathcal{S}_k}^r|$, condition (i) or (ii) means that *all* nodes in \mathcal{S}_k have at least r neighbors outside of \mathcal{S}_k . Given a pair $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ such that $0 < |\mathcal{S}_1| < r$ and $\mathcal{S}_2 = \mathcal{V} \setminus \mathcal{S}_1$, there can be no more than $|\mathcal{S}_1|$ nodes with at least r neighbors outside of their set. Hence, conditions (i) and (ii) quantify the maximum number of nodes with at least r neighbors outside of their set for such pairs, and must therefore be “enough”. Alternatively, if there are at least s nodes with at least r neighbors outside of their respective sets in the union $\mathcal{S}_1 \cup \mathcal{S}_2$, then condition (iii) is satisfied. For such pairs $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$, the parameter⁸ $1 \leq s \leq n$ quantifies what is meant by “enough” nodes.

In the next section, we will show that these concepts we have proposed above are, in fact, *the key properties* needed to characterize the performance of the class of local filtering algorithms given by W-MSR. That is, sufficiently robust digraphs guarantee resilient consensus.

V. RESILIENT CONSENSUS ANALYSIS

We start with the following result showing that W-MSR always satisfies the safety condition for resilient asymptotic consensus. Recall that $M[t]$ and $m[t]$ are the maximum and minimum values of the *normal* nodes at time-step t , respectively.

Lemma 1: Suppose each normal node updates its value according to the W-MSR algorithm with parameter F under the F -total or F -local Byzantine (or malicious) model, or with parameter

⁷The trivial graph is defined to be both (0,1)-robust and (1,1)-robust for consistency with properties shown to hold for larger digraphs in Section VII.

⁸Note that $s = 0$ is *not* allowed in (r, s) -robustness because in that case any digraph on $n \geq 2$ nodes satisfies the definition for any $r \in \mathbb{Z}_{\geq 0}$, which subverts the interpretation of the parameter r . At the other extreme, the maximal meaningful value of s is $s = n$ since condition (iii) can never be satisfied with $s > n$.

f under the f -fraction local Byzantine (or malicious) model. Then, for each node $i \in \mathcal{N}$, $x_i[t+1] \in [m[t], M[t]]$, regardless of the network topology.

Proof: The proof is straightforward and follows directly from the definitions and the facts that the values in $\mathcal{J}_i[t] \setminus \mathcal{R}_i[t]$ used in the W-MSR update rule lie in the interval $[m[t], M[t]]$ and that the update rule in (2) is a convex combination of these values. ■

Having guaranteed the safety condition, we now provide a characterization of networks where the agreement condition (and thus, the validity condition) will be satisfied for each of the threat models introduced in Section II.

A. F -Total Malicious Model

The following result is one of the major contributions of this paper and provides, for the first time, a *necessary and sufficient* condition for the W-MSR algorithm to succeed under the F -total malicious model.

Theorem 1: Consider a time-invariant network modeled by a digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ where each normal node updates its value according to the W-MSR algorithm with parameter F . Under the F -total malicious model, resilient asymptotic consensus is achieved if and only if the network topology is $(F+1, F+1)$ -robust.

Proof: (Necessity) If \mathcal{D} is not $(F+1, F+1)$ -robust, then there are nonempty, disjoint $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ such that none of the conditions (i) – (iii) hold. 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 $|\mathcal{X}_{\mathcal{S}_1}^{F+1}| + |\mathcal{X}_{\mathcal{S}_2}^{F+1}| \leq F$, suppose all nodes in $\mathcal{X}_{\mathcal{S}_1}^{F+1}$ and $\mathcal{X}_{\mathcal{S}_2}^{F+1}$ are malicious and keep their values constant. With this assignment of adversaries, there is still at least one normal node in both \mathcal{S}_1 and \mathcal{S}_2 since $|\mathcal{X}_{\mathcal{S}_1}^{F+1}| < |\mathcal{S}_1|$ and $|\mathcal{X}_{\mathcal{S}_2}^{F+1}| < |\mathcal{S}_2|$, respectively. Since these normal nodes remove the F or less values of in-neighbors outside of their respective sets, no consensus among normal nodes is reached.

(Sufficiency) Recall that \mathcal{N} is the set of normal nodes, and define $N = |\mathcal{N}|$. Furthermore, define $M[t]$ and $m[t]$ to be the maximum and minimum values of the normal nodes at time-step t , respectively. We know from Lemma 1 that both $M[t]$ and $m[t]$ are monotone and bounded functions of t and thus each of them has some limit, denoted by A_M and A_m , respectively. Note that if $A_M = A_m$, the normal nodes will reach consensus. We will now prove by contradiction that this must be the case.

Suppose that $A_M \neq A_m$ (note that $A_M > A_m$ by definition). We can then define some constant $\epsilon_0 > 0$ such that $A_M - \epsilon_0 > A_m + \epsilon_0$. At any time-step t and for any positive real number ϵ_i , let $\mathcal{X}_M(t, \epsilon_i) = \{i \in \mathcal{V} : x_i[t] > A_M - \epsilon_i\}$, which includes all normal and malicious nodes that have values larger than $A_M - \epsilon_i$, and let $\mathcal{X}_m(t, \epsilon_i) = \{i \in \mathcal{V} : x_i[t] < A_m + \epsilon_i\}$, which includes all normal and malicious nodes that have values smaller than $A_m + \epsilon_i$. Note that $\mathcal{X}_M(t, \epsilon_0)$ and $\mathcal{X}_m(t, \epsilon_0)$ are disjoint, by the definition of ϵ_0 .

Fix $\epsilon < \frac{\alpha^N}{1-\alpha^N}\epsilon_0$, which satisfies $\epsilon_0 > \epsilon > 0$. Let t_ϵ be such that $M[t] < A_M + \epsilon$ and $m[t] > A_m - \epsilon, \forall t \geq t_\epsilon$ (we know that such a t_ϵ exists by the definition of convergence). Consider the nonempty and disjoint sets $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ and $\mathcal{X}_m(t_\epsilon, \epsilon_0)$. Since the network is $(F + 1, F + 1)$ -robust and there are no more than F malicious nodes in the network (F -total model), there is a normal node in the union that has at least $F + 1$ neighbors outside of its set. Without loss of generality, suppose normal node $i \in \mathcal{X}_M(t_\epsilon, \epsilon_0) \cap \mathcal{N}$ has at least $F + 1$ neighbors outside of $\mathcal{X}_M(t_\epsilon, \epsilon_0)$. By definition, these neighbors have values at most equal to $A_M - \epsilon_0$, and at least one of these values will be used by node i (since node i removes at most F values lower than its own value). Note that at each time-step, every normal node's value is a convex combination of its own value and the values it uses from its neighbors, and each coefficient in the combination is lower bounded by α . Since the largest value that node i will use at time-step t_ϵ is $M[t_\epsilon]$, placing the largest possible weight on $M[t_\epsilon]$ produces

$$\begin{aligned} x_i[t_\epsilon + 1] &\leq (1 - \alpha)M[t_\epsilon] + \alpha(A_M - \epsilon_0) \\ &\leq (1 - \alpha)(A_M + \epsilon) + \alpha(A_M - \epsilon_0) \\ &\leq A_M - \alpha\epsilon_0 + (1 - \alpha)\epsilon. \end{aligned}$$

Note that this upper bound also applies to the updated value of any normal node that is not in $\mathcal{X}_M(t_\epsilon, \epsilon_0)$, because such a node will use its own value in its update. Similarly, if $j \in \mathcal{X}_m(t_\epsilon, \epsilon_0) \cap \mathcal{N}$ has at least $F + 1$ neighbors outside of $\mathcal{X}_m(t_\epsilon, \epsilon_0)$, then $x_j[t_\epsilon + 1] \geq A_m + \alpha\epsilon_0 - (1 - \alpha)\epsilon$. Again, any normal node that is not in $\mathcal{X}_m(t_\epsilon, \epsilon_0)$ will have the same lower bound.

Define $\epsilon_1 = \alpha\epsilon_0 - (1 - \alpha)\epsilon$, which satisfies $0 < \epsilon < \epsilon_1 < \epsilon_0$. Consider the sets $\mathcal{X}_M(t_\epsilon + 1, \epsilon_1)$ and $\mathcal{X}_m(t_\epsilon + 1, \epsilon_1)$. Since at least one of the normal nodes in $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ decreases at least to $A_M - \epsilon_1$ (or below), or one of the nodes in $\mathcal{X}_m(t_\epsilon, \epsilon_0)$ increases at least to $A_m + \epsilon_1$ (or above), it must be that either $|\mathcal{X}_M(t_\epsilon + 1, \epsilon_1) \cap \mathcal{N}| < |\mathcal{X}_M(t_\epsilon, \epsilon_0) \cap \mathcal{N}|$ or $|\mathcal{X}_m(t_\epsilon + 1, \epsilon_1) \cap \mathcal{N}| < |\mathcal{X}_m(t_\epsilon, \epsilon_0) \cap \mathcal{N}|$, or both. Since $\epsilon_1 < \epsilon_0$, $\mathcal{X}_M(t_\epsilon + 1, \epsilon_1)$ and $\mathcal{X}_m(t_\epsilon + 1, \epsilon_1)$ are still disjoint. For $j \geq 1$, define ϵ_j

recursively as $\epsilon_j = \alpha\epsilon_{j-1} - (1 - \alpha)\epsilon$, and observe that $\epsilon_j < \epsilon_{j-1}$. As long as there are still normal nodes in both $\mathcal{X}_M(t_\epsilon + j, \epsilon_j)$ and $\mathcal{X}_m(t_\epsilon + j, \epsilon_j)$, we can repeat the above analysis for time-steps $t_\epsilon + j$. Furthermore, at time-step $t_\epsilon + j$, either $|\mathcal{X}_M(t_\epsilon + j, \epsilon_j) \cap \mathcal{N}| < |\mathcal{X}_M(t_\epsilon + j - 1, \epsilon_{j-1}) \cap \mathcal{N}|$ or $|\mathcal{X}_m(t_\epsilon + j, \epsilon_j) \cap \mathcal{N}| < |\mathcal{X}_m(t_\epsilon + j - 1, \epsilon_{j-1}) \cap \mathcal{N}|$, or both. Since $|\mathcal{X}_M(t_\epsilon, \epsilon_0) \cap \mathcal{N}| + |\mathcal{X}_m(t_\epsilon, \epsilon_0) \cap \mathcal{N}| \leq N$, there must be some time-step $t_\epsilon + T$ (where $T \leq N$) where either $\mathcal{X}_M(t_\epsilon + T, \epsilon_T) \cap \mathcal{N}$ or $\mathcal{X}_m(t_\epsilon + T, \epsilon_T) \cap \mathcal{N}$ is empty. In the former case, all normal nodes in the network at time-step $t_\epsilon + T$ have value at most $A_M - \epsilon_T$, and in the latter case all normal nodes in the network at time-step $t_\epsilon + T$ have value no less than $A_m + \epsilon_T$. We will show that $\epsilon_T > 0$, which will contradict the fact that the largest value monotonically converges to A_M (in the former case) or that the smallest value monotonically converges to A_m (in the latter case). To do this, note that

$$\begin{aligned}
\epsilon_T &= \alpha\epsilon_{T-1} - (1 - \alpha)\epsilon \\
&= \alpha^2\epsilon_{T-2} - \alpha(1 - \alpha)\epsilon - (1 - \alpha)\epsilon \\
&\quad \vdots \\
&= \alpha^T\epsilon_0 - (1 - \alpha)(1 + \alpha + \dots + \alpha^{T-1})\epsilon \\
&= \alpha^T\epsilon_0 - (1 - \alpha^T)\epsilon \\
&\geq \alpha^N\epsilon_0 - (1 - \alpha^N)\epsilon.
\end{aligned}$$

Since $\epsilon < \frac{\alpha^N}{1 - \alpha^N}\epsilon_0$, we obtain $\epsilon_T > 0$, providing the desired contradiction. It must thus be the case that $\epsilon_0 = 0$, proving that $A_M = A_m$. ■

The above result establishes the notion of network robustness introduced in Definition 13 as the appropriate metric for reasoning about purely local distributed algorithms, supplanting the traditional metric of connectivity. We will discuss the relationship between connectivity and robustness in further detail later in the paper.

When the network is time-varying, one can state the following corollary of the above theorem. The proof is given in the Appendix.

Corollary 1: Consider a time-varying network modeled by a digraph $\mathcal{D}[t] = (\mathcal{V}, \mathcal{E}[t])$ where each normal node updates its value according to the W-MSR algorithm with parameter F . Let $\{t_k\}$ denote the set of time-steps in which $\mathcal{D}[t]$ is $(F + 1, F + 1)$ -robust. Then, under the F -total malicious model, resilient asymptotic consensus is achieved if $|\{t_k\}| = \infty$ and $|t_{k+1} - t_k| \leq c$, $\forall k$, where $c \in \mathbb{Z}_{>0}$.

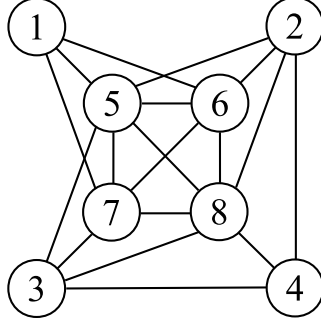


Fig. 4. A 3-robust graph that is *not* (3,2)-robust.

To illustrate these results consider the graphs in Figs. 1 and 4. These graphs can withstand the compromise of $F = 1$ malicious node in the network using the W-MSR algorithm with parameter $F = 1$ (each graph is (2,2)-robust but not (3,3)-robust). This is not to say that it is impossible for the normal nodes to reach consensus if there are, for example, two nodes that are compromised. Instead, these results say that there are two *specific* nodes that can be compromised by an adversary to prevent consensus (e.g., nodes 5 and 6 in Fig. 4).

B. F -Local and f -Fraction Local Malicious Models

The previous result fully characterizes those network topologies that are able to handle F -total malicious adversaries. In order to capture the case when the total number of adversaries is quite large (e.g., in large-scale networks), we now consider the F -local and f -fraction local malicious models. Due to the fact that there is no meaningful upper bound on the total number of adversaries under these models, we cannot rely on a ‘sufficiently large’ number of nodes in each set having $F + 1$ neighbors outside. Instead, we must return to the original definition of an r -robust network and increase the requirements on the number of external neighbors for a node in one out of any pair of disjoint sets.

Theorem 2: Consider a time-invariant network modeled by a digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ where each normal node updates its value according to the W-MSR algorithm with parameter F . Under the F -local malicious model, resilient asymptotic consensus is achieved if the topology of the network is $(2F + 1)$ -robust. Furthermore, a necessary condition is for the topology of the network to be $(F + 1)$ -robust.

Proof: (Necessity) If the network is not $(F + 1)$ -robust, there exist two disjoint subsets of nodes that are not $(F + 1)$ -reachable, i.e., each node in these two sets would have at most F neighbors outside the set. If we assign the maximum and minimum values in the network to these two sets, respectively, the nodes in these sets would never use any values from outside their own sets. Thus, their values would remain unchanged, and consensus will not be reached.

(Sufficiency) The proof of sufficiency is similar to the proof of Theorem 1. Note that when considering the nonempty, disjoint sets $\mathcal{X}_M(t_\epsilon, \epsilon_0) \cap \mathcal{N}$ and $\mathcal{X}_m(t_\epsilon, \epsilon_0) \cap \mathcal{N}$ defined in the proof of Theorem 1, at least one of these two sets must be $(2F + 1)$ -reachable due to the assumption of $(2F + 1)$ -robustness of the network. Thus, at least one of these two sets contains some normal node which will use at least one of its normal neighbors' values from outside. ■

Corollary 2: Consider a time-varying network modeled by a digraph $\mathcal{D}[t] = (\mathcal{V}, \mathcal{E}[t])$ where each normal node updates its value according to the W-MSR algorithm with parameter F . Let $\{t_k\}$ denote the set of time-steps in which $\mathcal{D}[t]$ is $(2F + 1)$ -robust. Then, under the F -local malicious model, resilient asymptotic consensus is achieved if $|\{t_k\}| = \infty$ and $|t_{k+1} - t_k| \leq c$, $\forall k$, where $c \in \mathbb{Z}_{>0}$.

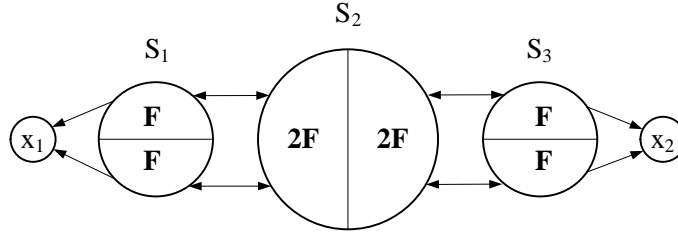


Fig. 5. Illustration of Proposition 2

Although the sufficient and necessary conditions in Theorem 2 do not coincide, the following result shows that the sufficient condition in the theorem is *sharp*.

Proposition 2: For every $F \in \mathbb{Z}_{>0}$, there exists a $2F$ -robust network that fails to reach consensus using the W-MSR algorithm with parameter F .

Proof: We will prove the result by giving a construction of such a graph, visualized in Figure 5. In Figure 5, S_1 , S_2 and S_3 are all complete components with $|S_1| = |S_3| = 2F$ and $|S_2| = 4F$. Each node in S_1 connects to $2F$ nodes of S_2 and each node in S_3 connects to the other $2F$ nodes of S_2 , and all of these connections are undirected. Node x_1 has incoming edges

from all nodes in \mathcal{S}_1 and similarly node x_2 has incoming edges from all nodes in \mathcal{S}_3 . This is an example of a graph that arises from the construction that we derive in Section VII-A, where we show that such a graph will be $2F$ robust. We choose F nodes of \mathcal{S}_1 and also F nodes of \mathcal{S}_3 to be malicious; note that this constitutes an F -local set of malicious nodes. Then we assign node x_1 with initial value m , node x_2 with initial value M and the other normal nodes with initial values c , such that $m < c < M$. Malicious nodes in \mathcal{S}_1 and \mathcal{S}_3 will keep their values unchanged at m and M , respectively. We can see that, by using the W-MSR algorithm, the values of nodes x_1 and x_2 will never change and thus consensus cannot be reached, completing the proof. ■

To illustrate these results, consider the 3-robust graph of Fig. 4. Recall that this graph cannot generally sustain 2 malicious nodes as specified by the 2-total model; it is not (3,3)-robust. However, under the 1-local model, it can sustain two malicious nodes if the *right* nodes are compromised. For example, nodes 1 and 4 may be compromised under the 1-local model and the normal nodes will still reach consensus. This example illustrates the advantage of the F -local model, where there is no concern about global assumptions. If a digraph is $(2F + 1)$ -robust, then up to F nodes may be compromised in any node's neighborhood, possibly resulting in more than F malicious nodes in the network (as in the previous example).

We now extend the discussion to the f -fraction local malicious model.

Theorem 3: Consider a time-invariant network modeled by a digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ where each normal node updates its value according to the W-MSR algorithm with parameter f . Under the f -fraction local malicious model, resilient asymptotic consensus is achieved if the topology of the network is p -fraction robust, where $2f < p \leq 1$. Furthermore, a necessary condition is for the topology of the network to be p' -fraction robust, where $p' > f$.

Proof: The proof is similar to the proof of Theorem 2. For the proof of sufficiency, note that under the f -fraction local model, each normal node will disregard at most $2 \times \lfloor fd_i \rfloor$ values from its neighborhood at each time-step. Thus, if the network is p -fraction robust, where $2f < p \leq 1$, at least one of these two sets $\mathcal{X}_M(t_\epsilon, \epsilon_0) \cap \mathcal{N}$ and $\mathcal{X}_m(t_\epsilon, \epsilon_0) \cap \mathcal{N}$ will adopt some normal node's value from outside. ■

Corollary 3: Consider a time-varying network modeled by a digraph $\mathcal{D}[t] = (\mathcal{V}, \mathcal{E}[t])$ where each normal node updates its value according to the W-MSR algorithm with parameter f . Let $\{t_k\}$ denote the set of time-steps in which $\mathcal{D}[t]$ is p -fraction robust, where $2f < p \leq 1$. Then, under the f -fraction local malicious model, resilient asymptotic consensus is achieved if $|\{t_k\}| = \infty$

and $|t_{k+1} - t_k| \leq c, \forall k$, where $c \in \mathbb{Z}_{>0}$.

C. F -Total, F -Local and f -Fraction Local Byzantine Models

Our above results have focused on the case of malicious (but not Byzantine) adversaries. The recent paper [28] investigates a similar algorithm in the context of F -total Byzantine adversaries, and provides necessary and sufficient conditions for the algorithm to succeed. While their proof techniques are different, the main result can be captured neatly by the notion of robustness as follows.

Definition 14: For a network $\mathcal{D} = (\mathcal{V}, \mathcal{E})$, define the *normal network* of \mathcal{D} , denoted by $\mathcal{D}_{\mathcal{N}}$, as the network induced by the normal nodes, i.e., $\mathcal{D}_{\mathcal{N}} = (\mathcal{N}, \mathcal{E}_{\mathcal{N}})$, where $\mathcal{E}_{\mathcal{N}}$ is the set of directed edges among the normal nodes.

Theorem 4 ([28]): Consider a time-invariant network modeled by a digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ where each normal node updates its value according to the W-MSR algorithm with parameter F . Under the F -total Byzantine model, resilient asymptotic consensus is achieved if and only if the topology of the *normal network* is $(F + 1)$ -robust.

Proof: To prove sufficiency, besides the method used in [28], [51], we can also use the approach proposed in the proof of Theorem 1. Note that when the original network is $(2F + 1)$ -robust, the normal network will be $(F + 1)$ -robust.

To prove necessity, if the normal network is not $(F + 1)$ -robust, we can assign the two disjoint sets that are not $(F + 1)$ -reachable the maximum and minimum values, respectively. Since the Byzantine nodes can send different values to different neighbors, suppose they send the maximum and minimum values to the maximum and minimum sets, respectively. Then, nodes in these two sets never use any values from outside their own sets and consensus is not reached. ■

The following results are straightforward extensions of the above result from [28] to the local models and time-varying networks.

Corollary 4: Consider a time-invariant network modeled by a digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ where each normal node updates its value according to the W-MSR algorithm with parameter F (or parameter f for the f -fraction local model). Under the F -local Byzantine model, resilient asymptotic consensus is achieved if and only if the topology of the normal network is $(F + 1)$ -robust. Under the f -fraction local Byzantine model, resilient asymptotic consensus is achieved if the

normal network is p -fraction robust, where $p > f$, and a necessary condition is for the normal network to be p' -fraction robust, where $p' \geq f$.

Proof: The proof is similar to the proof of Theorem 4. For the proof of necessity, note that the choice of Byzantine nodes should satisfy the F -local and f -fraction local properties, respectively. Further note that the only difference between the sufficient and necessary conditions for the f -fraction local model is $p = f$. When the normal network is f -fraction robust, we can choose two sets which are at most f -fraction reachable and each node i in these two sets has at most $\lceil fd_i \rceil$ neighbors outside. For certain choice of initial values (i.e., these two sets have the maximum and minimum initial values, respectively), consensus can be reached if $fd_i \notin \mathbb{Z}_{\geq 1}$ and cannot be reached if $fd_i \in \mathbb{Z}_{\geq 1}$. ■

Corollary 5: Consider a time-varying network modeled by a digraph $\mathcal{D}[t] = (\mathcal{V}, \mathcal{E}[t])$ where each normal node updates its value according to the W-MSR algorithm with parameter F (or parameter f for the f -fraction local model). Let $\{t_k\}$ denote the set of time steps in which $\mathcal{D}[t]$ is either (i) $(2F + 1)$ -robust, or (ii) p -fraction robust, where $2f < p \leq 1$. Then, under (i) the F -total or F -local Byzantine model, or (ii) the f -fraction local Byzantine model, respectively, resilient asymptotic consensus is achieved if $|\{t_k\}| = \infty$ and $|t_{k+1} - t_k| \leq c, \forall k$, where $c \in \mathbb{Z}_{>0}$.

VI. SIMULATION RESULTS

This section presents a numerical example to illustrate our results. In this example, the network is given by the (2,2)-robust graph shown in Fig. 6, in which the node set is $\mathcal{V} = \{1, 2, \dots, 14\}$ and node $i \in \mathcal{V}$ has initial value $x_i[0]$ shown in the circle representing the node. To verify that this graph is (2,2)-robust one must exhaustively check every nonempty, disjoint pair of subsets of nodes to make sure that either every node in one of the sets has at least 2 neighbors outside of its set, or that there are at least 2 nodes in the union of the subsets that have 2 or more neighbors outside of their respective set. For example, the pair of sets $\{6\}$ and $\mathcal{V} \setminus \{6\}$ passes this test since each node in the first set (just node 6) has at least 2 neighbors outside of its set (in this case just node 6's neighbors). As another example, the pair of sets $\{1, 2, 11, 12\}$ and $\{5, 6\}$ passes since node 11 and node 5 each have 2 or more neighbors outside of their respective sets.

Since the network is (2,2)-robust, Theorem 1 indicates it can sustain a single malicious node in the network under the 1-total model. Suppose that the node with the largest degree, node 14, is compromised and turns malicious. The normal nodes use either the LCP given in (1) or

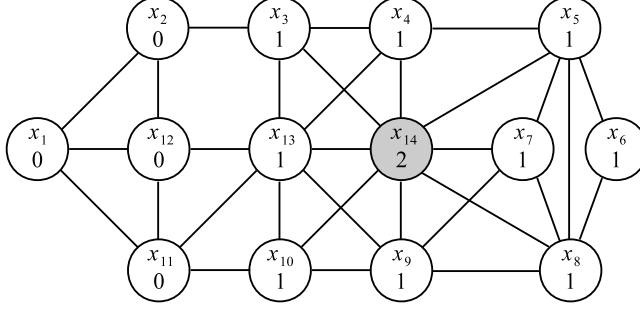


Fig. 6. (2,2)-Robust network topology.

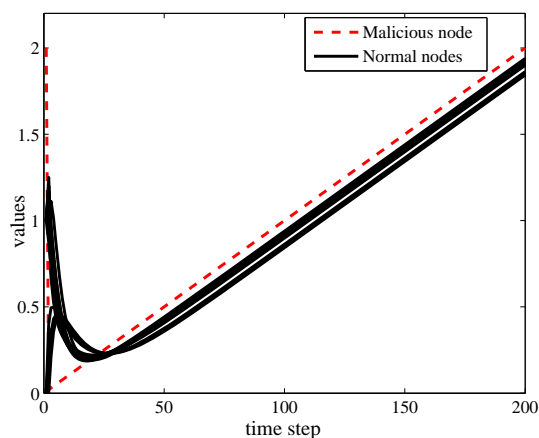
W-MSR for their update rule. Each normal node $i \in \mathcal{N}$ uses the weights $w_{ij}[t] = |\mathcal{J}_i[t]|^{-1}$ for each $j \in \mathcal{J}_i[t]$ with LCP and $w_{ij}[t] = (|\mathcal{J}_i[t] \setminus \mathcal{R}_i[t]|)^{-1}$ for each $j \in \mathcal{J}_i[t] \setminus \mathcal{R}_i[t]$ with W-MSR. The malicious node's objective is to prevent the normal nodes from reaching consensus and to drive the normal node values outside of the range of their initial values.

The results for the time-invariant network of Fig. 6 are shown in Fig. 7. It is clear in Fig. 7(a) that the malicious node is able to drive the values of the normal nodes outside of the range of initial values and prevent consensus whenever LCP is used. On the other hand, the malicious node is unable to achieve its goal whenever W-MSR is used. Note that although consensus can be reached, the malicious node still has the potential to drive the consensus process to any value in the interval $[0, 1]$ by choosing the desired value as its initial value and remaining constant. However, this is allowed with resilient asymptotic consensus (because the consensus value is within the range of the initial values held by normal nodes).

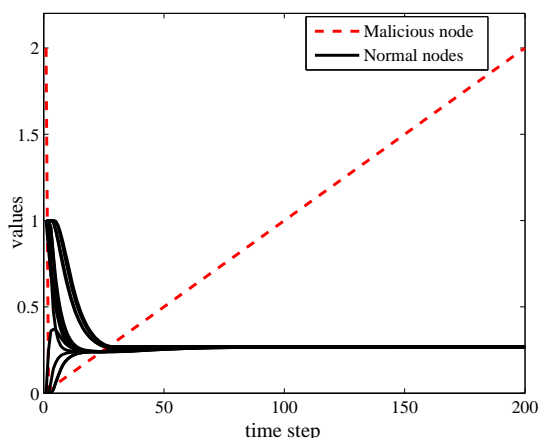
Finally, we illustrate the time-varying network result for the 1-total malicious model by removing approximately half of the directed edges in 9 out of every 10 consecutive time-steps. To do this, we check whether the time-step is equal to 0 modulo 10. If it is not, then we model directed edge removal by a Bernoulli process with parameter $p = 0.5$, so that approximately half of the directed edges are removed in these time-steps. The results are illustrated in Fig. 8(b), and show that only the speed of convergence is affected when using W-MSR.

VII. REVISITING NETWORK ROBUSTNESS: CONSTRUCTION AND PROPERTIES

Having established network robustness as the key metric for characterizing the efficacy of the W-MSR algorithm, we now provide more insight into robust networks. First, we provide a

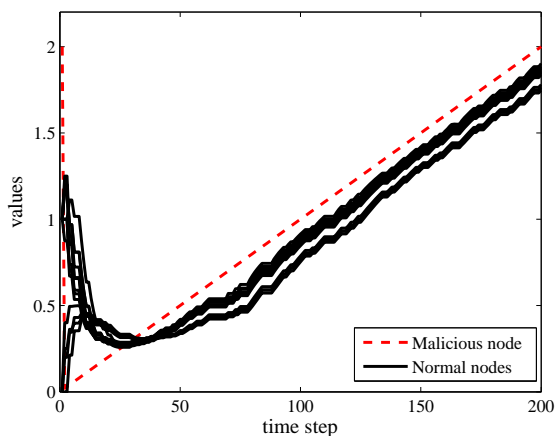


(a) LCP.

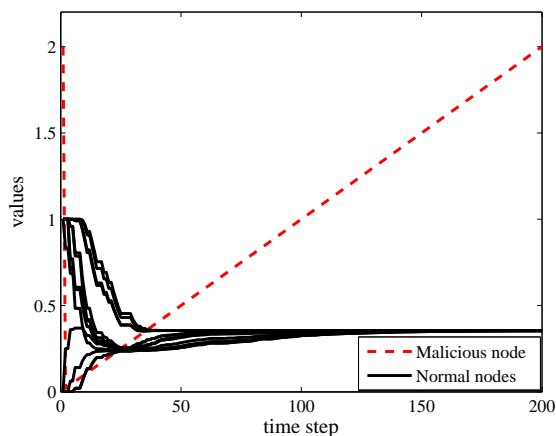


(b) W-MSR.

Fig. 7. Malicious node attempts to prevent the normal nodes from reaching consensus and drive their values away from the convex hull containing their initial values. The malicious node succeeds whenever LCP is used, but fails whenever W-MSR is used.



(a) LCP.



(b) W-MSR.

Fig. 8. The malicious node has the same objective as before, but now the network is time-varying with only one time-step out of every ten guaranteed to be $(2,2)$ -robust.

method for constructing robust digraphs, and show that scale-free networks constructed using the preferential-attachment model are robust. We then explore more properties of robust digraphs.

A. Construction of Robust Digraphs

Note that robustness requires checking every possible nonempty disjoint pair of subsets of nodes in the digraph for certain conditions. Currently, we do not have a computationally efficient method to check whether these properties hold in arbitrary digraphs. However, in [2] it is shown that the common *preferential-attachment* model for complex networks (e.g., [52]) produces r -robust graphs, provided that a sufficient number of links are added to new nodes as they are attached. Here we show that preferential attachment also leads to (r, s) -robust graphs.

Theorem 5: Let $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ be an (r, s) -robust digraph (with $s \geq 1$). Then the digraph $\mathcal{D}' = (\mathcal{V} \cup \{v_{\text{new}}\}, \mathcal{E} \cup \mathcal{E}_{\text{new}})$, where v_{new} is a new vertex added to \mathcal{D} and \mathcal{E}_{new} is the directed edge set related to v_{new} , is (r, s) -robust if $d_{v_{\text{new}}} \geq r + s - 1$.

Proof: For a pair of nonempty, disjoint sets \mathcal{S}_1 and \mathcal{S}_2 , there are three cases to check: $v_{\text{new}} \notin \mathcal{S}_i$, $\{v_{\text{new}}\} = \mathcal{S}_i$ and $v_{\text{new}} \in \mathcal{S}_i$, for some $i \in \{1, 2\}$. In the first case, since \mathcal{D} is (r, s) -robust, the conditions in Definition 13 must hold. In the second case, $\mathcal{X}_{\mathcal{S}_i}^r = \mathcal{S}_i$, and we are done. In the third case, suppose, without loss of generality, $\mathcal{S}_2 = \mathcal{S}'_2 \cup \{v_{\text{new}}\}$. Since \mathcal{D} is (r, s) -robust, at least one of the following conditions hold: $|\mathcal{X}_{\mathcal{S}_1}^r| + |\mathcal{X}_{\mathcal{S}'_2}^r| \geq s$, $|\mathcal{X}_{\mathcal{S}_1}^r| = |\mathcal{S}_1|$, or $|\mathcal{X}_{\mathcal{S}'_2}^r| = |\mathcal{S}'_2|$. If either of the first two hold, then the corresponding conditions hold for the pair $\mathcal{S}_1, \mathcal{S}_2$ in \mathcal{D}' . So assume only $|\mathcal{X}_{\mathcal{S}'_2}^r| = |\mathcal{S}'_2|$ holds. Then, the negation of the first condition $|\mathcal{X}_{\mathcal{S}_1}^r| + |\mathcal{X}_{\mathcal{S}'_2}^r| \geq s$ implies $|\mathcal{X}_{\mathcal{S}'_2}^r| = |\mathcal{S}'_2| < s$. Hence, $|\mathcal{V}_{v_{\text{new}}} \setminus \mathcal{S}_2| \geq r$, and $|\mathcal{X}_{\mathcal{S}_2}^r| = |\mathcal{S}_2|$, completing the proof. ■

The above result indicates that to construct an (r, s) -robust digraph with n nodes (where $n > r$), we can start with an (r, s) -robust digraph with relatively smaller order (such as a complete graph), and continually add new nodes with incoming edges from at least $r + s - 1$ nodes in the existing digraph. Note that this method does not specify *which* existing nodes should be chosen. The preferential-attachment model corresponds to the case when the nodes are selected with a probability proportional to the number of edges that they already have. This leads to the formation of so-called *scale-free* networks [52], and is cited as a plausible mechanism for the formation of many real-world complex networks. Theorem 5 indicates that a certain class of scale-free networks is resilient to the threat models studied in this paper (provided the number of edges added in each round is sufficiently large when the network is forming).

For example, Fig. 9 illustrates a $(3, 2)$ -robust graph constructed using preferential attachment by starting with the complete graph on 5 nodes K_5 – which is also $(3, 3)$ -robust and is the only

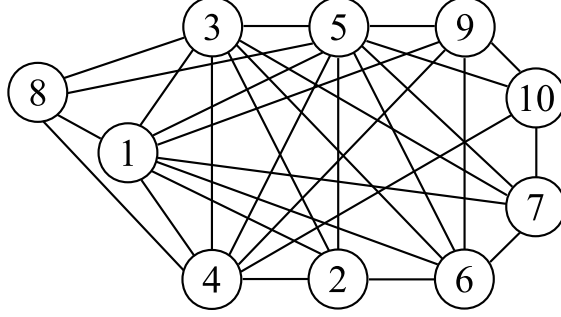


Fig. 9. A $(3, 2)$ -robust graph constructed from K_5 using preferential attachment.

$(3, 2)$ -robust digraph on 5 nodes (c.f., Lemma 4 in the sequel)) – and by adding 4 new edges to each new node in each step. Note that this graph is also 4-robust, which could *not* be predicted from Theorem 5 since K_5 is not 4-robust. Therefore, it is possible (but not guaranteed) to end up with a *more* robust digraph than the initial one using the growth model from Theorem 5.

B. Properties of Robust Networks

In this subsection, we begin with the important observation that $(r, 1)$ -robustness is equivalent to r -robustness. This holds because conditions (i) – (iii) in Definition 13 for $(r, 1)$ -robustness collapse to the condition that at least one of \mathcal{S}_1 and \mathcal{S}_2 is r -reachable. We next establish an inheritance property of (r, s) -robust digraphs. Note that all the proofs of the results in this subsection can be found in the Appendix.

Lemma 2: Every (r, s) -robust digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is also (r', s') -robust when $0 \leq r' \leq r$, $1 \leq s' \leq s$.

It follows from Lemma 2 that a digraph is r -robust whenever it is (r, s) -robust. The converse, however, is not true. Consider the graph in Fig. 4. This graph is 3-robust, but is not $(3, 2)$ -robust. For example, let $\mathcal{S}_1 = \{1, 3, 5, 6, 7\}$ and $\mathcal{S}_2 = \{2, 4\}$. Only node 2 has at least 3 nodes outside of its set, so all of the conditions (i) – (iii) fail. Therefore, (r, s) -robustness is a strict generalization of r -robustness.

The following result formalizes the intuition that adding links to a robust network can never reduce the robustness of the network.

Lemma 3 (Monotonicity): Suppose $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is an (r, s) -robust spanning subdigraph of $\mathcal{D}' = (\mathcal{V}, \mathcal{E}')$, where $\mathcal{E}' = \mathcal{E} \cup \mathcal{E}''$ and $|\mathcal{E}''| \geq 0$. Then \mathcal{D}' is (r, s) -robust.

Next, we look at the maximum amount of robustness one can expect from a network with n nodes. As expected, the complete digraph K_n is the most robust topology on n nodes.

Lemma 4 (Maximum robustness): No digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ on n nodes is $(\lceil n/2 \rceil + 1)$ -robust. Conversely, the complete digraph, denoted $K_n = (\mathcal{V}, \mathcal{E}_{K_n})$, with $\mathcal{E}_{K_n} = \{(i, j) \in \mathcal{V} \times \mathcal{V} : i \neq j\}$, is $(\lceil n/2 \rceil, s)$ -robust, for $1 \leq s \leq n$. Furthermore, whenever $n > 1$ is odd, K_n is the only digraph on n nodes that is $(\lceil n/2 \rceil, s)$ -robust with $s \geq \lfloor n/2 \rfloor$.

The next property relates robustness of the network to its minimum in-degree.

Lemma 5 (Minimum In-Degree): Given an (r, s) -robust digraph $\mathcal{D} = (\mathcal{V}, \mathcal{E})$, with $0 \leq r \leq \lceil n/2 \rceil$ and $1 \leq s \leq n$, the *minimum in-degree* of \mathcal{D} , $\delta^{\text{in}}(\mathcal{D})$, is at least

$$\delta^{\text{in}}(\mathcal{D}) \geq \begin{cases} r + s - 1 & \text{if } s < r; \\ 2r - 2 & \text{if } s \geq r. \end{cases}$$

The following result provides a lower bound on the amount of robustness that can be maintained in a digraph after removing incoming edges from nodes in the network.

Lemma 6 (Directed Edge Removal): Given an (r, s) -robust (p -fraction robust) digraph \mathcal{D} , let \mathcal{D}' be the digraph produced by removing up to k (q -fraction of) incoming edges of each node in \mathcal{D} , where $0 \leq k < r$ ($0 \leq q < p \leq 1$). Then \mathcal{D}' is $(r - k, s)$ -robust ($(p - q)$ -fraction robust).

Recall that when there are no misbehaving nodes, the Linear Consensus Protocol given in (1) achieves consensus if and only if the network contains a rooted out-branching. The following result shows that 1-robustness is equivalent to containing a rooted out-branching.

Lemma 7: A digraph \mathcal{D} is 1-robust if and only if \mathcal{D} contains a rooted-out branching.

Next, we relate the robustness of the underlying graph to its connectivity.

Theorem 6 (Connectivity of Robust Graphs): Suppose $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is an r -robust digraph, with $0 \leq r \leq \lceil n/2 \rceil$. Then the underlying graph $\mathcal{G}_{\mathcal{D}}$ is at least r -connected. Furthermore, if \mathcal{D} is (r, r) -robust, with $3 \leq r \leq \lceil n/2 \rceil$, then $\mathcal{G}_{\mathcal{D}}$ is at least $(\lceil 3r/2 \rceil - 1)$ -connected.

Finally, we discuss how to compare the robustness of different networks. Clearly, if digraph \mathcal{D}_1 is (r_1, s_1) -robust and digraph \mathcal{D}_2 is (r_2, s_2) -robust with maximal r_k and s_k for $k \in \{1, 2\}$, where $r_1 > r_2$ and $s_1 > s_2$, then one can conclude that \mathcal{D}_1 is more robust than \mathcal{D}_2 . However, in cases where $r_1 > r_2$ but $s_1 < s_2$, which digraph is more robust? For example, consider the graphs of Figs. 1 and 4. The graph in Fig. 1 can be shown to be $(2, s)$ -robust, for all $1 \leq s \leq n = 8$. This follows because *all* nodes in at least one of the sets \mathcal{S}_1 and \mathcal{S}_2 have at least 2 neighbors

outside of their set, for any nonempty and disjoint $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$. Therefore, condition (iii) in Definition 13 is never needed, and the definition is satisfied with $r = 2$ for all valid values of s . However, this graph is *not* 3-robust. This can be shown by selecting $\mathcal{S}_1 = X$ and $\mathcal{S}_2 = Y$. The graph in Fig. 4 is 3-robust, but is not (2,5)-robust (e.g., let $\mathcal{S}_1 = \{1, 5, 6\}$ and $\mathcal{S}_2 = \{2, 3, 4\}$).

In general, the parameter r in (r, s) -robustness takes precedence in the partial order that determines relative robustness, and the maximal s is used for ordering the robustness of networks with the same value of r . This choice is motivated by the dependence of the properties outlined in this section on the value of r . Therefore, the graph in Fig. 4 is more robust than the graph of Fig. 1. Yet, the graph of Fig. 4 is only 3-connected, whereas the graph of Fig. 1 is 5-connected. Hence, it is possible that a digraph with *less* connectivity is *more* robust.

VIII. CONCLUSION

The notion of graph connectivity has long been the backbone of investigations into fault tolerant and secure distributed algorithms. Indeed, under the assumption of full knowledge of the network topology, connectivity is *the key* metric in determining whether a fixed number of malicious adversaries can be overcome. However, in large scale systems and complex networks, it is not practical for the various nodes to obtain knowledge of the global network topology. This necessitates the development of algorithms that allow the nodes to be agnostic of the topology and identities of non-neighbor nodes, and operate on purely local information. This paper continues and extends the work started in [1], [2], [26]–[29], [31], [45], [46], and represents a step in this direction for the particular application of distributed consensus. Using the W-MSR algorithm and the notion of robust digraphs introduced in [2], and the extensions of each presented here, we characterize necessary/sufficient conditions for the normal nodes in large-scale networks to mitigate the influence of adversaries. We show that the notions of robust digraphs are the appropriate analogues to graph connectivity when considering purely local filtering rules at each node in the network. Just as connectivity has played a central role in the existing analysis of reliable distributed algorithms with global topological knowledge, we believe that robustness will play an important role in the investigation of purely local algorithms.

APPENDIX

A. Proof of Proposition 1

Proof: For simplicity, we focus on the case when n is even. Construct an *undirected* graph as follows. Let \mathcal{X} and \mathcal{Y} be two complete graphs on $\frac{n}{2}$ nodes. Number nodes in \mathcal{X} and \mathcal{Y} as $x_1, x_2, \dots, x_{\frac{n}{2}}$ and $y_1, y_2, \dots, y_{\frac{n}{2}}$, respectively. For any node $x_i \in \mathcal{X}$, if $i \leq |\mathcal{Y}| - F + 1$, connect x_i with nodes $y_i, y_{i+1}, \dots, y_{i+F-1}$; otherwise, connect x_i with nodes $y_i, \dots, y_{\frac{n}{2}}$ and nodes $y_1, \dots, y_{i+F-\frac{n}{2}-1}$. Then each node in \mathcal{X} and \mathcal{Y} has exactly F neighbors in the other set.

Next we will prove that the connectivity of this graph is $\frac{n}{2} + F - 1$. Let $\mathcal{C} = \{\mathcal{C}_\mathcal{X}, \mathcal{C}_\mathcal{Y}\}$ be a vertex cut, where $\mathcal{C}_\mathcal{X} = \mathcal{C} \cap \mathcal{X}$ and $\mathcal{C}_\mathcal{Y} = \mathcal{C} \cap \mathcal{Y}$. Without loss of generality, assume that $\mathcal{C}_\mathcal{X} = \{x_1, x_2, \dots, x_{|\mathcal{C}_\mathcal{X}|}\}$; other ways of choosing $\mathcal{C}_\mathcal{X}$ are equivalent to this situation by renumbering the nodes. By the definition of a vertex cut, we know $|\mathcal{C}_\mathcal{X}| \geq F$; otherwise, each node in $\mathcal{Y} \setminus \mathcal{C}_\mathcal{Y}$ still has at least one neighbor in \mathcal{X} , and since $\mathcal{X} \setminus \mathcal{C}_\mathcal{X}$ and $\mathcal{Y} \setminus \mathcal{C}_\mathcal{Y}$ each induce fully-connected subgraphs, we see that the graph will be connected (contradicting the fact that \mathcal{C} is a vertex cut). When $F \leq |\mathcal{C}_\mathcal{X}| < \frac{n}{2}$, the remaining nodes of \mathcal{X} collectively still have $k = \frac{n}{2} - |\mathcal{C}_\mathcal{X}| + F - 1$ neighbors in \mathcal{Y} , which implies we need to remove at least k nodes from \mathcal{Y} to disconnect the graph. When $\mathcal{C}_\mathcal{X} = \mathcal{X}$, since \mathcal{Y} is complete, we know $|\mathcal{C}_\mathcal{Y}| = \frac{n}{2} - 1$. Thus the connectivity of this graph is $\frac{n}{2} + F - 1$.

In this graph, assume that all nodes in \mathcal{X} have initial value a , and all nodes in \mathcal{Y} have initial value b , where $a < b$. When any node x_i applies the W-MSR algorithm, all of its F neighbors in \mathcal{Y} have the highest values in x_i 's neighborhood, and thus they are all disregarded. Similarly, all of y_i 's neighbors in \mathcal{X} are disregarded as well. Thus, each node in each set only uses the values from its own set, and no node ever changes its value, which shows that consensus will never be reached in this network. ■

B. Proof of Corollary 1

Proof: As in the proof of Theorem 1, we define the same terms and argue by contradiction. In this case, fix $\epsilon < \frac{\alpha^{Nc}}{1-\alpha^{Nc}}\epsilon_0$, which satisfies $\epsilon_0 > \epsilon > 0$. Let t_ϵ be such that $M[t] < A_M + \epsilon$ and $m[t] > A_m - \epsilon, \forall t \geq t_\epsilon$. By hypothesis, there exists $\tau_1 \in \{t_\epsilon, t_\epsilon + 1, \dots, t_\epsilon + c - 1\}$ such that $\mathcal{D}[\tau_1]$ is $(F+1, F+1)$ -robust. As in the proof of Theorem 1, there either exists $i \in \mathcal{X}_M(\tau_1, \epsilon_0) \cap \mathcal{N}$ such that $x_i[\tau_1 + 1] \leq A_M - \epsilon_1$ or $j \in \mathcal{X}_m(\tau_1, \epsilon_0) \cap \mathcal{N}$ such that $x_j[\tau_1 + 1] \geq A_m + \epsilon_1$, or both, where we

have defined $\epsilon_1 = \alpha\epsilon_0 - (1 - \alpha)\epsilon$. Note that as before, these inequalities hold for all normal nodes outside of the sets $\mathcal{X}_M(\tau_1, \epsilon_0)$ and $\mathcal{X}_m(\tau_1, \epsilon_0)$, respectively, and $0 < \epsilon < \epsilon_1 < \epsilon_0$ by the choice of ϵ . Furthermore, $|\mathcal{X}_M(\tau_1 + 1, \epsilon_1) \cap \mathcal{N}| < |\mathcal{X}_M(\tau_1, \epsilon_0) \cap \mathcal{N}|$ or $|\mathcal{X}_m(\tau_1 + 1, \epsilon_1) \cap \mathcal{N}| < |\mathcal{X}_m(\tau_1, \epsilon_0) \cap \mathcal{N}|$, or both.

Define recursively $\epsilon_k = \alpha\epsilon_{k-1} - (1 - \alpha)\epsilon$ for $1 \leq k \leq Nc$. Regardless of the network topology, we can show that any normal node i satisfying $x_i[\tau_1 + 1] \leq A_M - \epsilon_1$ will satisfy $x_i[\tau_1 + k] \leq A_M - \epsilon_k$ at time $\tau_1 + k$, for all $1 \leq k \leq Nc$. This holds because each normal node uses its own value with weight no smaller than α . Likewise, any normal node j satisfying $x_j[\tau_1 + 1] \geq A_m + \epsilon_1$ will satisfy $x_j[\tau_1 + k] \geq A_m + \epsilon_k$ at time $\tau_1 + k$, for all $1 \leq k \leq Nc$. Because of these relationships, we have that $|\mathcal{X}_M(\tau_1 + k, \epsilon_k) \cap \mathcal{N}| \leq |\mathcal{X}_M(\tau_1 + k - 1, \epsilon_{k-1}) \cap \mathcal{N}|$ and $|\mathcal{X}_m(\tau_1 + k, \epsilon_k) \cap \mathcal{N}| \leq |\mathcal{X}_m(\tau_1 + k - 1, \epsilon_{k-1}) \cap \mathcal{N}|$, for each time-step regardless of the network topology. However, we are interested in the time-steps τ_1, τ_2, \dots , in which $|\mathcal{X}_M(\tau_j + 1, \epsilon_{(1+\tau_j-\tau_1)}) \cap \mathcal{N}| < |\mathcal{X}_M(\tau_j, \epsilon_{(\tau_j-\tau_1)}) \cap \mathcal{N}|$ or $|\mathcal{X}_m(\tau_j + 1, \epsilon_{(1+\tau_j-\tau_1)}) \cap \mathcal{N}| < |\mathcal{X}_m(\tau_j, \epsilon_{(\tau_j-\tau_1)}) \cap \mathcal{N}|$. These time-steps correspond to the times at which $\mathcal{D}[\tau_j]$ is $(F + 1, F + 1)$ -robust and both $\mathcal{X}_M(\tau_j, \epsilon_{(\tau_j-\tau_1)})$ and $\mathcal{X}_m(\tau_j, \epsilon_{(\tau_j-\tau_1)})$ have at least one normal node, for $j \geq 1$ (by the argument made in the proof of Theorem 1). Since $|\mathcal{X}_M(\tau_1, \epsilon_0) \cap \mathcal{N}| + |\mathcal{X}_m(\tau_1, \epsilon_0) \cap \mathcal{N}| \leq N$ and $|\tau_N - \tau_1| \leq Nc$, there must be some time-step $\tau = \tau_1 + T$ (where $T \leq Nc$) where either $\mathcal{X}_M(\tau_1 + T, \epsilon_T) \cap \mathcal{N}$ or $\mathcal{X}_m(\tau_1 + T, \epsilon_T) \cap \mathcal{N}$ is empty. In the former case, all normal nodes in the network at time-step $\tau_1 + T$ have value at most $A_M - \epsilon_T$, and in the latter case all normal nodes in the network at time-step $\tau_1 + T$ have value no less than $A_m + \epsilon_T$. Since $\epsilon < \frac{\alpha^{Nc}}{1 - \alpha^{Nc}}\epsilon_0$, we can show that $\epsilon_T > 0$, producing the desired contradiction. ■

C. Proof of Lemma 2

Proof: If \mathcal{D} is empty or trivial, there is nothing to prove, so assume \mathcal{D} is nonempty and nontrivial. For any nonempty, disjoint pair $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$, at least one of the three conditions (i)–(iii) of Definition 13 holds. Observe that $|\mathcal{X}_{\mathcal{S}_k}^{r'}| \geq |\mathcal{X}_{\mathcal{S}_k}^r|$ for $k = 1, 2$. Hence if (i) or (ii) hold, then $|\mathcal{X}_{\mathcal{S}_k}^{r'}| \geq |\mathcal{X}_{\mathcal{S}_k}^r| = |\mathcal{S}_k| \geq |\mathcal{X}_{\mathcal{S}_k}^r|$, which implies $|\mathcal{X}_{\mathcal{S}_k}^{r'}| = |\mathcal{S}_k|$. If (iii) holds, then

$$|\mathcal{X}_{\mathcal{S}_1}^{r'}| + |\mathcal{X}_{\mathcal{S}_2}^{r'}| \geq |\mathcal{X}_{\mathcal{S}_1}^r| + |\mathcal{X}_{\mathcal{S}_2}^r| \geq s \geq s'.$$

Thus, any pair of nonempty, disjoint subsets of nodes in \mathcal{D} satisfy Definition 13 with r and s replaced by r' and s' . Therefore, \mathcal{D} is (r', s') -robust. ■

D. Proof of Lemma 3

Proof: Suppose \mathcal{D}' is not (r, s) -robust. Then there exists a pair of nonempty, disjoint subsets $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ such that all of the conditions (i)-(iii) in Definition 13 fail to hold with r and s . By removing directed edges in \mathcal{E}'' , the number of nodes in $\mathcal{X}_{\mathcal{S}_1}^r$ and $\mathcal{X}_{\mathcal{S}_2}^r$ can only decrease, and therefore none of conditions (i)-(iii) hold for the pair $\mathcal{S}_1, \mathcal{S}_2$ in \mathcal{D} . Hence, \mathcal{D} is not (r, s) -robust, which is a contradiction. ■

E. Proof of Lemma 4

Proof: Assume \mathcal{D} is nonempty and nontrivial (otherwise, the result holds by definition). Pick \mathcal{S}_1 and \mathcal{S}_2 by taking any bipartition of \mathcal{V} such that $|\mathcal{S}_1| = \lceil n/2 \rceil$ and $|\mathcal{S}_2| = \lfloor n/2 \rfloor$. Neither \mathcal{S}_1 nor \mathcal{S}_2 have $\lceil n/2 \rceil + 1$ nodes; therefore, neither one is $(\lceil n/2 \rceil + 1)$ -reachable. Hence, \mathcal{D} is not $(\lceil n/2 \rceil + 1)$ -robust. Now suppose $\mathcal{D} = K_n$. For any nonempty, disjoint $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$, $|\mathcal{V} \setminus \mathcal{S}_i| \geq \lceil n/2 \rceil$ holds for at least one of $i \in \{1, 2\}$. For whichever i this holds, $|\mathcal{X}_{\mathcal{S}_i}^{\lceil n/2 \rceil}| = |\mathcal{S}_i|$, so that K_n is $(\lceil n/2 \rceil, s)$ -robust, for $1 \leq s \leq n$. For the last statement, we show that whenever $n > 1$ is odd, removing any directed edge from K_n causes the resulting digraph to lose $(\lceil n/2 \rceil, \lfloor n/2 \rfloor)$ -robustness. Suppose $e = (i, j)$ is the directed edge removed from \mathcal{E}_{K_n} to form $\mathcal{D}'' = (\mathcal{V}, \mathcal{E}'')$, with $\mathcal{E}'' = \mathcal{E}_{K_n} \setminus \{e\}$. Choose \mathcal{S}_1 and \mathcal{S}_2 by taking any bipartition of \mathcal{V} in \mathcal{D}'' such that $|\mathcal{S}_1| = \lceil n/2 \rceil$, $|\mathcal{S}_2| = \lfloor n/2 \rfloor$, $i \in \mathcal{S}_1$, and $j \in \mathcal{S}_2$. Then, $|\mathcal{X}_{\mathcal{S}_1}^{\lceil n/2 \rceil}| = 0$ and $|\mathcal{X}_{\mathcal{S}_2}^{\lceil n/2 \rceil}| = \lfloor n/2 \rfloor - 1 < |\mathcal{S}_2|$. Therefore, \mathcal{D}'' is not $(\lceil n/2 \rceil, s)$ -robust for $s \geq \lfloor n/2 \rfloor$. This is sufficient to prove the statement because of the monotonicity result of Lemma 3, combined with the fact that any spanning subdigraph of K_n , $\mathcal{D}' = (\mathcal{V}, \mathcal{E}') \subset K_n$, can be obtained from a directed edge removal process starting with some directed edge $e = (i, j) \notin \mathcal{E}'$. ■

F. Proof of Lemma 5

Proof: Whenever $r \in \{0, 1\}$, there is nothing to prove. Also, if $n \leq 2$, then $r \leq 1$. Therefore, assume $n \geq 3$ and $2 \leq r \leq \lceil n/2 \rceil$. Fix $j \in \mathcal{V}$. First, let $\mathcal{S}_1 = \{j\}$ and $\mathcal{S}_2 = \mathcal{V} \setminus \mathcal{S}_1$. Then, $|\mathcal{X}_{\mathcal{S}_2}^r| = 0$ so that $|\mathcal{X}_{\mathcal{S}_1}^r| = |\mathcal{S}_1|$. This proves $d_j \geq r$. Next, whenever $s < r$, form \mathcal{S}_1 by choosing $s - 1$ of node j 's in-neighbors along with j itself. Take $\mathcal{S}_2 = \mathcal{V} \setminus \mathcal{S}_1$ as before. Since $|\mathcal{S}_1| = s < r$, again $|\mathcal{X}_{\mathcal{S}_2}^r| = 0$ so that $|\mathcal{X}_{\mathcal{S}_1}^r| = |\mathcal{S}_1|$. This implies j has an additional r in-neighbors outside of \mathcal{S}_1 , thereby guaranteeing $d_j \geq r + s - 1$. On the other hand, whenever $s \geq r$, form \mathcal{S}_1 by choosing $r - 2$ of node j 's in-neighbors along with j itself. Again, choose $\mathcal{S}_2 = \mathcal{V} \setminus \mathcal{S}_1$. Since

$|\mathcal{S}_1| < r$ and $s \geq r$, again $|\mathcal{X}_{\mathcal{S}_2}^r| = 0$ so that $|\mathcal{X}_{\mathcal{S}_1}^r| = |\mathcal{S}_1|$. This implies j has an additional r in-neighbors outside of \mathcal{S}_1 , thereby guaranteeing $d_j \geq 2r - 2$. Since $j \in \mathcal{V}$ is arbitrary, the bound on $\delta^{\text{in}}(\mathcal{D})$ follows. ■

G. Proof of Lemma 6

Proof: From the definition of (r, s) -reachable (p -fraction reachable) set, we know that if a set is (r, s) -reachable (p -fraction reachable), then by removing up to k (q -fraction of) incoming edges of each node in \mathcal{D} , where $0 \leq k < r$ ($0 \leq q < p < 1$), the set is $(r - k, s)$ -reachable ($(p - q)$ -fraction reachable). Thus, by the definition of (r, s) -robustness (p -fraction robustness), the result follows. ■

H. Proof of Lemma 7

Proof: If \mathcal{D} is 1-robust, we will prove that \mathcal{D} has a rooted out-branching by contradiction. Assume that \mathcal{D} does not have a rooted out-branching. Decompose \mathcal{D} into its strongly connected components, and note that since \mathcal{D} does not have a rooted out-branching, there must be at least two components that have no incoming edges from any other components. However, this contradicts the assumption that \mathcal{D} is 1-robust (at least one of the two subsets must have a neighbor outside the set), so it must be true that there exists a rooted out-branching.

Assume \mathcal{D} contains a rooted out-branching, but is not 1-robust. Then we can find two subsets of nodes which do not have neighbors from outside, which contradicts with the assumption that \mathcal{D} contains a rooted out-branching, completing the proof. ■

Remark 1: The proof of Lemma 7 is a more direct version of the proof of Theorem 5 in [41].

I. Proof of Theorem 6

Proof: If $r = 0$, the first statement is vacuously true, and if $r = 1$, it holds by Lemma 7. Therefore, assume $r \geq 2$. By Lemma 3, the underlying graph $\mathcal{G}_{\mathcal{D}} = (\mathcal{V}, \mathcal{E}_{\mathcal{G}})$ is r -robust. By Lemmas 2 and 7, the graph is connected. Suppose there is a vertex cut $\mathcal{K} \subset \mathcal{V}$ such that $|\mathcal{K}| < r$, and denote the $k \geq 2$ connected components remaining after the removal of \mathcal{K} by $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$. Let $\mathcal{S}_1 = \mathcal{C}_1$ and $\mathcal{S}_2 = \mathcal{C}_2$. Since $\mathcal{G}_{\mathcal{D}}$ is r -robust, either \mathcal{S}_1 or \mathcal{S}_2 is r -reachable, which contradicts the fact that \mathcal{K} is a vertex cut. Hence, any vertex cut \mathcal{K} must satisfy $|\mathcal{K}| \geq r$, so that $\mathcal{G}_{\mathcal{D}}$ is at least r -connected.

For the second statement, suppose there is a vertex cut $\mathcal{K} \subset \mathcal{V}$ such that $r \leq |\mathcal{K}| \leq \lceil 3r/2 \rceil - 2$, and denote the $k \geq 2$ connected components remaining after the removal of \mathcal{K} by $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$. Partition \mathcal{K} into $\mathcal{K} = \mathcal{K}_1 \cup \mathcal{K}_2 \cup \mathcal{K}_3$ such that $|\mathcal{K}_1| = |\mathcal{K}_2| = \lceil r/2 \rceil - 1$ and the remaining nodes go to \mathcal{K}_3 ; i.e., $1 \leq |\mathcal{K}_3| \leq \lfloor r/2 \rfloor$. Then form $\mathcal{S}_1 = \mathcal{C}_1 \cup \mathcal{K}_1$ and $\mathcal{S}_2 = \mathcal{C}_2 \cup \mathcal{K}_2$. Since $\mathcal{G}_{\mathcal{D}}$ is (r, r) -robust by Lemma 3, $\delta(\mathcal{G}_{\mathcal{D}}) \geq 2r - 2$ by Lemma 5, so that $|\mathcal{C}_i| \geq \lfloor r/2 \rfloor + 1$ (since there are at most $\lceil 3r/2 \rceil - 2$ neighbors in \mathcal{K}). It follows that $|\mathcal{S}_1|, |\mathcal{S}_2| \geq r$, and we are guaranteed $|\mathcal{X}_{\mathcal{S}_1}^r| + |\mathcal{X}_{\mathcal{S}_2}^r| \geq r$. Because $|\mathcal{K}_1 \cup \mathcal{K}_2| \leq r - 1$ and $r \geq 3$, there is $v \in \mathcal{C}_1 \cup \mathcal{C}_2$ such that v has at least r neighbors outside of its set. Without loss of generality, assume $v \in \mathcal{C}_1$. Since $|\mathcal{K}_2| + |\mathcal{K}_3| \leq r - 1$, $\exists j \in \mathcal{C}_2 \cup \dots \cup \mathcal{C}_k$ such that $(j, v) \in \mathcal{E}$, which contradicts the fact that \mathcal{K} is a vertex cut whose removal results in components $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k$. Hence, $\mathcal{G}_{\mathcal{D}}$ is at least $(\lceil 3r/2 \rceil - 1)$ -connected. ■

ACKNOWLEDGMENTS

The authors thank Nitin Vaidya for helpful discussions and for the pointer to related work. H. J. LeBlanc and X. Koutsoukos are 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. H. Zhang and S. Sundaram are supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada (NSERC), and by a grant from the Waterloo Institute for Complexity and Innovation (WICI).

REFERENCES

- [1] 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.
- [2] 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.
- [3] A. Giridhar and P. R. Kumar, "Toward a theory of in-network computation in wireless sensor networks," *IEEE Communications Magazine*, vol. 44, no. 4, pp. 98–107, April 2006.
- [4] 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.
- [5] N. A. Lynch, *Distributed Algorithms*. San Francisco, California: Morgan Kaufmann Publishers Inc., 1997.
- [6] 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.

- [7] J. N. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Department of EECS, MIT, 1984.
- [8] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation and consensus using linear iterative strategies," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 4, pp. 650–660, May 2008.
- [9] D. Kempe, A. Dobra, and J. Gehrke, "Gossip-based computation of aggregate information," in *44th Annual IEEE Symposium on Foundations of Computer Science*, Cambridge, MA, Oct. 2003, pp. 482–491.
- [10] 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.
- [11] J. N. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803–812, 1986.
- [12] P. A. Forero, A. Cano, and G. B. Giannakis, "Consensus-based distributed support vector machines," *The Journal of Machine Learning Research*, vol. 11, pp. 1663–1707, 2010.
- [13] 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.
- [14] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 2, pp. 382–401, 1982.
- [15] M. O. Rabin, "Randomized Byzantine generals," in *24th Annual Symposium on Foundations of Computer Science*, Nov. 1983, pp. 403–409.
- [16] M. Ben-Or, "Another advantage of free choice: Completely asynchronous agreement protocols," in *Proceedings of the 2nd Annual ACM Symposium on Principles of Distributed Computing*, ser. (PODC), Montréal, Quebec, Canada, 1983, pp. 27–30.
- [17] A. D. Fekete, "Asymptotically optimal algorithms for approximate agreement," *Distributed Computing*, vol. 4, pp. 9–29, Mar. 1990.
- [18] M. Ben-Or, D. Dolev, and E. N. Hoch, "Simple gradecast based algorithms," *CoRR*, vol. abs/1007.1049, 2010.
- [19] J. Hromkovic, R. Klasing, A. Pelc, P. Ruzicka, and W. Unger, *Dissemination of Information in Communication Networks*. Springer-Verlag, 2005.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] 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.

- [26] 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.
- [27] —, "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.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] —, "Low cost approximate agreement in partially connected networks," *Journal of Computing and Information*, vol. 3, no. 1, pp. 53–85, 1993.
- [32] H. J. LeBlanc, "Resilient cooperative control of networked multi-agent systems," Ph.D. dissertation, Department of EECS, Vanderbilt University, 2012.
- [33] H. Zhang, "Network Robustness: Diffusing Information Despite Adversaries," M.S. thesis, Department of ECE, University of Waterloo, 2012.
- [34] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, June 2006.
- [35] A. Pelc and D. Peleg, "Broadcasting with locally bounded Byzantine faults," in *Information Processing Letters*, 2005, pp. 109–115.
- [36] 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.
- [37] D. Easley and J. Kleinberg, *Networks, Crowds and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, 2010.
- [38] 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.
- [39] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Systems & Control Letters*, vol. 53, pp. 65–78, 2004.
- [40] W. Ren and R. W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 655–661, May 2005.
- [41] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 169–182, Feb. 2005.
- [42] B. Touri and A. Nedić, "On ergodicity, infinite flow, and consensus in random models," *IEEE Transactions on Automatic Control*, vol. 56, no. 7, pp. 1593–1605, July 2011.
- [43] J. Lorenz and D. A. Lorenz, "On conditions for convergence to consensus," *IEEE Transactions on Automatic Control*, vol. 55, no. 7, pp. 1651–1656, July 2010.
- [44] 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.

- [45] A. H. Azadmanesh and H. Bajwa, "Global convergence in partially fully connected networks (pfcn) with limited relays," *The 27th Annual Conference of the IEEE Industrial Electronics Society*, vol. 3, pp. 2022–2025, 2001.
- [46] M. H. Azadmanesh and R. M. Kieckhafer, "Asynchronous approximate agreement in partially connected networks," *International Journal of Parallel and Distributed Systems and Networks*, vol. 5, no. 1, pp. 26–34, 2002.
- [47] J. Li, E. Elhamifar, I. J. Wang, and R. Vidal, "Consensus with robustness to outliers via distributed optimization," in *IEEE Conference on Decision and Control*, Atlanta, GA, Dec. 2010, pp. 2111–2117.
- [48] E. Montijano, S. Martínez, and S. Sagués, "De-RANSAC: robust distributed consensus in sensor networks," *European Journal of Control*, submitted 2012.
- [49] S. Hoory, N. Linial, and A. Wigderson, "Expander graphs and their applications," *Bulletin of the American Mathematical Society*, vol. 43, no. 4, pp. 439–561, Oct. 2006.
- [50] A. Lubotzky, "Expander graphs in pure and applied mathematics," *Bulletin of the American Mathematical Society*, vol. 49, no. 1, pp. 113–162, Jan. 2012.
- [51] N. H. Vaidya, "Matrix representation of iterative approximate Byzantine consensus in directed graphs," *CoRR*, vol. abs/1201.1888, 2012.
- [52] R. Albert and A. L. Barabási, "Statistical mechanics of complex networks," *Rev. Mod. Phys.*, vol. 74, no. 1, pp. 47–97, Jan. 2002.