Network Robustness: Diffusing Information Despite Adversaries

Haotian Zhang[1]*, Shreyas Sundaram[1]*

Abstract
In this report, we consider the problem of diffusing information resiliently in networks that contain misbehaving nodes. Previous strategies to achieve resilient information diffusion typically require the normal nodes to hold some global information, such as the topology of the network and the identities of non-neighboring nodes. However, these assumptions are not suitable for large-scale networks and this necessitates our study of resilient algorithms based on only local information.

We propose a consensus algorithm where, at each time-step, each normal node removes the extreme values in its neighborhood and updates its value as a weighted average of its own value and the remaining values. We show that traditional topological metrics (such as connectivity of the network) fail to capture such dynamics. Thus, we introduce a topological property termed as *network robustness* and show that this concept, together with its variants, is the key property to characterize the behavior of a class of resilient algorithms that use purely local information.

We then investigate the robustness properties of complex networks. Specifically, we consider common random graph models for complex networks, including the preferential attachment model, the Erdős-Rényi model, and the geometric random graph model, and compare the metrics of connectivity and robustness in these models. While connectivity and robustness are greatly different in general (i.e., there exist graphs which are highly connected but with poor robustness), we show that the notions of robustness and connectivity are equivalent in the preferential attachment model, cannot be very different in the geometric random graph model, and share the same threshold functions in the Erdős-Rényi model, which gives us more insight about the structure of complex networks. Finally, we provide a construction method for robust graphs.

Keywords
Network robustness — Resilient consensus — Complex networks

[1] *Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, Canada*
***Corresponding authors**: h223zhan, ssundaram@uwaterloo.ca

Contents

1. Introduction

A core question in the study of large-scale and complex networks (both natural and engineered) is: how do the global

behaviors emerge from local interactions? For instance, the fields of sociology and epidemiology examine the spread of ideas, decisions and diseases through populations of people, based on the patterns of contact between the individuals in the population [1–3]. In this context, one can ask whether a few stubborn individuals (who do not change their beliefs) are able to affect the decisions reached by the rest of the population [4, 5]. Similarly, the efficacy of engineered networks (such as communication networks, or multi-agent systems) is often predicated on their ability to disseminate information throughout the network [6, 7]. For example, the 'broadcast' operation is used as a building block for more complex functions, allowing certain nodes to inform all other nodes of pertinent information [6].

The ability of a few individuals to affect the global behavior of the system is clearly a double-edged sword. When the network contains legitimate leaders or experts, it is beneficial to ensure that the innovations introduced by these small groups spread throughout the population. On the other hand, networks that facilitate diffusion are also vulnerable to disruption by individuals that are not acting with the best interests of the society in mind. In engineering applications, these individuals could correspond to faulty or malicious nodes that do not follow preprogrammed strategies due to malfunctions or attacks, respectively. Thus, a fundamental challenge is to identify network properties and diffusion dynamics that allow legitimate information to propagate throughout the network, while limiting the effects of illegitimate individuals and actions.

Another fundamental challenge in large-scale networks is that the quantities of interest must be computed using only *local information*, i.e., information obtained by each node through sensor measurements, calculations, or communication only with neighbors in the network. To obtain the desired computational result, it is important to design the fault tolerant algorithms to be able to withstand the compromise of a subset of the nodes and still ensure some level of correctness (possibly at a degraded level of performance). We refer to such a networked system as being *resilient* to adversaries.

Distributed consensus is one of the most important objectives in large-scale networks and has applications in various areas such as data aggregation [8], distributed optimization [9], and flocking [10]. In the consensus problem, we want the nodes to reach agreement on some value corresponding to some function of their initial states. A fundamental challenge for reaching distributed consensus in these networks is that they are vulnerable to attacks or failures at one or more of the nodes in the network. The problem of reaching consensus resiliently in the presence of misbehaving nodes has been studied extensively by various communities (e.g., see [6, 11] and the references therein). Among other things, it has been shown that given F (worst-case) adversarial nodes, there exist strategies for these nodes to disrupt consensus if the network connectivity[1] is $2F$ or less. Conversely, if the network connectivity is at least $2F + 1$, then there exist strategies for the normal nodes to use that ensure consensus is reached (under the local broadcast model of communication) [11–13]. However, these consensus algorithms either require that normal nodes have at least some global information (e.g., topology of the network) or assume that the network is *complete*, i.e., all-to-all communication or sensing [14–18]. Moreover, these algorithms lead to expensive computation and communication costs and need the normal nodes to store large amounts of data. Therefore, there is a need for resilient consensus algorithms that are more lightweight and operate using only local information (i.e., without knowledge of the network topology and the identities of non-neighboring nodes).

In this report, we study resilient information diffusion dynamics that use only local information, and characterize topological properties that facilitate such diffusion. The rest of this report will be organized as follows. In Section 2, we focus on the consensus problem and 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 [11, 12]), are *not* adequate (i.e., there exist graphs with large connectivity and minimum degree but fail to reach consensus) 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 introduce an efficient resilient consensus algorithm (the W-MSR algorithm), and provide a comprehensive characterization of the network topologies where algorithms such as W-MSR (which uses only local information) can succeed despite the presence of a broad class of adversaries. In Section 3, we describe two extensions of network robustness – (r, s)-robustness and strong robustness. The concept of (r, s)-robustness is a strict generalization of network robustness and characterizes two types of information redundancy; by using this concept, we give *necessary and sufficient* condition for the W-MSR algorithm to achieve resilient consensus under some specific fault model. In addition to consensus, we also consider the problem of fault tolerant broadcast, and use robustness to provide conditions for the operations to succeed. Although connectivity and robustness are very different in general, in Section 4, we answer the question: how robust are complex networks? More specifically, how do the metrics of connectivity and robustness compare for various random graph constructions that are commonly used to model complex networks? We study three random graph models (preferential attachment, Erdős-Rényi, and geometric random graphs) for complex networks; we show that these models demonstrate a threshold behavior for robustness, whereby a certain degree of robustness is almost surely inherent in the network if the links are added with a probability above a certain value. Our analysis reveals that the notions of robustness and connectivity 'coincide' on these

[1] The network connectivity is defined as the number of nodes that have to be removed before the network becomes disconnected.

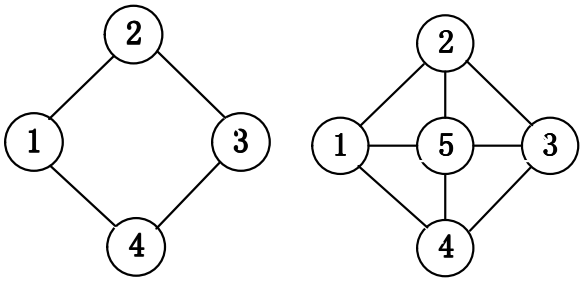

(a) 2-connected graph. (b) 3-connected graph.

Figure 1. Illustration of the role of connectivity in resilient consensus.

random graph models,[2] and this indicates that local filtering dynamics will be effective at facilitating resilient agreement on complex networks. Inspired by the preferential attachment mechanism, we also provide a construction method for robust networks. Finally, some conclusions and future research directions are given in Section 5.

1.1 Background on Resilient Information Diffusion

As mentioned earlier, the *connectivity* of the network has traditionally been viewed as the key metric with regard to resilience of consensus algorithms (and information diffusion algorithms in general). The intuition behind the role of connectivity in resilient consensus is illustrated by the following example. Consider the undirected networks in Figure 1. In both networks, suppose that node 1 wishes to obtain some information about the value of node 3, and that node 2 is malicious. In Figure 1(a), where the graph is only 2-connected, node 2 can prevent node 1 from getting information from node 3 by pretending that node 3's value is something that it is not. Node 1 would never know whether to trust the information it receives from node 2, or from node 4, and thus could never resolve the ambiguity caused by node 2's deception. However, in Figure 1(b), when the graph is 3-connected, there are three disjoint paths from node 3 to node 1, and thus node 1 will receive matching information about node 3 from at least two different paths; it could then use an appropriate scheme (such as majority voting) to eliminate node 2's influence.

More generally, if the connectivity of the network is $2F$ or less (for some nonnegative integer F), then there exists at least one set of F coordinated malicious nodes that can prevent the network from reaching consensus *regardless* of the mechanism that is used to achieve consensus [6, 11, 19]. On the other hand, if the connectivity is $2F + 1$ or higher, various algorithms have been proposed to overcome certain class of misbehaving nodes [11–13].

Remark 1. *Note that the network connectivity is fundamental when characterizing the ability of the network to tolerate the removal of nodes (either accidentally or by intent). The resilience of complex networks to such structural perturbations was studied in [20, 21]. In such cases, the network is required to have connectivity $F+1$ if one wishes to tolerate the removal of up to F nodes. Note the difference in these latter scenarios from the ones outlined above; node removal corresponds to a structural attack, whereas malicious behavior corresponds to attacks on the dynamical process running on the network.*

While the above connectivity bounds provide fundamental limitations on the resilience of networks to misbehaving nodes, the mechanisms proposed to overcome malicious behavior typically make the prohibitive assumption that all nodes know the entire network topology. While this assumption allows the 'good' agents in the network to *completely* overcome the effects of worst-case behavior by the malicious nodes, the bookkeeping burden is unrealistic in real-world complex networks with large number of nodes. To remedy this, we need efficient mechanisms to facilitate consensus that require only local information, such as the W-MSR algorithm proposed in this report.

1.2 Notation and Terminology

Throughout this report, 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 set difference of $\mathcal{S}_1$ by $\mathcal{S}_2$ is denoted $\mathcal{S}_1 \setminus \mathcal{S}_2 = \{x \in \mathcal{S}_1 \colon x \notin \mathcal{S}_2\}$. Furthermore, we will use the following notation regarding the asymptotic behavior of functions:

- $f(x) = O(g(x))$ if there exist constants c and x_0 such that $f(x) \leq cg(x)$ for any $x \geq x_0$;
- $f(x) = \Omega(g(x))$ if $g(x) = O(f(x))$;
- $f(x) = \Theta(g(x))$ if $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$;
- $f(x) = o(g(x))$ if $\frac{f(x)}{g(x)} \to 0$ as $x \to \infty$.

1.2.1 Graph Theory

A finite simple (directed) graph is denoted $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$, in which $\mathcal{V}$ is the *node set* and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the *directed edge set*. With a slight abuse of terminology, we will often refer to the network and the graph 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 $\mathcal{E}_{\mathcal{G}}$. We may also simply use $\mathcal{G}$ to represent an *undirected* graph when the context is clear. A graph $\mathcal{D}_1 = \{\mathcal{V}_1, \mathcal{E}_1\}$ is a *subgraph* of $\mathcal{D}$, written $\mathcal{D}_1 \subseteq \mathcal{D}$, if $\mathcal{V}_1 \subseteq \mathcal{V}$ and $\mathcal{E}_1 \subseteq \mathcal{E}$. A graph $\mathcal{D}' = \{\mathcal{V}', \mathcal{E}'\}$ is *isomorphic* to $\mathcal{D}$ if there exists a bijection $\psi \colon \mathcal{V} \to \mathcal{V}'$ such that $(i, j) \in \mathcal{E}$ if and only if $(\psi(i), \psi(j)) \in \mathcal{E}'$.

[2] More precisely, robustness and connectivity are equivalent in the preferential attachment model, cannot be very different in the geometric random graph model, and share the same threshold functions in the Erdős-Rényi model.

A *path* is a sequence of distinct vertices $i_0, i_1, \ldots, i_k$ such that $(i_j, i_{j+1}) \in \mathcal{E}$, $j = 0, 1, \ldots, k-1$. We use the notion of a 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 graph has a *directed spanning tree* if there exists a node r, the root, such that for each $i \in \mathcal{V}$, there exists a path from r to i.

2. Network Robustness

2.1 Introduction

As explored in Section 1.1, 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 adversaries can be overcome. However, in large-scale systems and complex networks, it is not practical for the nodes to obtain knowledge of the global network topology.[3] This necessitates the development of algorithms that allow the nodes to operate on purely local information. In this section, we focus on distributed consensus, which is an important operation in networks, and introduce a resilient consensus algorithm - the W-MSR algorithm. In order to characterize the performance of algorithms using only local information, such as W-MSR, we develop the notion of *network robustness* and provide necessary/sufficient conditions for the normal nodes in large-scale networks to mitigate the influence of adversaries. We show that the notion of robustness is the appropriate analog to graph connectivity when considering purely local filtering rules at each node in the network.

2.2 Problem Formulation

2.2.1 Network Model

Consider a time-varying network modeled by the *(directed) graph* $\mathcal{D}[t] = \{\mathcal{V}, \mathcal{E}[t]\}$, where $\mathcal{V} = \{1, ..., n\}$ 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 (or misbehaving)* 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 $\mathcal{V}_i[t] = \{j \in \mathcal{V} \colon (j, i) \in \mathcal{E}[t]\}$ 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 $\mathcal{V}_i^{\text{out}}[t] = \{j \in \mathcal{V} \colon (i, j) \in \mathcal{E}[t]\}$. 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.

[3] Note that only knowing the connectivity is not sufficient for previous strategies to operate.

2.2.2 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.). In order to achieve some specified objective, 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

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

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 (potentially time-varying) function of the values from node i's inclusive neighborhood, 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. Note that the strategy can be easily extended to the case $x_i[0] \in \mathbb{R}^m$ by doing componentwise iterations. 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.

2.2.3 Fault 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 malicious nodes. 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 nodes, we simply say those nodes are *misbehaving*.

It is clear that we cannot deal with networks that only contain misbehaving nodes and thus it is necessary to restrict the *number* of misbehaving nodes. We 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 each node which is not in $\mathcal{S}$ for all t, i.e., $|\mathcal{V}_i[t] \bigcap \mathcal{S}| \leq F$, $\forall i \in \mathcal{V} \setminus \mathcal{S}$, $\forall t \in \mathbb{Z}_{\geq 0}$, $F \in \mathbb{Z}_{\geq 0}$.*

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 each node which is not in $\mathcal{S}$ for all t, i.e., $|\mathcal{V}_i[t] \bigcap \mathcal{S}| \leq f|\mathcal{V}_i[t]|$, $\forall i \in \mathcal{V} \setminus \mathcal{S}$, $\forall t \in \mathbb{Z}_{\geq 0}$, $0 \leq f \leq 1$.*

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 steps. These definitions facilitate the following fault models.

Definition 6. *A set of misbehaving 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 fault models as the **F-total**, **F-local**, and **f-fraction local models**, respectively.*

F-totally bounded faults have been studied in distributed computing [11, 14, 22] 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 [23, 24]. However, to the best of our knowledge, there are no prior works discussing the f-fraction local model; our investigation of this motif is inspired by ideas pertaining to *contagion* in social and economic networks [3], 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 interpretation placed in previous work.

2.2.4 Resilient Asymptotic Consensus

Given the fault models, 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 $\lim_{t\to\infty} x_i[t] = L$ 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.

The resilient asymptotic consensus problem has three 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 safety condition is important when the current estimate of the consensus value is used in a safety critical process and 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.

2.3 Resilient Consensus Algorithm: W-MSR

While there are various approaches to facilitate consensus, a class of *linear algorithms* have attracted significant interest in recent years [25, 26]. This has been largely due to the applicability of linear systems theory and matrix theory to analyzing such strategies, but is also motivated by the resulting low communication overhead and simplicity. In such strategies, at each time-step t, each node senses or receives information from its neighbors, and changes its value according to

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

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 [10, 27–30]. 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 graph has a *directed spanning tree* [26]. 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 graphs across each interval has a directed spanning tree [28]. 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 [31]. Finally, the lower bound on the weights is needed because there are examples of asymptotically vanishing weights in which consensus is not reached [32].

One problem with the linear update given in (1) is that it is not resilient to misbehaving nodes. In the rest of this subsection, we will discuss the sensitivity of LCP to misbehaving nodes (and more generally, the topic of resilient consensus), and introduce an efficient algorithm to remedy this using only local information.

2.3.1 Background on Resilient Consensus

We first provide more background on resilient consensus. Note that the problem of resilient consensus has been investigated in the computer science community for several decades [11], but here we focus on LCP which has been studied extensively by the control community. While various conditions have been provided to guarantee consensus in the absence of misbehaving nodes, it was shown in [10, 33] that consensus can be disrupted by even a single node that updates its values arbitrarily. The paper [10] studied the use of linear iterative strategies as a mechanism for achieving *flocking* behavior in multi-agent systems. They showed that if a 'leader' node in the network does not update its value at each time-step (i.e., it maintains a constant value), then the linear iterative strategy will cause all nodes to asymptotically converge to the value of the leader. While this may be acceptable behavior when the network has a legitimate leader, it also seems to indicate that a simple asymptotic consensus scheme can be easily disrupted by just a single malicious node. A similar analysis was done in [33], where it was argued that since the asymptotic consensus scheme can be disrupted by a single node that maintains a constant value, it can also be disrupted by a single node that updates its values arbitrarily (since maintaining a constant value is a special case of arbitrary updates). Both of these works only considered a straightforward application of the linear iteration for asymptotic consensus, without having the normal nodes perform any operations to avoid the influence of malicious behavior.

In [12], the authors provided a comprehensive analysis of linear iterative strategies in the presence of malicious nodes. They demonstrated that linear iterative strategies are able to achieve the minimum bound required to disseminate information reliably; specifically, when a network is $2F+1$ connected, F malicious nodes will be unable to prevent any node from calculating any function of the initial values (under the broadcast model of communication). Variations of this problem were studied in [13, 34]. The result was extended in [13] to analyze linear iterative strategies for asymptotic consensus in the presence of faulty agents (in addition to malicious agents), and [34] studied the problem of detecting attacks in networks of linear continuous-time systems. While these results require minimal connectivity, they also require each normal node to have full knowledge of the network topology, along with strong computational and storage capabilities. The paper [35] considered the problem of reducing the influence of external intruders on asymptotic consensus in tree networks. They proposed a rewiring scheme whereby each node changes its parent node in an effort to slow down the effect of externally connected adversaries. While the approach presented in that paper is distributed, it only applies to tree topologies and requires that the location and intention of the adversaries to be known by the nodes.

2.3.2 Description of W-MSR

As argued before in Section 2.3.1, a single 'leader' node can cause all agents to reach consensus on an arbitrary value of its choosing (potentially resulting in a dangerous situation) simply by holding its value constant. Thus 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. We focus first on the F-local and F-total models, and then describe the modifications required for the f-fraction local model. 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:

- At each time-step t, each normal node i obtains the values of its neighbors, and forms a sorted list.
- 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.
- 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

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

where the weights $w_{ij}[t]$ satisfy the conditions stated before, but with $\mathcal{J}_i[t]$ replaced by $\mathcal{J}_i[t] \setminus \mathcal{R}_i[t]$.[4]

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).[5]

Remark 2. *Consensus algorithms with state-dependent switching have drawn increased attention in recent years in the context of opinion dynamics [36, 37]. For example, the following model was introduced in [36] to capture opinion dynamics in networks:*

$$x_i[t+1] = \frac{\sum_{j:|x_i[t]-x_j[t]|<1} x_j[t]}{|\{j : |x_i[t] - x_j[t]| < 1\}|}.$$

The constraint $|x_i[t] - x_j[t]| < 1$ represents 'bounded confidence' among these nodes: each node considers one of its neighbors' opinions as reasonable and accepts it if their opinions differ by less than 1. There are various differences in the analysis in these previous works in comparison with this paper. First, the above updating scheme assumes that the underlying graph is complete, so that each node sees all other nodes and selects only those whose values are close to its own. Second, there exists a fixed threshold (1 in the above scheme) to represent 'bounded confidence', and this might cause the agents to converge to different clusters for certain choices of initial states [37]. Most importantly, these previous works on state-dependent connectivity do not consider the presence of misbehaving nodes; we posit that the fixed threshold in the update rule still allows a misbehaving node to draw all of the other nodes to any desired consensus value, simply by waiting until all node values have converged sufficiently close together, and then slowly inducing drift by keeping its value near the edge of the fixed threshold. The algorithm considered in this report, on the other hand, applies to general topologies and inherently limits the amount of bias that can be introduced by a broad class of misbehaving nodes.

The above algorithm is more lightweight than previous strategies, 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 is: in what networks will the above algorithm facilitate resilient asymptotic consensus?

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

[5] Note that from the view of the whole system, the W-MSR algorithm is nonlinear.

2.3.3 Related Algorithms in Previous Work

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 [38], 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 so-called *Mean-Subsequence-Reduced (MSR)* algorithms [39]. There are a few key differences between the operations used in the W-MSR algorithm and the MSR algorithm of [39]. First, W-MSR does not always remove the largest and smallest F values as in the MSR algorithm [39]. 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. However, because the lower bound on the weights, $\alpha > 0$, may be arbitrarily small, W-MSR can come arbitrarily close to selecting only a subsequence of remaining values by setting the appropriate weights to α (instead of 0 as would be done in MSR). Finally, MSR averages the remaining values instead of allowing for weighted averages as in W-MSR.

Besides Byzantine faults, some works also consider other fault models and a combination of these faults [39]. However, few papers have addressed the convergence of MSR algorithms in less restrictive (non-complete) networks. Some exceptions include [40–42]. In [40], the authors studied *local convergence* (convergence of a subset of nodes) in undirected regular graphs[6]; the results are extended to asynchronous networks in [42] and global convergence of a class of undirected regular graphs, named *Partially Fully Connected Networks (PFCN)*, in [41]. More recently, [22] 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 subsections, 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 fault models described in Section 2.2.3.

2.4 Network Robustness

While network connectivity has been the key metric for studying robustness of distributed algorithms (as described in Section 1.1), the following proposition suggests that in general

[6] A regular graph is a graph where each vertex has the same number of neighbors.

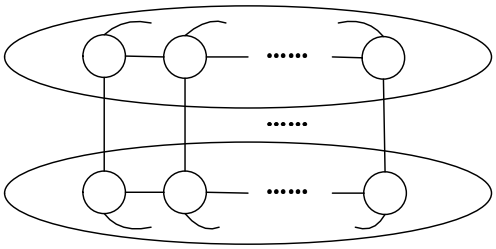

Figure 2. Example of a $\frac{n}{2}$-connected graph which fails to reach consensus.

networks, connectivity is no longer useful for characterizing the behavior of purely local algorithms, such as W-MSR.

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 Appendix A. Figure 2 illustrates an example of this kind of graph with $F = 1$. This network is undirected and there are two complete subgraphs (the upper and lower sets) each with degree $\frac{n}{2}$ (suppose n is even). Each node in the upper set has one and only one neighbor from the lower set. Note that the graph is $\frac{n}{2}$-connected and has minimum degree $\frac{n}{2}$. Suppose that nodes in the upper and lower sets have initial values a and b, respectively. When $a \neq b$, consensus will not be reached by using the W-MSR algorithm with parameter $F = 1$. This is because each node will throw away the value of its neighbor from the opposite set and thus its own value will remain unchanged, even when there are no misbehaving nodes.

Thus, even a relatively large connectivity (or minimum in-degree) in graphs 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. 2, 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 novel graph-theoretic properties termed *reachable sets* and *network robustness*. We also provide a variation on these properties for the f-fraction local model.

Definition 8 (r-reachable set)**.** *Given a graph $\mathcal{D}$ 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 graph $\mathcal{D}$ 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 $|\mathcal{V}_i \backslash \mathcal{S}| \geq p|\mathcal{V}_i|$, 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 graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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.*

Definition 11 (p-fraction robustness)**.** *A graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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 $\mathcal{V}$, at least one of the subsets is p-fraction reachable.*

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. 2. If either the upper or the lower set is 2-reachable ($r = F + 1 = 2$), then at least one node would be sufficiently influenced by a node outside of its set (because each node only removes up to $F = 1$ 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. In the next subsection, we will show that network robustness is essential to characterize the performance of the W-MSR algorithm.

2.5 Resilient Consensus Using Only Local Information

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 model, or with parameter f under the f-fraction local Byzantine 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 the update rule in (2) is a convex combination of these values. □

Note that since the malicious model is a special case of the Byzantine model, the above result also holds for related malicious models. 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 2.2.3.

2.5.1 F-Local and f-Fraction Local Malicious Models

The following key result provides a condition on the network that will guarantee that the algorithm achieves resilient consensus. The proof of Theorem 1 can be found in Appendix B.

Theorem 1. *Consider a time-invariant network modeled by a graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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.*

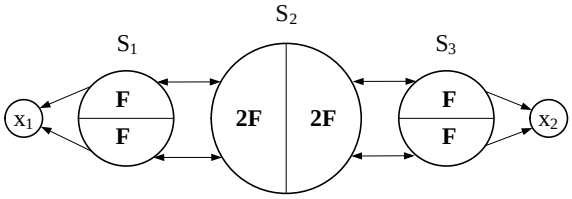

Figure 3. Illustration of Proposition 2

Although the sufficient and necessary conditions in Theorem 1 do not coincide, the following results show that both the sufficient and necessary conditions in the theorem are *sharp*, i.e., relaxing the sufficient condition leads to examples in which consensus is not achieved and there exist graphs satisfying the necessary condition where consensus is guaranteed.

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 under the F-local malicious model.*

Proof. We will prove the result by giving a construction of such a graph, visualized in Figure 3. In Figure 3, $\mathcal{S}_1$, $\mathcal{S}_2$ and $\mathcal{S}_3$ are all complete subgraphs with $|\mathcal{S}_1| = |\mathcal{S}_3| = 2F$ and $|\mathcal{S}_2| = 4F$. Each node in $\mathcal{S}_1$ connects to $2F$ nodes of $\mathcal{S}_2$ and each node in $\mathcal{S}_3$ connects to the other $2F$ nodes of $\mathcal{S}_2$, and all these connections are undirected.[7] 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 $\mathcal{S}_3$. 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. □

[7] This is an example of a graph that arises from the construction that we will derive in Section 4.4, where we will show that such a graph is $2F$ robust.

Proposition 3. *For every $F \in \mathbb{Z}_{>0}$, there exists a $(F+1)$-robust network where resilient consensus is achieved using the W-MSR algorithm with parameter F under the F-local malicious model.*

Proof. For simplicity, we focus on the case when n is even and construct an *undirected* graph which is similar to the one constructed in the proof of Proposition 1. Let $\mathcal{X}$ and $\mathcal{Y}$ be two complete graphs on $\frac{n}{2}$ nodes. Number nodes in $\mathcal{X}$ and $\mathcal{Y}$ as $x_1, x_2, \ldots, x_{\frac{n}{2}}$ and $y_1, y_2, \ldots, y_{\frac{n}{2}}$, respectively. For any node $x_i \in \mathcal{X}$, if $i \leq |\mathcal{Y}| - F$, connect x_i with nodes $y_i, y_{i+1}, \ldots, y_{i+F}$; otherwise, connect x_i with nodes $y_i, \ldots, y_{\frac{n}{2}}$ and nodes $y_1, \ldots, y_{i+F-\frac{n}{2}}$. Then each node in $\mathcal{X}$ and $\mathcal{Y}$ has exactly $F+1$ neighbors in the other set and thus, the graph is $(F+1)$-robust. Note that the minimum degree of this graph is $\frac{n}{2} + F$. Further note that, the F-local model is equivalent to the F-total model in this graph, i.e., there does not exist a F-local set with more than F nodes. In [43], the authors have shown that if the minimum degree of the graph is $\lfloor \frac{n}{2} \rfloor + F$, then consensus can be reached under the F-total malicious model; by using this result, we complete the proof. □

While the above discussions have been for an underlying time-invariant network $\mathcal{D}$, it is relatively straightforward to extend the results to time-varying networks as follows. The proof of Corollary 1 can be found in Appendix C.

Corollary 1. *Consider a time-varying network modeled by a graph $\mathcal{D}[t] = \{\mathcal{V}, \mathcal{E}[t]\}$ 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 $c \in \mathbb{Z}_{>0}$.*

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

Theorem 2. *Consider a time-invariant network modeled by a graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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 1. For the proof of sufficiency, note that under the f-fraction local model, each normal node will disregard at most $2 \times \lfloor f d_i \rfloor$ 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 $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ and $\mathcal{X}_m(t_\epsilon, \epsilon_0)$ defined in the proof of Theorem 1 will adopt some normal node's value from outside. □

Corollary 2. *Consider a time-varying network modeled by a graph $\mathcal{D}[t] = \{\mathcal{V}, \mathcal{E}[t]\}$ 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 $c \in \mathbb{Z}_{>0}$.

2.5.2 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 [22] investigates a similar algorithm in the context of F-total Byzantine faults, 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 12. *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 edges among the normal nodes.*

Theorem 3 ([22])**.** *Consider a time-invariant network modeled by a graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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 [22, 44], we can also use the approach proposed in the proof of Theorem 1. Consider the two disjoint sets $\mathcal{X}_M(t_\epsilon, \epsilon_i)$ and $\mathcal{X}_m(t_\epsilon, \epsilon_i)$ defined in the proof of Theorem 1. If the normal network is $(F+1)$-robust, then one of the two sets (or both) contains some normal node which has at least $F+1$ normal neighbors outside.

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 [22] to the local models and time-varying networks.

Corollary 3. *Consider a time-invariant network modeled by a graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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 3. 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 network is f-fraction robust, we can choose two sets which are f-fraction reachable and these two sets contain some node i which has $\lceil fd_i \rceil$ neighbors outside. 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 4. *Consider a time-varying network modeled by a graph $\mathcal{D}[t] = \{\mathcal{V}, \mathcal{E}[t]\}$ 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 the normal network of $\mathcal{D}[t]$ is either (i) $(F+1)$-robust, or (ii) p-fraction robust, where $f < p \leq 1$. Then, under the (i) F-local Byzantine model, or (ii) 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 $c \in \mathbb{Z}_{>0}$.*

Remark 3. *Note that when the original network is $(2F+1)$-robust (or p-fraction robust, where $2f < p \leq 1$), the normal network will be $(F+1)$-robust (or p-fraction robust, where $f < p \leq 1$). Thus, the results in Section 2.5.1 also hold for related Byzantine models. Further note that, the necessary conditions described in this section do not apply for the related malicious models, which implies that the restriction of the ability of misbehaving nodes results in extra complexity. Thus, the results in Section 2.5.1 are nontrivial.*

3. Extensions of Network Robustness

3.1 Introduction

In the previous section, we introduced the concept of network robustness and showed that this concept is the key property to characterize the performance of algorithms using only local information, such as W-MSR. In this section, we will study two extensions of network robustness – (r, s)-*robustness* and *strong robustness*. We first introduce the concept of (r, s)-robustness, where s represents the total number of nodes in a pair of sets that each have at least r neighbors outside their own set and characterizes another type of information redundancy; we will provide a *necessary and sufficient* condition for the W-MSR algorithm to achieve resilient consensus under the F-total malicious model using this concept.[8] Then we will turn to fault tolerant broadcast, which is another important operation in networks, and use strong robustness to show that broadcast will succeed in certain networks that do not meet the conditions studied before.

3.2 (r, s)-Robustness for Resilient Consensus

Consider the network modeled by the graph in Fig. 4. One can verify that the graph is 3-robust by checking every possible

[8] This part of work is done in conjunction with H. LeBlanc and X. Koutsoukos from Vanderbilt University.

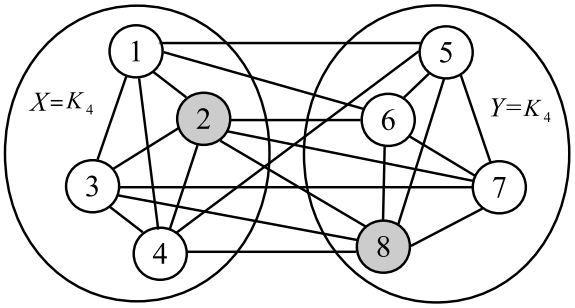

Figure 4. A 3-robust graph in which sets X and Y are each 3-reachable. Nodes 2 and 8 are malicious (shown in grey).

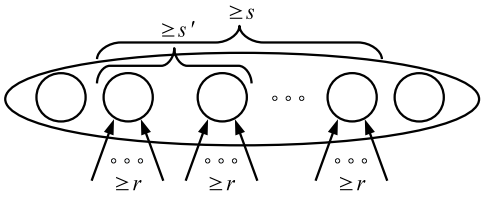

Figure 5. Illustration of an (r, s)-reachable set of nodes.

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. Given a set of nodes, this type of redundancy may also reduce the requirement on the number of neighbors from outside each set. In order to capture this intuition, we define the following concept.

Definition 13 ((r, s)-reachable set)**.** *Given a graph $\mathcal{D}$ 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 $\mathcal{X}_{\mathcal{S}} = \{i \in \mathcal{S} \colon |\mathcal{V}_i \setminus \mathcal{S}| \geq r\}$, then $|\mathcal{X}_{\mathcal{S}}| \geq s$.*

An illustration of an (r, s)-reachable set of nodes is shown in Fig. 5. 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 14 ((r, s)-robustness)**.** *A graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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 $\mathcal{V}$ such that $\mathcal{S}_1$ is $(r, s_{r,1})$-reachable and $\mathcal{S}_2$ is $(r, s_{r,2})$-reachable with $s_{r,1}$ and $s_{r,2}$ maximal (i.e., $s_{r,k} = |\mathcal{X}_{\mathcal{S}_k}|$ where $\mathcal{X}_{\mathcal{S}_k} = \{i \in \mathcal{S}_k \colon |\mathcal{V}_i \setminus \mathcal{S}_k| \geq r\}$ for $k \in \{1, 2\}$), then at least one of the following hold:*

- $s_{r,1} = |\mathcal{S}_1|$;
- $s_{r,2} = |\mathcal{S}_2|$;
- $s_{r,1} + s_{r,2} \geq s$.

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 $s_{r,k} = |\mathcal{X}_{\mathcal{S}_k}|$, 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 parameter[9] $1 \leq s \leq n$ quantifies what is meant by "enough" nodes.

In the rest of this subsection, we will show that these concepts we have proposed above are the key properties needed to characterize the performance of the W-MSR algorithm under the F-Total Malicious Model. We will also explore properties of of (r, s)-robust graphs.

3.2.1 F-Total Malicious Model

The following result provides a *necessary and sufficient* condition for the W-MSR algorithm to succeed under the F-total malicious model.

Theorem 4. *Consider a time-invariant network modeled by a graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ 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.*

[9] Note that $s = 0$ is *not* allowed in (r, s)-robustness because in that case any graph 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 $s > n$.

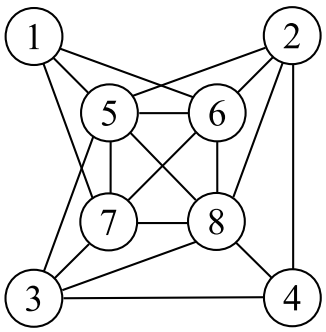

Figure 6. A 3-robust graph that is *not* (3,2)-robust.

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 $s_{F+1,1} + s_{F+1,2} \leq F$, suppose all nodes in $\mathcal{X}_{\mathcal{S}_1}$ and $\mathcal{X}_{\mathcal{S}_2}$ 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 $s_{F+1,1} < |\mathcal{S}_1|$ and $s_{F+1,2} < |\mathcal{S}_2|$, 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) The proof of sufficiency is similar to the proof of Theorem 1. Note that here we need to modify the definitions of $\mathcal{X}_M(t_\epsilon, \epsilon_i)$ and $\mathcal{X}_m(t_\epsilon, \epsilon_i)$ defined in the proof of Theorem 1 to be $\mathcal{X}_M(t, \epsilon_i) = \{i \in \mathcal{V} : x_i[t] > A_M - \epsilon_i\}$, which includes all normal and malicious nodes that have values larger than $A_M - \epsilon_i$, and $\mathcal{X}_m(t, \epsilon_i) = \{i \in \mathcal{V} : x_i[t] < A_m + \epsilon_i\}$, which includes all normal and malicious nodes that have values smaller than $A_m + \epsilon_i$. 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. □

This result establishes the notion of (r, s)-robustness introduced in Definition 14 as the appropriate metric for reasoning about purely local distributed algorithms, supplanting the traditional metric of connectivity. When the network is time-varying, one can state the following corollary of the above theorem.

Corollary 5. *Consider a time-varying network modeled by a graph $\mathcal{D}[t] = \{\mathcal{V}, \mathcal{E}[t]\}$ 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 $c \in \mathbb{Z}_{>0}$.*

To illustrate these results consider the graph in Figure 6. This graph can withstand the compromise of $F = 1$ malicious node in the network using the W-MSR algorithm with parameter $F = 1$ (the 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 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. 6).

3.2.2 Properties of (r, s)-Robust Graphs

Now we explore more properties of (r, s)-robust graphs. We begin with the important observation that $(r, 1)$-*robustness is equivalent to r-robustness*. This holds because conditions $(i) - (iii)$ in Definition 14 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 graphs.

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

Proof. 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 14 holds. Observe that $s_{r',k} \geq s_{r,k}$ for $k = 1, 2$. Hence if (i) or (ii) hold, then $s_{r',k} \geq s_{r,k} = |\mathcal{S}_k| \geq s_{r',k}$, which implies $s_{r',k} = |\mathcal{S}_k|$. If (iii) holds, then

$$s_{r',1} + s_{r',2} \geq s_{r,1} + s_{r,2} \geq s \geq s'.$$

Thus, any pair of nonempty, disjoint subsets of nodes in $\mathcal{D}$ satisfy Definition 14 with r and s replaced by r' and s'. Therefore, $\mathcal{D}$ is (r', s')-robust. □

It follows from Lemma 2 that a graph is r-robust whenever it is (r, s)-robust. The converse, however, is not true. Consider the graph in Fig. 6. 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 subgraph 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.*

Proof. Suppose $\mathcal{D}'$ is not (r, s)-robust. Then there exists a pair of nonempty, disjoint subsets $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ that are $(r, s_{r,k})$-reachable with maximal $s_{r,k}$ for $k \in \{1, 2\}$, but all of conditions (i)-(iii) in Definition 14 fail to hold. By removing directed edges in $\mathcal{E}''$, the magnitude of $s_{r,1}$ and $s_{r,2}$ 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. □

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

Lemma 4 (Minimum In-Degree)**.** *Given an (r, s)-robust graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$, with $0 \leq r \leq \lceil n/2 \rceil$ and $1 \leq s \leq n$, the minimum in-degree of $\mathcal{D}$, $\delta^{in}(\mathcal{D})$, is at least*

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

Proof. Whenever $r = 0, 1$, there is nothing to prove. Therefore, assume $2 \leq r \leq \lceil n/2 \rceil$, and fix $j \in \mathcal{V}$. First, let $\mathcal{S}_1 = \{j\}$ and $\mathcal{S}_2 = \mathcal{V} \backslash \mathcal{S}_1$. Then, $s_{r,2} = 0$ so that $s_{r,1} = |\mathcal{S}_1|$. This proves $\mathcal{V}_j \geq r$. 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 $s_{r,2} = 0$ so that $s_{r,1} = |\mathcal{S}_1|$. This implies j has an additional r in-neighbors outside of $\mathcal{S}_1$, thereby guaranteeing $\mathcal{V}_j \geq r+s-1$. 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 $\mathcal{S}_2 = \mathcal{V} \backslash \mathcal{S}_1$. Since $|\mathcal{S}_1| < r$ and $s \geq r$, again $s_{r,2} = 0$ so that $s_{r,1} = |\mathcal{S}_1|$. This implies j has an additional r in-neighbors outside of $\mathcal{S}_1$, thereby guaranteeing $\mathcal{V}_j \geq 2r-2$. Since $j \in \mathcal{V}$ is arbitrary, the bound on $\delta^{\text{in}}(\mathcal{D})$ follows. □

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

Lemma 5 (Directed Edge Removal)**.** *Given an (r, s)-robust (p-fraction robust) graph $\mathcal{D}$, let $\mathcal{D}'$ be the graph 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).*

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. □

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 directed spanning tree. The following result shows that 1-robustness is equivalent to containing a directed spanning tree.

Lemma 6. *A graph $\mathcal{D}$ is 1-robust if and only if $\mathcal{D}$ contains a directed spanning tree.*

Proof. If $\mathcal{D}$ is 1-robust, we will prove that $\mathcal{D}$ has a directed spanning tree by contradiction. Assume that $\mathcal{D}$ does not have a directed spanning tree. Decompose $\mathcal{D}$ into its strongly connected components, and note that since $\mathcal{D}$ does not have a directed spanning tree, 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 directed spanning tree.

Assume $\mathcal{D}$ contains a directed spanning tree, 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 directed out-branching, completing the proof. □

Remark 4. *The proof of Lemma 6 is a more direct and simpler version of the proof of Theorem 5 in [29].*

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

Lemma 7 (Connectivity of Robust Graphs)**.** *Suppose $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ is an r-robust graph, 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.*

Proof. If $r = 0$, the first statement is vacuously true, and if $r = 1$, it holds by Lemma 6. 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 6, 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 $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_k$. 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 $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_k$. 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 4, so that $|\mathcal{C}_i| \geq \lfloor r/2 \rfloor + 1$, $i \in \{1, \ldots, k\}$ (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 $s_{r,1} + s_{r,2} \geq r$. 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 $|\mathcal{K}_2| + |\mathcal{K}_3| \leq r-1$, $\exists j \in \mathcal{C}_2 \cup \cdots \cup \mathcal{C}_k$ such that $(j, v) \in \mathcal{E}$, which contradicts the fact that $\mathcal{K}$ is a vertex cut whose removal results in components $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_k$. Hence, $\mathcal{G}_\mathcal{D}$ is at least $(\lceil 3r/2 \rceil - 1)$-connected. □

3.3 Strong Robustness for Resilient Broadcasting

Having characterized the behavior of the consensus algorithm in terms of the network robustness, we now turn our attention to another important objective in networks: broadcasting a single value throughout the network. We focus on the following problem, studied in [23, 45]. Consider a time-invariant communication network $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$, with a specially designated source node $s \in \mathcal{V}$. The source has a value $x_s[0]$ that it wishes to broadcast to every other node in the network. However,

there may be various misbehaving nodes scattered throughout the network that wish to prevent certain nodes from obtaining the correct value of the source. To achieve broadcast (i.e., all normal nodes receive the source's message) under the F-local model, [45] proposes the following so-called *Certified Propagation Algorithm (CPA)*:

- At time-step 0, the source broadcasts its value to all of its neighbors, and maintains its value for all subsequent time-steps.
- At time-step 1, all normal neighbors of the source receive the source's value and broadcast it to all of their neighbors. The normal neighbors of the source maintain this value for all subsequent time-steps.
- At each time-step t, if a normal node has received an identical value from $F+1$ neighbors, then it accepts that value and broadcasts it to all of its neighbors. This normal node keeps this value for all subsequent time-steps.

Under the F-local model, it is easy to see that a normal node will only ever accept a value if it is the actual value broadcast by the source. For CPA, the following result from [23] provides a sufficient condition for all normal nodes in the network to eventually accept the value broadcast by the source.

Theorem 5 ([23]). *For a graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ and nodes $v, s \in \mathcal{V}$, let $X(v,s)$ denotes the number of nodes that are in v's neighborhood and are closer to s than v. Let $X(\mathcal{D}) = \min\{X(v,s) | v, s \in \mathcal{V}, (v,s) \notin \mathcal{E}\}$. Then CPA succeeds if $X(\mathcal{D}) > 2F$.*

This is only a sufficient condition; we will now provide a different sufficient condition for CPA to succeed, in terms of the robust-graph property that we have defined. We will first introduce another variation of the concept of an r-robust graph.

Definition 15 (Strongly r-robustness). *For a positive integer r, a graph $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ is **strongly r-robust** if for any nonempty subset $\mathcal{S} \subseteq \mathcal{V}$, either $\mathcal{S}$ is r-reachable or there exists a node $i \in \mathcal{S}$ such that $\mathcal{V}_i = \mathcal{V} \setminus \mathcal{S}$.*

Note that the difference between a strongly r-robust graph and the standard r-robust graph is that the former requires *every* subset of nodes to be either r-reachable, or have a node that connects to every node outside the set, whereas the latter only requires that one of any two sets satisfies the property of being r-reachable. Any strongly r-robust graph is r-robust, but not vice versa.

Theorem 6. *Under the F-local Byzantine (or malicious) model, CPA succeeds for any source if the network is strongly $(2F+1)$-robust.*

Proof. All normal neighbors of the source receive the message directly, and thus they all accept it. We will use contradiction to prove that all other nodes receive the broadcast message. Suppose that CPA fails to deliver the message to all normal nodes, and let $\mathcal{S}$ denote the set of all such normal nodes. By the definition of a strongly $(2F+1)$-robust graph, we know that some node i in $\mathcal{S}$ must have $2F+1$ neighbors outside $\mathcal{S}$ or connects to all nodes outside. For the former situation, at most F of these nodes can be malicious, and all other nodes are normal nodes that have received the message and re-broadcasted it; for the latter, this node would directly connect to the source and thus get the message. In either case, this contradicts the assumption that node i would fail to get the message, and thus the algorithm achieves broadcast. □

The following Proposition shows that CPA succeeds in certain networks which do not satisfy the condition proposed in Theorem 5.

Proposition 4. *For some F, there exist graphs with $X(\mathcal{D}) \leq 2F$ but that are strongly $(2F+1)$-robust.*

Proof. For $F = 1$, construct an undirected graph $\mathcal{G}$ as follows. Start with a fully-connected graph of five nodes, denoted $1, 2, 3, 4, 5$. Add two nodes 6 and 7 and connect them to nodes $2, 3, 4$ and $3, 4, 5$ respectively. Finally, add a node 8 and connect it to nodes $3, 4, 6, 7$. If we take node 1 as the source, it's easy to check that in the neighborhood of node 8, there are only two nodes that are closer to the source. Thus $X(\mathcal{G}) \leq 2F$ here, but the graph is still strongly $(2F+1)$-robust, and CPA will succeed. □

Finally, it will be of interest to note that the notions of r-reachable sets and and p-fraction reachable sets are similar to the notion of 'clusters', which are topological structures identified in [3] as being impediments to information cascades in networks. While the topic of information cascades (and more specifically, behavior adoption) is closely related to the problem that we considered in this subsection, the presence of misbehaving nodes in our setup significantly complicates the analysis; in the rest of this subsection, we provide more discussions on the differences between these two scenarios.

3.3.1 Behavior Adoption

We first give some background on the behavior adoption problem discussed in [3]. Consider a fixed (time-invariant) social network where every node i can adopt one of two possible behaviors at any given time-step t; we model this by saying that $\mathbf{x}[t] \in \{0, 1\}$ for all $t \in \mathbb{N}$. At each time-step t, each node i plays a game against each of its neighbors, where the strategy set available to each node is $\{0, 1\}$ (i.e., it chooses how to behave at that time-step), and the payoffs are as follows: node i and its neighbor j both get a payoff of a if both play strategy 0, both get a payoff of b of both play strategy 1, and they get no payoff if they play different strategies. The total payoff of any node i is the sum of the payoffs from the games played against each of its neighbors. Let $A_i[t]$ denote the number of

neighbors of node i that are playing strategy 0 at time-step t, and let $B_i[t] = \deg_i - A_i[t]$ be the number of neighbors that are playing strategy 1. In this case, it was shown in [3] that there exists a constant q satisfying $0 < q < 1$ (and depending on the payoffs in the game) such that the best strategy for each node to follow is as follows:

$$x_i[t+1] = \begin{cases} 0 & \text{if } \frac{A_i[t]}{\deg_i} > q, \\ 1 & \text{otherwise.} \end{cases} \quad (3)$$

In other words, each node should follow a *threshold* rule: if a sufficiently large fraction of neighbors plays strategy 0, the node should also play 0, and it should play 1 otherwise. Based on this interaction rule, the authors studied the following problem: suppose that all nodes in the network start out playing strategy 1, except for a small subset $\mathcal{A}$ of nodes that start playing strategy 0, and commit to 0 forevermore (i.e., $x_i[0] = 1$ for all $i \in \mathcal{V} \setminus \mathcal{A}$, and $x_i[t] = 0$ for all $i \in \mathcal{A}$ and for all $t \in \mathbb{N}$). Under what conditions will every node in the network eventually adopt strategy 0? The answer depends on the nature of the set $\mathcal{A}$ and the topology of the network, and is given by the following theorem.

Theorem 7 ([3]). *All nodes in the network will eventually switch to strategy* 0 *if and only if the rest of the graph does not contain a set of nodes where each node (in this set) has at least a fraction* $1 - q$ *of their neighbors inside the set.*

At this point, we highlight the following comparisons between the fault-tolerant broadcasting scenario and the behavior adoption scenario described previously. In both cases, there is a small set of nodes in the network that wish to have their values adopted by all other nodes. Furthermore, both scenarios feature *threshold*-based interactions; in the broadcasting case, the value is accepted if more than a fixed number of neighbors carry that value, whereas in the behavior adoption case, the value is adopted if more than a fixed *fraction* of the neighbors carry the value. However, there is one key difference between the two scenarios: there is no notion of a misbehaving node in the behavior adoption scenario. In other words, each node executes only the strategy 0 or 1, and each node follows the prescribed dynamics. The presence of misbehaving nodes in the broadcasting scenario complicates matters, and prevents us from directly applying the result in [3].

4. Robustness of Complex Networks

4.1 Introduction

In Section 2.5 (see Remark 3), we show that $(2F+1)$-robustness is sufficient to achieve resilient asymptotic consensus under the F-local Byzantine model. Since the F-total Byzantine and malicious models and the F-local malicious model are special cases of the F-local Byzantine model, the condition of being $(2F+1)$-robust is sufficient for those fault models as well. Although the robustness condition presented in Theorem 1 is not necessary for all of these fault models, we will show later in this section that this metric is conducive for deriving threshold functions under which random graphs will be resilient to all four fault models.[10] Furthermore, these threshold functions will 'coincide' with the threshold functions for $(2F+1)$-connectivity (which is a necessary condition for all four models); in other words, we sandwich the various conditions for resilient consensus by our sufficient condition ($(2F+1)$-robustness) and the fundamental necessary condition ($(2F+1)$-connected), thereby implicitly providing threshold functions for those conditions as well. Although we can construct 'worst-case' networks where even very large connectivity can not guarantee sufficient robustness, we show that the story is different in complex networks. Specifically, in the rest of this section, we will demonstrate that in complex networks, connectedness and robustness cannot differ too much (the meaning will be clear in the following subsections).

4.2 Robustness of Erdős-Rényi Random Graphs

In this subsection, we study robustness in Erdős-Rényi Random graphs [46–48], one of the most common models for large-scale complex networks. Erdős and Rényi proposed a number of versions of their model and the most commonly studied is the one denoted as $\mathcal{G}_{n,p}$. In this model, the graph consists of n vertices and each possible (undirected) edge between two vertices is present with independent probability p (which may be a function of n), and absent with probability $q = 1 - p$. Let the probability of an event be denoted by $\mathbb{P}(\cdot)$. Recall that a *graph property* can be regarded as a class of graphs that is closed under isomorphism. A key feature of the $\mathcal{G}_{n,p}$ model is that we can explore properties that are shared by *almost all* graphs, a notion that is defined as follows.

Definition 16. *Assume $\mathcal{P}$ is a graph property and $p = p(n)$ is a fixed function (possibly constant). We say that* ***almost all*** *$G \in \mathcal{G}_{n,p}$ have the property $\mathcal{P}$ if $\mathbb{P}(G \in \mathcal{P}) \to 1$ as $n \to \infty$; and* ***almost no*** *$G \in \mathcal{G}_{n,p}$ has the property $\mathcal{P}$ if $\mathbb{P}(G \in \mathcal{P}) \to 0$ as $n \to \infty$.*

We are interested in the evolution of random graphs as p evolves and normally we consider the case when $p \to 0$. One important feature of $\mathcal{G}_{n,p}$, which was demonstrated by Erdős and Rényi, is that the model shows a 'phase transition' phenomenon. More precisely, we define a *threshold function* as follows.

Definition 17 (Threshold Function). *A* ***threshold function*** *for a graph property $\mathcal{P}$ is a function $t(n)$ such that $p(n) = o(t(n))$ implies that almost no $G \in \mathcal{G}_{n,p}$ has the property $\mathcal{P}$ and $t(n) = o(p(n))$ implies that almost all $G \in \mathcal{G}_{n,p}$ have the property $\mathcal{P}$.*

Note that the assumption that $p(n) = o(t(n))$ or $t(n) = o(p(n))$ is conservative in some cases and we define the following property for threshold functions.

[10] Namely, the F-total Byzantine, F-local Byzantine, F-total malicious and F-local malicious models.

Definition 18 (Sharp Threshold Function)**.** *A threshold function $t(n)$ with the form $\frac{g(n)}{n}$ for a graph property $\mathcal{P}$ is **sharp** if $p(n) = \frac{g(n)+x}{n}$ implies that almost no $G \in \mathcal{G}_{n,p}$ has the property $\mathcal{P}$ and $p(n) = \frac{g(n)-x}{n}$ implies that almost every $G \in \mathcal{G}_{n,p}$ has the property $\mathcal{P}$, where $g(n)$ is some function of n, $x = o(g(n))$ and $x \to \infty$ as $n \to \infty$.*

Note that the reason we consider the threshold function with the form $\frac{g(n)}{n}$ is that all properties we are going to study in this section have this form. In the rest of this subsection, when we refer to a threshold function, we will use Definition 17 if we do not emphasis its sharpness. Now we focus on the following properties.

Definition 19. *For $G \in \mathcal{G}_{n,p}$ and constants $r, F \in \mathbb{Z}_{\geq 1}$, define the properties of being r-**connected**, r-**robust** and **achieving resilient asymptotic consensus (RAC)** using the W-MSR algorithm with parameter F (under the F-local/total Byzantine/malicious models) by $\mathcal{C}_r$, $\mathcal{R}_r$ and $\mathcal{RAC}_F$, respectively.*

Lemma 8 ([47, 48])**.** *For any constant $r \in \mathbb{Z}_{\geq 1}$, $t(n) = \frac{\ln n + (r-1)\ln\ln n}{n}$ is a sharp threshold function for the property $\mathcal{C}_r$.*

The following is one of our main results: it establishes a threshold function for r-robustness in Erdős-Rényi random graphs.

Theorem 8. *For every constant $r \in \mathbb{Z}_{\geq 1}$, $t(n) = \frac{c \ln n}{n}$ is a threshold function for the property $\mathcal{R}_r$, where $c > 0$ is some constant.*

Proof. Let $\gamma \equiv \frac{p(n)}{t(n)}$. For the first part of the proof, we show that almost all $G \in \mathcal{G}_{n,p}$ are r-robust if $\gamma \to \infty$ as $n \to \infty$. Denote the probability that some set of cardinality up to $n_c = \lceil \frac{n}{2} \rceil$ is not r-reachable as $\mathbb{P}_0$ and the probability that some vertex set $\mathcal{S} \subset \mathcal{V}$ with cardinality k (i.e., $|\mathcal{S}| = k$) is not r-reachable as $\mathbb{P}_k$. By the union bound, we know that $\mathbb{P}_0 \leq \sum_{k=1}^{n_c} \mathbb{P}_k$. Note that if $\mathbb{P}_0 \to 0$, then all sets of cardinality up to $\lceil \frac{n}{2} \rceil$ are r-reachable and by the definition of robustness, graphs $G \in \mathcal{G}_{n,p}$ are r-robust.

For fixed $\mathcal{S}$, the probabilities that a vertex $v \in \mathcal{S}$ has less than r neighbors outside and $\mathcal{S}$ is not r-reachable are $\sum_{i=0}^{r-1} \binom{n-k}{i} q^{n-k-i} p^i$ and $(\sum_{i=0}^{r-1} \binom{n-k}{i} q^{n-k-i} p^i)^k$, respectively. Since there are $\binom{n}{k}$ such sets $\mathcal{S}$, we know that $\mathbb{P}_k \leq \binom{n}{k} (\sum_{i=0}^{r-1} \binom{n-k}{i} q^{n-k-i} p^i)^k$. We obtain the following upper bound for $\mathbb{P}_k$:

$$\begin{aligned}
\mathbb{P}_k &\leq \binom{n}{k} (\sum_{i=0}^{r-1} \binom{n-k}{i} q^{n-k-i} p^i)^k \\
&\leq n^k (\sum_{i=0}^{r-1} (n-k)^i q^{n-k-i} p^i)^k \\
&\leq (n(n-k)^{r-1} \sum_{i=0}^{r-1} q^{n-k-i} p^i)^k \\
&\leq (n(n-1)^{r-1} q^{n-k} \frac{1 - \left(\frac{p}{q}\right)^r}{1 - \frac{p}{q}})^k.
\end{aligned}$$

Note that $\binom{n}{k} \leq \frac{n^k}{k!} \leq n^k$. Let $f(k) \equiv n(n-1)^{r-1} q^{n-k} \frac{1-(\frac{p}{q})^r}{1-\frac{p}{q}}$ be a function of k. We can check that $f'(k) > 0$ for all $k \geq 1$. Thus, $f(k) \leq f(n_c)$ and

$$\begin{aligned}
\mathbb{P}_0 &\leq \sum_{k=1}^{n_c} \mathbb{P}_k \\
&\leq \sum_{k=1}^{n_c} f(n_c)^k \\
&= f(n_c) \frac{1 - f(n_c)^{n_c}}{1 - f(n_c)}.
\end{aligned}$$

Now we get the following approximation for $f(n_c)$:

$$\begin{aligned}
f(n_c) &= O\big(n(n-1)^{r-1} q^{n-n_c}\big) && (4) \\
&= O\big(n^r (1 - \frac{\gamma c \ln n}{n})^{n_c}\big) \\
&= O\big(n^r \exp\{\frac{-\gamma c}{2} \ln n\}\big) && (5) \\
&= O\big(n^{r - \frac{\gamma c}{2}}\big).
\end{aligned}$$

Note that (4) is due to the fact that $\frac{1-(\frac{p}{q})^r}{1-\frac{p}{q}} \to 1$ as[11] $p \to 0$ and (5) is due to the fact that $1 - m < e^{-m}$ for $m > 0$. Thus, we know that $f(n_c) \to 0$ as $n, \gamma \to \infty$ and the first part of the proof is obtained by noting that $\mathbb{P}_0 \to 0$ as $n, \gamma \to \infty$.

For the second part of the proof, we need to show that almost no $G \in \mathcal{G}_{n,p}$ are r-robust if $\gamma \to 0$ as $n \to \infty$. By Lemma 7 and Lemma 8, we know that if $\gamma \to 0$ as $n \to \infty$, almost no $G \in \mathcal{G}_{n,p}$ are r-connected and thus are not r-robust, completing the proof. □

Note that in the first part of the proof, for fixed r, $\gamma \to \infty$ is not necessary and we just need $\gamma c > 2r$. From Theorem 8, we know that any function $t(n) = \Theta(\frac{\ln n}{n})$ is a threshold function for the property $\mathcal{R}_r$ and since $\frac{\ln n + (r-1)\ln\ln n}{n} = \Theta(\frac{\ln n}{n})$, the threshold function for r-connectivity is also a threshold function for r-robustness. Note that in [49], we

[11] Note if $p \not\to 0$, then p (and thus q) is some constant and $f(n_c) \to 0$ as $n \to \infty$, and we will still get the same result.

have shown that the threshold function for r-connectivity is actually a *sharp* threshold function for r-robustness.

Finally, we characterize the threshold function for the W-MSR algorithm to achieve resilient asymptotic consensus in the $\mathcal{G}_{n,p}$ model.

Theorem 9. *For every constant $F \in \mathbb{Z}_{\geq 1}$, $t(n) = \frac{c \ln n}{n}$ is a threshold function for the property $\mathcal{RAC}_F$, where $c > 0$ is some constant.*

Proof. As discussed in Section 4.1, Theorem 1 and Remark 3, $(2F+1)$-connectivity is necessary and $(2F+1)$-robust is sufficient, respectively, for the W-MSR algorithm to achieve resilient asymptotic consensus under the F-local/total Byzantine/malicious models. Thus, by Lemma 8 and Theorem 8, the result follows. □

4.3 Robustness of Geometric Random Graphs

In addition to the Erdős-Rényi model, another widely used model is the *geometric random graph*, which captures edges between nodes that are in close (spatial) proximity to each other. In Section 4.2, we showed that the properties of connectedness and robustness have the same threshold function in Erdős-Rényi graphs. In this subsection, we will prove similar results for geometric random graphs.

We consider the geometric random graph model $\mathcal{G}^d_{n,\rho,l}$, which is a random *undirected* graph generated by first placing n vertices at random (uniformly and independently) in a region $\Omega_d = [0,l]^d$, where $d = 1, 2, 3$. Two vertices in the graph are then connected by an edge if and only if the distance between them is at most a threshold ρ. In the more widely-studied model $\mathcal{G}^d_{n,\rho}$, in which the vertices are distributed on $[0,1]^d$, where $d = 1, 2, 3$, graph properties are typically explored when $n \to \infty$ and $\rho \to 0$ [50]. Thus, this model is more suitable for dense random networks [51]. In the more general model $\mathcal{G}^d_{n,\rho,l}$, however, the density $\frac{n}{l^d}$ can converge to 0 or some constant, making it suitable for capturing both dense and sparse random networks. Furthermore, the model $\mathcal{G}^d_{n,\rho,l}$ is more convenient for us to deal with.

We define properties for *almost all* graphs in $\mathcal{G}^d_{n,\rho,l}$ as follows, similar to the $\mathcal{G}_{n,p}$ model.

Definition 20. *Assume $\mathcal{P}$ is a graph property. We say that **almost all** $G \in \mathcal{G}^d_{n,\rho,l}$ have the property $\mathcal{P}$ if $\mathbb{P}(G \in \mathcal{P}) \to 1$ as $l \to \infty$; and **almost no** $G \in \mathcal{G}^d_{n,\rho,l}$ has the property $\mathcal{P}$ if $\mathbb{P}(G \in \mathcal{P}) \to 0$ as $l \to \infty$.*

Note that we study these properties in $\mathcal{G}^d_{n,\rho,l}$ as $l \to \infty$ and n and ρ are functions of l, i.e., $n = n(l)$ and $\rho = \rho(l)$. In the rest of this section, we focus on the one-dimensional case and consider the line $\Omega_1 = [0,l]$. We start by providing a result showing that connectivity and robustness cannot be very different in one-dimensional geometric graphs (regardless of how they are generated).

Lemma 9. *In $\Omega_1 = [0,l]$, if the graph is $\lfloor \frac{3}{2} r \rfloor$-connected, then the graph is r-robust.*

Proof. Denote vertex i's value on the line $\Omega_1 = [0,l]$ by $x(i)$, $i \in \{1,\dots,n\}$. Without loss of generality, we assume if $i < j$, then $x(i) < x(j)$, $\forall i,j \in \{1,\dots,n\}$; otherwise, we can just renumber the vertices. When the graph is $\lfloor \frac{3}{2} r \rfloor$-connected, any interval $[x, x+\rho] \subseteq (x(1), x(n))$ contains at least $\lfloor \frac{3}{2} r \rfloor$ vertices; otherwise, removing the vertices in $[x, x+\rho]$ will disconnect the graph, which contradicts the assumption that the graph is $\lfloor \frac{3}{2} r \rfloor$-connected.

For a set $\mathcal{S}$ of vertices, we say the set of consecutive vertices $\{i_1,\dots,i_k\}$,[12] where $k \in \mathbb{Z}_{\geq 1}$, is a *cluster* of $\mathcal{S}$, denoted as $\mathcal{C}_\mathcal{S}$, if $\{i_1,\dots,i_k\} \subseteq \mathcal{S}$ and $i_1 - 1, i_k + 1 \notin \mathcal{S}$. Denote the i-th cluster of $\mathcal{S}$ by $\mathcal{C}^i_\mathcal{S}$ (which is ordered by the positions of the vertices inside the cluster). We say two disjoint clusters $\mathcal{C}^i_\mathcal{S} = \{i_1,\dots,i_k\}$ and $\mathcal{C}^j_\mathcal{S} = \{j_1,\dots,j_t\}$ are *connected* if either $|x(i_k) - x(j_1)|$ or $|x(i_1) - x(j_t)|$ is smaller than ρ. Denote $d(i,\rho)$ as the interval within distance ρ of $x(i)$, i.e., $d(i,\rho) = [x(i) - \rho, x(i) + \rho]$. We say a set $\mathcal{S}$ of vertices is a *full coverage* of the graph if $[x(1), x(n)] \subseteq \bigcup_{i \in \mathcal{S}} d(i,\rho)$.

For any set $\mathcal{S}$ of vertices, consider the following two cases:

- Case 1: there exist two clusters $\mathcal{C}^i_\mathcal{S}$ and $\mathcal{C}^{i+1}_\mathcal{S}$ which are not connected;
- Case 2: $\mathcal{S}$ is not a full coverage of the graph.

Assume $\mathcal{S} = \{i_1,\dots,i_k\}$. If Case 1 is true, there will be r vertices which are not in $\mathcal{S}$ belonging to $d(i_k,\rho)$ and $\mathcal{S}$ is r-reachable. If Case 2 is true, then either there exist two clusters of set $\mathcal{S}$ which are not connected, or $|x(i_1) - x(1)|$ or $|x(i_k) - x(n)|$ (or both) will be bigger than ρ and $\mathcal{S}$ is r-reachable.

Thus, consider any pair of sets $\mathcal{S}_1$ and $\mathcal{S}_2$; if either Case 1 or Case 2 is true for either of the two sets, then at least one of them is r-reachable. Otherwise, we can choose a vertex $i \in \mathcal{S}_2$ such that $d(i,\rho) \subset [x(1), x(n)]$.[13] Since Case 2 is not true for the set $\mathcal{S}_1$, there must exist at least one vertex of $\mathcal{S}_1$ in both $d(i,\rho) \cap \big(x(1), x(i)\big)$ and $d(i,\rho) \cap \big(x(i), x(n)\big)$. The situation is illustrated in Figure 7, where $j,k \in \mathcal{S}_1$ and $x(j) < x(i) < x(k)$. The interval $[x(j) - \rho, x(k) + \rho]$ can be divided into eight segments based on $x(j), x(i), x(k)$ and $d(j,\rho), d(i,\rho), d(k,\rho)$, and let $a, b, \dots, h$ be the number of vertices which are not in $\mathcal{S}_1$ within their related intervals, respectively. If $\mathcal{S}_1$ is not r-reachable, then $a+b+c+d+e+f < r-1$ and $c+d+e+f+g+h < r-1$. Thus, we get $a+b+2(c+d+e+f)+g+h < 2r-2$ and $b+c+d+e+f+g < 2r-2-(a+c+d+e+f+h)$. Let the number of vertices in $d(i,\rho)$ which are not in $\mathcal{S}_2$ be s. Then $s \geq 2\lfloor \frac{3}{2} r \rfloor - (b+c+d+e+f+g) \geq 3r-1-(2r-2)+(a+c+d+e+f+h) > r$. So if $\mathcal{S}_1$ is not r-reachable, $\mathcal{S}_2$ must be r-reachable.

Thus, if the graph is $\lfloor \frac{3}{2} r \rfloor$-connected, for any pair of sets, at least one of them is r-reachable and the graph is r-robust, completing the proof. □

[12] In the rest of the proof, when we refer to a set of vertices $\{i_1,\dots,i_k\}$ we assume it is ordered, i.e., $i_1 < \cdots < i_k$.

[13] Note that the case when $\big(x(i), x(n)\big) \leq 2\rho$ is trivial.

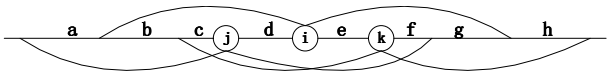

Figure 7. Illustration of the proof of Lemma 9.

Once again, note that the result in Lemma 9 does not depend on *how* the positions of the nodes are generated. Unfortunately, the proof of the lemma does not extend to geometric graphs in higher-dimensions. For example, the graph shown in Figure 2 can be viewed as a geometric graph in two dimensions, where the nodes in each set are all clustered horizontally within a distance ρ, and the two sets are vertically separated by a distance just below ρ, so that each node is within a distance ρ of exactly one node in the opposite set. Clearly that graph is only 1-robust, despite having a connectivity of $\frac{n}{2}$.

Next we will present an asymptotic approach to analyzing one-dimensional random graphs (complementary to the analysis in Lemma 9). We will be using the following result from [51].

Theorem 10 ([51]). *Assume that $\rho n = kl \ln l$ for some $k > 0$.*

- *If $k > 2$, or $k = 2$ and $\rho \to \infty$, then almost all $G \in \mathcal{G}^1_{n,\rho,l}$ are connected.*
- *If $k \leq (1-\epsilon)$ and $\rho \in \Theta(l^\epsilon)$ for some $0 < \epsilon < 1$, then almost no $G \in \mathcal{G}^1_{n,\rho,l}$ is connected.*

We now present the following conditions under which the one-dimensional random graph becomes r-robust.[14]

Theorem 11. *Assume that $\rho n = kl \ln l$ for some $k > 0$.*

- *If $k > \lfloor \frac{3}{2}r \rfloor + 1$, or $k = \lfloor \frac{3}{2}r \rfloor + 1$ and $\rho \to \infty$, then almost all $G \in \mathcal{G}^1_{n,\rho,l}$ are r-robust.*
- *If $k \leq (1-\epsilon)$ and $\rho \in \Theta(l^\epsilon)$ for some $0 < \epsilon < 1$, then almost no $G \in \mathcal{G}^1_{n,\rho,l}$ is r-robust.*

Proof. In order to prove the first part, we know it is sufficient to show that any interval of length ρ contains at least $\lfloor \frac{3}{2}r \rfloor$ vertices (as argued in the proof of Lemma 9). Let $\Omega_1 = [0, l]$ be subdivided into $c = \frac{(\lfloor \frac{3}{2}r \rfloor + 1)l}{\rho}$ non-overlapping segments of length $\frac{\rho}{\lfloor \frac{3}{2}r \rfloor + 1}$. Then any interval of length ρ will contain at least $\lfloor \frac{3}{2}r \rfloor$ whole segments and thus we just need to show every segment contains at least one vertex.

Let ω be a random variable representing the number of empty segments. Since ω is a nonnegative integer random variable, by Markov's inequality we know $\mathbb{P}(\omega > 0) < \mathbb{E}(\omega)$, where $\mathbb{E}(\omega) = c(1 - \frac{1}{c})^n$ is the expected value of ω. Since $1 - x < \exp(-x)$ for $x > 0$, we have

$$\begin{aligned}\mathbb{E}(\omega) &= c\big(1 - \frac{1}{c}\big)^n \\ &< c \exp\big(-\frac{n}{c}\big) \\ &= \frac{(\lfloor \frac{3}{2}r \rfloor + 1)l}{\rho} \exp\big(-\frac{k}{\lfloor \frac{3}{2}r \rfloor + 1} \ln l\big) \\ &= \frac{\lfloor \frac{3}{2}r \rfloor + 1}{\rho} l^{1 - \frac{k}{\lfloor \frac{3}{2}r \rfloor + 1}}.\end{aligned}$$

If $k > \lfloor \frac{3}{2}r \rfloor + 1$, or $k = \lfloor \frac{3}{2}r \rfloor + 1$ and $\rho \to \infty$, then $\mathbb{E}(\omega) \to 0$ as $l \to \infty$ and completing the proof for the first part.

The second part is obvious, because under the given conditions, Theorem 10 indicates that the graph will not be connected with high probability, and thus the result follows. □

4.4 Robustness of Preferential Attachment Networks

In this subsection, we provide a construction for (r, s)-robust graphs, and show that our construction contains the preferential-attachment model of scale-free networks as a special case [52]. Note that in the previous subsections we focused on undirected graphs, but in this subsection, we considered *directed* graphs.

Theorem 12. *Let $\mathcal{D} = \{\mathcal{V}, \mathcal{E}\}$ be an (r, s)-robust graph (with $s \geq 1$). Then the graph $\mathcal{D}' = \{\mathcal{V} \cup \{v_{new}\}, \mathcal{E} \cup \mathcal{E}_{new}\}$, 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 $d_{v_{new}} \geq r + s - 1$.*

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 $v_{\text{new}} \in \mathcal{S}_i$, $i \in \{1, 2\}$. In the first case, since $\mathcal{D}$ is (r, s)-robust, the conditions in Definition 14 must hold. In the second case, $\mathcal{X}_{\mathcal{S}_i} = \mathcal{S}_i$, 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: $s_{r,1} + s'_{r,2} \geq s$, $s_{r,1} = |\mathcal{S}_1|$, or $s'_{r,2} = |\mathcal{S}'_2|$. 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 $s'_{r,2} = |\mathcal{S}'_2|$ holds. Then, the negation of the first condition $s_{r,1} + s'_{r,2} \geq s$ implies $s'_{r,2} = |\mathcal{S}'_2| < s$. Hence, $|\mathcal{V}_{v_{\text{new}}} \setminus \mathcal{S}_2| \geq r$, and $s_{r,2} = |\mathcal{S}_2|$, completing the proof. □

The above result indicates that to construct an (r, s)-robust graph with n nodes (where $n > r$), we can start with an (r, s)-robust graph with relatively smaller order (such as some complete graph), and continually add new nodes with incoming edges from at least $r + s - 1$ nodes in the existing graph. The theorem does not specify *which* existing nodes should be chosen as neighbors. A particularly interesting case is when the nodes are selected with a probability proportional to the number of edges that they already have; this is known as *preferential-attachment*, and leads to the formation of so-called *scale-free* networks. Specifically, the construction process in Theorem 12 coincides with the *Barabási-Albert (BA) model* [52]: start with a network of k_0 nodes and add new nodes to the network one at a time; each new node connects to

[14] The proof of this result is inspired by the prior work on $\mathcal{G}^d_{n,\rho,l}$ in [51].

k existing nodes chosen by the preferential-attachment mechanism. The BA model is cited as a plausible mechanism for the formation of many real-world complex networks, and thus our analysis indicates that these networks will also be resilient to locally-bounded Byzantine or malicious nodes (provided that r is sufficiently large when the network is forming).

Recall that for resilient consensus, the fundamental necessary condition is $(2F+1)$-connectivity and we have shown the sufficiency of $(2F+1)$-robustness. In order to sandwich various conditions by connectivity and robustness, we now focus on r-robust graphs ($s=1$). Note that the network constructed in Theorem 12 is (at most) r-connected, since each node only connects to r existing nodes. In other words, in scale-free networks generated by this process, r-connected implies r-robustness. This leads to the following result.

Theorem 13. *In the BA model, when the initial network is r-robust, then the generated scale-free network is r-connected if and only if the network is r-robust.*

Proof. Note that in the BA model, if there exists some new node which connects to less than r existing nodes, then the network will be neither r-connected or r-robust; on the other hand, if all the new nodes connect to r existing nodes, then the network will be both r-connected and r-robust. □

5. Conclusions and Future Research

In this report, we have studied the problem of disseminating information in networks that contain misbehaving nodes, where each normal node has no knowledge of the global topology of the network. Under the assumption of full knowledge of the network topology by every node, it has been well established that connectivity is the key metric for success. However, we have shown that connectivity is no longer an appropriate metric for an algorithm that uses a purely local filtering strategy. Instead, we introduced the notions of network robustness and its variants, and showed that these concepts allow us to provide conditions for achieving the objectives of resilient distributed consensus and fault-tolerant broadcast.

While in the worst-case, networks with very large connectivity cannot guarantee sufficient robustness, we showed in this report that the notions of robustness and connectivity 'coincide' in complex networks. Specifically, we considered three common models for the complex networks. In Erdős-Rényi random graphs, we showed that connectedness and robustness share the same threshold functions. In geometric random graphs, we focused on the one-dimensional case and proved that $\frac{3}{2}r$-connected implies r-robust and robustness exhibits similar thresholds as connectivity. In preferential attachment networks, we showed that when the initial network is robust, connectivity and robustness are equivalent in the BA model. We also provided a construction method for robust graphs based on the preferential attachment mechanism.

There are many interesting directions for future research. First, the necessary and sufficient condition for resilient consensus under the F-local malicious model is still an open problem. Secondly, we mainly focus on the consensus dynamics in this report and it will be of interest to use the concept of network robustness (and its invariants) to capture the performance of other resilient algorithms. Another topic is to explore threshold results for geometric random graphs in higher dimensions. Finally, it will be of interest to relate strong robustness defined in this report to other recent characterizations of network topologies that facilitate fault-tolerant broadcast [24].

Just as the notion of connectivity has played a central role in the existing analysis of reliable distributed algorithms with global topological knowledge, we believe that robust graphs (and its variants) will play an important role in the investigation of purely local algorithms.

Acknowledgments

The authors would like to thank Heath J. LeBlanc and Xenofon Koutsoukos for the great discussions and collaboration and Nitin Vaidya et al. for pointing out related research of Approximate Byzantine Consensus. The authors are also thankful to the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Waterloo Institute for Complexity and Innovation (WICI) for sponsoring this work.

APPENDICES

1. Proof of Proposition 1

Proof. For simplicity, we focus on the case when n is even and 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 $x_1, x_2, \ldots, x_{\frac{n}{2}}$ and $y_1, y_2, \ldots, y_{\frac{n}{2}}$, respectively. For any node $x_i \in \mathcal{X}$, if $i \leq |\mathcal{Y}| - F + 1$, connect x_i with nodes $y_i, y_{i+1}, \ldots, y_{i+F-1}$; otherwise, connect x_i with nodes $y_i, \ldots, y_{\frac{n}{2}}$ and nodes $y_1, \ldots, y_{i+F-\frac{n}{2}-1}$. 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 $\mathcal{C}_\mathcal{X} = \{x_1, x_2, \ldots, x_{|\mathcal{C}_\mathcal{X}|}\}$; 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 $|\mathcal{C}_\mathcal{X}| > F$; 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 $F < |\mathcal{C}_\mathcal{X}| < \frac{n}{2}$, the remaining nodes of $\mathcal{X}$ 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 $|\mathcal{C}_\mathcal{X}| = \frac{n}{2} - 1$. 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. □

2. Proof of Theorem 1

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) Recall that $\mathcal{N}$ is the set of normal nodes, and define $N = |\mathcal{N}|$. Furthermore, recall that $M[t]$ and $m[t]$ are the maximum and minimum values of the normal nodes at time-step t, respectively. From Lemma 1, we know 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 $\mathcal{X}_M(t, \epsilon_i)$ denote the set of all *normal* nodes that have values in the range $(A_M - \epsilon_i, A_M + \epsilon_i)$, and let $\mathcal{X}_m(t, \epsilon_i)$ denote the set of all *normal* nodes that have values in the range $(A_m - \epsilon_i, A_m + \epsilon_i)$. Note that $\mathcal{X}_M(t, \epsilon_0)$ and $\mathcal{X}_m(t, \epsilon_0)$ are disjoint, by the definition of ϵ_0.

For some ϵ (which we will show how to choose later) satisfying $\epsilon_0 > \epsilon > 0$, 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 disjoint sets $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ and $\mathcal{X}_m(t_\epsilon, \epsilon_0)$. At least one of these two sets must be $(2F+1)$-reachable due to the assumption of $(2F+1)$-robustness of the network. If $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ is $(2F+1)$-reachable, there exists some normal node $i \in \mathcal{X}_M(t_\epsilon, \epsilon_0)$ that has at least $F+1$ normal neighbors outside $\mathcal{X}_M(t_\epsilon, \epsilon_0)$. By definition, all of 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 $\mathcal{X}_m(t_\epsilon, \epsilon_0)$ is $(2F+1)$-reachable, there exists some normal node $j \in \mathcal{X}_m(t_\epsilon, \epsilon_0)$ that will satisfy $x_j[t_\epsilon + 1] \geq A_m + \alpha\epsilon_0 - (1-\alpha)\epsilon$. Again, any normal node that is not in $\mathcal{X}_m(t_\epsilon, \epsilon_0)$ will have the same lower bound.

Define $\epsilon_1 = \alpha\epsilon_0 - (1-\alpha)\epsilon$, and 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 sets $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ and $\mathcal{X}_m(t_\epsilon, \epsilon_0)$ was $(2F+1)$-reachable, it must be that either $|\mathcal{X}_M(t_\epsilon + 1, \epsilon_1)| < |\mathcal{X}_M(t_\epsilon, \epsilon_0)|$ or $|\mathcal{X}_m(t_\epsilon + 1, \epsilon_1)| < |\mathcal{X}_m(t_\epsilon, \epsilon_0)|$, or both. Further note that $\epsilon_1 < \epsilon_0$, and thus $\mathcal{X}_M(t_\epsilon + 1, \epsilon_1)$ and $\mathcal{X}_m(t_\epsilon + 1, \epsilon_1)$ are still disjoint. We can repeat this analysis for time-steps $t_\epsilon + j$, $j \geq 2$, to define sets $\mathcal{X}_M(t_\epsilon + j, \epsilon_j)$ and $\mathcal{X}_m(t_\epsilon + j, \epsilon_j)$, where ϵ_j is defined recursively as $\epsilon_j = \alpha\epsilon_{j-1} - (1-\alpha)\epsilon$. Furthermore, at time-step $t_\epsilon + j$, either $|\mathcal{X}_M(t_\epsilon + j, \epsilon_j)| < |\mathcal{X}_M(t_\epsilon + j - 1, \epsilon_{j-1})|$ or $|\mathcal{X}_m(t_\epsilon + j, \epsilon_j)| < |\mathcal{X}_m(t_\epsilon + j - 1, \epsilon_{j-1})|$, or both. Since $|\mathcal{X}_M(t_\epsilon, \epsilon_0)| + |\mathcal{X}_m(t_\epsilon, \epsilon_0)| \leq N$, there must be some time-step $t_\epsilon + T$ (where $T \leq N$) where either $\mathcal{X}_M(t_\epsilon + T, \epsilon_T)$ or $\mathcal{X}_m(t_\epsilon + T, \epsilon_T)$ is empty. In the former case, all nodes in the network at time-step $t_\epsilon + T$ have value less than $A_M - \epsilon_T$, and in the latter case all nodes in the network at time-step $t_\epsilon + T$ have value greater 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

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

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

3. 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 $\tau_1 \in \{t_\epsilon, t_\epsilon + 1, \ldots, t_\epsilon + c - 1\}$ such that $\mathcal{D}[\tau_1]$ is $(2F+1)$-robust. As in the proof of Theorem 1, there either exists $i \in \mathcal{X}_M(\tau_1, \epsilon_0)$ such that $x_i[\tau_1 + 1] \leq A_M - \epsilon_1$ or $j \in \mathcal{X}_m(\tau_1, \epsilon_0)$ 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,

$|\mathcal{X}_M(\tau_1+1,\epsilon_1)| < |\mathcal{X}_M(\tau_1,\epsilon_0)|$ or $|\mathcal{X}_m(\tau_1+1,\epsilon_1)| < |\mathcal{X}_m(\tau_1,\epsilon_0)|$, 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 τ_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 τ_1+k, for all $1 \leq k \leq Nc$. Because of these relationships, we have that $|\mathcal{X}_M(\tau_1+k,\epsilon_k)| \leq |\mathcal{X}_M(\tau_1+k-1,\epsilon_{k-1})|$ and $|\mathcal{X}_m(\tau_1+k,\epsilon_k)| \leq |\mathcal{X}_m(\tau_1+k-1,\epsilon_{k-1})|$, for each time-step regardless of the network topology. However, we are interested in the time-steps $\tau_1, \tau_2, \ldots$, in which $|\mathcal{X}_M(\tau_j+1,\epsilon_{(1+\tau_j-\tau_1)})| < |\mathcal{X}_M(\tau_j,\epsilon_{(\tau_j-\tau_1)})|$ or $|\mathcal{X}_m(\tau_j+1,\epsilon_{(1+\tau_j-\tau_1)})| < |\mathcal{X}_m(\tau_j,\epsilon_{(\tau_j-\tau_1)})|$. These time-steps correspond to the times at which $\mathcal{D}[\tau_j]$ is $(2F+1)$-robust. Since $|\mathcal{X}_M(\tau_1,\epsilon_0)| + |\mathcal{X}_m(\tau_1,\epsilon_0)| \leq N$ and $|\tau_N - \tau_1| \leq Nc$, there must be some time-step $\tau = \tau_1 + T$ (where $T \leq Nc$) where either $\mathcal{X}_M(\tau_1+T,\epsilon_T)$ or $\mathcal{X}_m(\tau_1+T,\epsilon_T)$ is empty. In the former case, all normal nodes in the network at time-step τ_1+T have value at most $A_M - \epsilon_T$, and in the latter case all normal nodes in the network at time-step τ_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. □

References

[1] S. Morris. Contagion. *The Review of Economic Studies*, 67(1):57–78, 2000.

[2] M. E. J. Newman. Spread of epidemic disease on networks. *Phys. Rev. E*, 66:016128, July 2002.

[3] D. Easley and J. Kleinberg. *Networks, Crowds and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, 2010.

[4] D. Stauffer and M. Sahimi. Can a few fanatics influence the opinion of a large segment of a society? *The European Physical Journal B - Condensed Matter and Complex Systems*, 57:147–152, 2007.

[5] J. Xie, S. Sreenivasan, G. Korniss, W. Zhang, C. Lim, and B. K. Szymanski. Social consensus through the influence of committed minorities. *Phys. Rev. E*, 84:011130, July 2011.

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

[7] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: a survey. *Computer Networks*, 38(4):393–422, 2002.

[8] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proc. 44th Annual IEEE Symp. Foundations of Comp. Sci.*, pages 482–491, oct. 2003.

[9] J. Tsitsiklis, D. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.

[10] A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, June 2003.

[11] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, California, 1996.

[12] S. Sundaram and C. N. Hadjicostis. Distributed function calculation via linear iterative strategies in the presence of malicious agents. *IEEE Transactions on Automatic Control*, 56(7):1495–1508, July 2011.

[13] F. Pasqualetti, A. Bicchi, and F. Bullo. Consensus computation in unreliable networks: A system theoretic approach. *IEEE Transactions on Automatic Control*, 57(1):90–104, Jan. 2012.

[14] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(2):382–401, 1982.

[15] 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*, (HSCC '11), pages 281–290, Chicago, IL, 2011.

[16] N. Agmon and D. Peleg. Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM J. Comput.*, 36(1):56–82, 2006.

[17] X. Défago, M. Gradinariu, S. Messika, and P. Raipin-Parvédy. Fault-tolerant and self-stabilizing mobile robots gathering. In Shlomi Dolev, editor, *Distributed Computing*, volume 4167 of *Lecture Notes in Computer Science*, pages 46–60. Springer Berlin, Heidelberg, 2006.

[18] Z. Bouzid, M. Gradinariu, and S. Tixeuil. Optimal byzantine-resilient convergence in uni-dimensional robot networks. *Theoretical Computer Science*, 411(34-36):3154–3168, 2010.

[19] D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *Journal of the Association for Computing Machinery*, 40(1):17–47, Jan. 1993.

[20] R. Albert, H. Jeong, and A. L. Barabasi. Error and attack tolerance of complex networks. *Letters to Nature*, 406:378–382, June 2000.

[21] B. Bollobas and O. Riordan. Robustness and vulnerability of scale-free random graphs. *Internet Mathematics*, 1(1):1–35, 2004.

[22] N. H. Vaidya, L. Tseng, and G. Liang. Iterative approximate Byzantine consensus in arbitrary directed graphs. In *Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing*, pages 365–374, 2012.

[23] A. Pelc and D. Peleg. Broadcasting with locally bounded Byzantine faults. In *Information Processing Letters*, pages 109–115, 2005.

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

[25] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

[26] W. Ren, R. W. Beard, and E. M. Atkins. Information consensus in multivehicle cooperative control. *IEEE Control Systems Magazine*, 27(2):71–82, April 2007.

[27] L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Systems and Control Letters*, 53:65–78, 2004.

[28] W. Ren and R. W. Beard. Consensus seeking in multiagent systems under dynamically changing interation topologies. *IEEE Transactions on Automatic Control*, 50(5):655–661, May 2005.

[29] L. Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–181, Feb. 2005.

[30] J. Tsitsiklis. *Problems in Decentralized Decision Making and Computation*. PhD thesis, Department of EECS, MIT, 1984.

[31] B. Touri and A. Nedić. On ergodicity, infinite flow, and consensus in random models. *IEEE Transactions on Automatic Control*, 56(7):1593–1605, July 2011.

[32] J. Lorenz and D. A. Lorenz. On conditions for convergence to consensus. *IEEE Transactions on Automatic Control*, 55(7):1651–1656, July 2010.

[33] V. Gupta, C. Langbort, and R. M. Murray. On the robustness of distributed algorithms. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 3437–3478, 2006.

[34] A. Teixeira, H. Sandberg, and K. H. Johansson. Networked control systems under cyber attacks with applications to power networks. In *Proc. of the American Control Conference*, pages 3690–3696, 2010.

[35] A. Chapman and M. Mesbahi. Semi-autonomous networks: Network resilience and adaptive trees. In *Proceedings of the 49th IEEE Conference on Decision and Control*, pages 7473–7478, 2010.

[36] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3):1–24, 2002.

[37] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis. On Krause's multi-agent consensus model with state-dependent connectivity. *IEEE Transactions on Automatic Control*, 54(11):2586–2597, November 2009.

[38] 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*, 33(3):499 – 516, 1986.

[39] R. M. Kieckhafer and M. H. Azadmanesh. Reaching approximate agreement with mixed mode faults. *IEEE Transactions on Parallel and Distributed Systems*, 5(1):53–63, 1994.

[40] R. M. Kieckhafer and M. H. Azadmanesh. Low cost approximate agreement in partially connected networks. *Journal of Computing and Information*, 3(1):53–85, 1993.

[41] A. H. Azadmanesh and H. Bajwa. Global convergence in partially fully connected networks (pfcn) with limited relays. *The 27th Annual Conference of the IEEE*, 3:2022–2025, 2001.

[42] M. H. Azadmanesh and R. M. Kieckhafer. Asynchronous approximate agreement in partially connected networks. *International Journal of Parallel and Distributed Systems and Networks*, 5(1):26–34, 2002.

[43] H. J. LeBlanc and X. D. Koutsoukos. Low complexity resilient consensus in networked multi-agent systems with adversaries. In *Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control*, HSCC '12, pages 5–14, 2012.

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

[45] C. Y. Koo. Broadcast in radio networks tolerating byzantine adversarial behavior. In *Proc. of 23rd Annual ACM Symp. on Principles of Distributed Computing*, pages 275–282, 2004.

[46] P. Erdos and A. Renyi. On random graphs. *Publ. Math. Debrecen*, 6:290–297, 1959.

[47] P. Erdos and A. Renyi. On the strength of connectedness of a random graph. *Acta Math. Acad. Sci. Hungar.*, 12:261–267, 1961.

[48] B. Bollobas. *Random Graphs*. 2nd Edition. Cambridge University Press, 2001.

[49] H. Zhang and S. Sundaram. Robustness of complex networks with implications for consensus and contagion. In *Proceedings of the 51st IEEE Conference on Decision and Control*, pages 3426–3432, 2012.

[50] M. Penrose. *Random Geometric Graphs*. Oxford University Press, 2003.

[51] P. Santi and D. M. Blough. The critical transmitting range for connectivity in sparse wireless ad hoc networks. *IEEE Transactions on Mobile Computing*, 2(1):25–39, 2003.

[52] R. Albert and A. L. Barabasi. Statistical mehanics of complex networks. *Review. Mod. Phys.*, 74(1):47–97, Jan 2002.