A Simple Median-Based Resilient Consensus Algorithm

Haotian Zhang and Shreyas Sundaram

***Abstract*— This paper studies the problem of reaching consensus resiliently in the presence of misbehaving nodes. We design a consensus algorithm where, at each time-step, each node updates its value as a weighted average of its own value and the median of its neighbors' values. This algorithm requires no global information about the network, and is computationally lightweight. We develop a novel graph property that we term *excess robustness*, and use this property to characterize the ability of the median algorithm to succeed under various fault models. We also provide a construction for excess robust graphs. We show that the sensitivity of this algorithm varies greatly under different fault models, and make connections to related ideas from the literature on contagion and graphical games.**

I. Introduction

There has been a great deal of research devoted to the study of large-scale networks (both natural and engineered). Examples of such networks abound in both the natural world (e.g., ecological systems, biological systems, and social systems), and in engineered applications (e.g., the Internet, power grid, and large-scale sensor networks). Distributed consensus is a common objective in these networks, and has applications in various areas such as data aggregation [1], distributed optimization [2], and flocking [3]. A fundamental challenge for reaching consensus in large-scale 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 [4], [5] and the references therein). It has been shown that given F misbehaving nodes, there exist strategies for the misbehaving nodes to prevent certain normal nodes from correctly obtaining any information about other normal nodes' values 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 ensure consensus (under the wireless broadcasting model of communication) [4], [6], [7]. However, these strategies either require the normal nodes to hold some global information (e.g., the topology of the network) or assume that the network is *complete*, i.e., all-to-all communication or sensing [8], [9], [10]. 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 computationally lightweight and operate using only local information (i.e., without knowledge of the network topology and the identities of non-neighboring nodes).

Various researchers have investigated consensus dynamics where the nodes attempt to minimize the influence of malicious agents without knowledge of the global network topology [11], [12], [13]. In [14], [13], it was shown that the traditional metric of network connectivity is no longer sufficient to characterize convergence to agreement when the nodes use a certain class of local filtering rules (where each node removes the largest F and smallest F values from its neighborhood at each time-step, for some nonnegative integer F). Instead, in [14], we introduced a new topological property termed as *network robustness*, and showed that consensus can be reached (resiliently, and without requiring global information) in graphs that are sufficiently robust. Although connectivity and robustness are quite different in general, in that one can construct graphs that have extremely high connectivity and very poor robustness (as argued in [14], [13], [15]), we showed that the notions of connectivity and robustness actually coincide on various random graph models and provided more insights about (purely local) diffusion dynamics over complex networks [16]. These ideas were extended to necessary and sufficient conditions for tolerating various kinds of adversarial behavior in [15], [17].

In this paper, we develop another novel resilient consensus algorithm – the *Median Consensus Algorithm (MCA)* – where each node uses only its own value and the median of its neighbors' values in its update. This algorithm requires no global information (e.g., the topology of the network and the number of misbehaving nodes) and is computationally simple. We show that previously developed graph-theoretic properties (such as connectivity) are not directly applicable to characterize the performance of this algorithm, i.e., even highly connected graphs cannot guarantee consensus. Thus, we introduce an extension of robustness termed as *excess robustness* and show that this concept is the key property to capture the dynamics of MCA. We give a construction method for excess robust graphs, and make connections to graph properties that arise in the study of contagion dynamics on networks [18]. Finally, we analyze the sensitivity of MCA to different fault models in certain networks and show that MCA is sensitive to Byzantine behavior but is relatively robust to malicious behavior.[2]

This material is based upon work supported in part by the Natural Sciences and Engineering Research Council of Canada.

The authors are with the Department of Electrical and Computer Engineering at the University of Waterloo. E-mail for corresponding author: ssundara@uwaterloo.ca.

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

[2] The definitions of Byzantine behavior and malicious behavior will be clear in the following section.

II. Problem Formulation

A. Network Model

Consider a time-varying network modeled by the directed graph $\mathcal{G}[t] = \{\mathcal{V}, \mathcal{E}[t]\}$, 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 *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*, or just *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 $\mathcal{V}_i^{\text{out}}[t] = \{j \in \mathcal{V} \colon (i, j) \in \mathcal{E}[t]\}$ and the *out-degree* of i is denoted $d_i^{\text{out}}[t] =|\mathcal{V}_i^{\text{out}}[t]|$. To account for the fact that 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.

B. Update Model

Each normal 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.) and interacts synchronously with its 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 function, and may be different for each node, depending on its role in the network. These functions are designed *a priori* so that the normal nodes compute some desired function. However, some of the nodes may not follow the prescribed strategy if they are compromised by an adversary. Thus, 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.

C. Fault Model

Definition 1 (Byzantine and malicious nodes): A node $i \in \mathcal{A}$ is said to be *Byzantine* if it does not send the same value to all of its out-neighbors at some time-step, or if it applies some other function $f_i'(\cdot)$ at some time-step; it is said to be *malicious* if it sends $x_i[t]$ to all of its out-neighbors at each time-step, but applies some other function $f_i'(\cdot)$ at some time-step. □

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

It is clear that we cannot deal with networks that only contain misbehaving nodes and thus it is necessary to restrict the *number* of such nodes. We consider upper bounds on the number of compromised nodes either in the whole network (F-total) or in each node's neighborhood (F-local).

Definition 2 (F-total and F-local sets): For some $F \in \mathbb{Z}_{\geq 0}$, a set $\mathcal{S} \subset \mathcal{V}$ is F*-total* if it contains at most F nodes in the network, i.e., $|\mathcal{S}| \leq F$; it 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., $|\mathcal{V}_i[t] \bigcap \mathcal{S}| \leq F$, $\forall i \in \mathcal{V} \setminus \mathcal{S}$, $\forall t \in \mathbb{Z}_{\geq 0}$. □

It should be noted that in time-varying network topologies, the local properties defining an F-local set must hold at all time instances. These definitions facilitate the following fault models.

Definition 3 (F-total and F-local models): A set of misbehaving nodes is F*-totally bounded* or F*-locally bounded* if it is an F-total set or F-local set, respectively. We refer to these fault models as the F*-total* and F*-local models*, respectively. □

We have now defined four different fault models, i.e., F-total or F-local, each with Byzantine or malicious behavior. Note that the F-local Byzantine model contains all of the other models as special cases; the benefit of considering the more specialized cases is that we will be able to obtain tighter results for those models.

D. Resilient Asymptotic Consensus

Given the fault models, we now formally define the objective of resilient asymptotic consensus [15]. Let $M[t]$ and $m[t]$ be the *maximum* and *minimum* values of the *normal* nodes at time-step t, respectively.

Definition 4 (Resilient Asymptotic Consensus): Under any of the fault models, the normal nodes are said to achieve *resilient asymptotic consensus* if *both* of the following conditions are satisfied for *any* choice of initial values.

- *Agreement Condition*: there exists $C \in \mathbb{R}$ such that $\lim_{t\to\infty} x_i[t] = C, \forall i \in \mathcal{N}$.
- *Safety Condition*: $x_i[t] \in [m[0], M[0]], \forall i \in \mathcal{N}, \forall t \in \mathbb{Z}_{\geq 0}$.

□

If both of these conditions are satisfied, the consensus value will be in the interval defined by the initial values of the normal nodes. Note that this definition allows misbehaving nodes to influence the consensus value, but to a limited extent. This 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.

III. Median Consensus Algorithm

While there are various approaches to facilitate consensus, a class of *linear iterative strategies* have attracted significant interest in recent years (e.g., see [19], [20] and the references therein). In such strategies, at each time-step t, each node senses or receives information from its neighbors, and updates its value to be a weighted average of its own value and its neighbors' values. One problem with this strategy is that it is not resilient to misbehaving nodes. In fact, it was shown in [3], [21], [22] that a single 'leader' node can cause all agents to reach consensus on an arbitrary value of its choosing (potentially resulting in a harmful situation) simply by holding its value constant. To remedy this, it was shown in [6], [7] that these strategies can be made robust against attacks, at the cost of requiring the normal nodes to know the entire network topology and to perform extensive computations.

In order to achieve resilient consensus based on purely local information, in [14], we studied a simple modification to the update rule introduced in [11], where each node removes the extreme values with respect to its own value. More specifically, we investigated the following *Weighted-Mean-Subsequence-Reduced (W-MSR) algorithm* [14], [15]:[3]

1) At each time-step t, each normal node i obtains the values of its neighbors, and forms a sorted list.
2) If there are less than F values strictly larger than its own value, $x_i[t]$, then normal node i removes all values that are strictly larger than its own. Otherwise, it removes precisely the largest F values in the sorted list (breaking ties arbitrarily). Likewise, if there are less than F values strictly smaller than its own value, then node i removes all values that are strictly smaller than its own. Otherwise, it removes precisely the smallest F values.
3) Let $\mathcal{R}_i[t]$ denote the set of nodes whose values were removed by normal node i in step 2 at time-step t. Each normal node i applies the update

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

where the weights form a convex combination at each time step.

In the above algorithm, each normal node executes a *local filtering strategy* to eliminate potential misbehavior. This algorithm is computationally lightweight and requires no information of the network topology. In the next section, we will briefly review results from [14], [15], [17] describing network conditions under which the W-MSR algorithm can facilitate resilient consensus. In this paper, we will explore another form of local filtering strategy for the normal nodes to resist misbehavior. Specifically, we build on the insight in [23], where the author studied data aggregation in sensor networks and showed that the *median operation* is more robust than other operations (e.g., the average operation) to resist attacks; however, the author only considered a centralized setting and did not consider in-network aggregation. Here, we design a novel resilient consensus algorithm based on the median operation, called the *Median Consensus Algorithm (MCA)*:

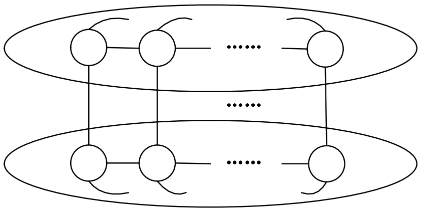

Fig. 1. A $\frac{n}{2}$-connected graph which fails to reach consensus under W-MSR with $F \geq 1$.

1) At each time-step t, each normal node i obtains the values of its neighbors, and forms a *sorted* list $\mathcal{L}_i$.[4] Let $L = |\mathcal{L}_i|$ and $\mathcal{L}_i(j)$ denote the j-th element of $\mathcal{L}_i$.
2) Node i calculates the median $\tilde{x}_i[t]$ of $\mathcal{L}_i$, i.e.,

$$\tilde{x}_i[t] = \frac{\mathcal{L}_i(\lceil \frac{L+1}{2} \rceil) + \mathcal{L}_i(\lfloor \frac{L+1}{2} \rfloor)}{2}. \qquad (2)$$

3) Each normal node i applies the update

$$x_i[t+1] = w_i[t] x_i[t] + \tilde{w}_i[t] \tilde{x}_i[t], \qquad (3)$$

where $w_i[t] + \tilde{w}_i[t] = 1$ and there exists a constant $\alpha > 0$ such that $w_i[t], \tilde{w}_i[t] > \alpha$, $\forall t$.

In MCA, the normal nodes use the median operation as a form of local filtering to attain resilience. Note that the set of nodes filtered away by node i can change over time, depending on their relative values. Thus, even if the network topology itself is fixed, the algorithm effectively induces a time-varying network. Compared with the W-MSR algorithm, to execute MCA, the normal nodes do not need to obtain or estimate the parameter F. Not surprisingly, there is a tradeoff between how much each node knows about the overall network and the conditions required for those nodes to overcome adversaries. In the next section, we will see that even highly connected graphs cannot guarantee consensus by using MCA, and we develop a new concept to capture such dynamics.

IV. Excess Robustness

In [14], [13], it was shown that traditional metrics (such as network connectivity) are not suitable to characterize the behavior of local filtering dynamics (such as W-MSR). To see why, consider the graph shown in Fig. 1. Each of the two subsets shown in that graph induce a complete graph on $\frac{n}{2}$ nodes (suppose n is even), and each node has exactly one neighbor in the opposite subset. This graph has connectivity (and minimum degree) $\frac{n}{2}$. Now suppose that all nodes in the top set start with value 1, and all nodes in the bottom set start with value 0. For each node in the top set, all but one of its neighbors have value 1, and thus the value of the node from the bottom set gets filtered out by using W-MSR. The same holds true for each node in the bottom set. Thus no node in either set ever uses the value of a node from the

[3] Similar dynamics have also been studied in [12], [13], [17].

[4] Note that $\mathcal{L}_i$ depends on the time-step t, and does not include node i's own value.

opposite set, and consensus is never reached in this graph, even when there are no misbehaving nodes.

The problem in the above network is that no node has enough neighbors *outside* its own set. In order to capture this idea, we developed the following concept of robustness in [14].

Definition 5 (r-reachable set): Given a graph $\mathcal{G}$ and a nonempty subset $\mathcal{S}$ of nodes of $\mathcal{G}$, 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 6 (r-robustness): A graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ 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. □

Robustness captures the idea that each pair of subsets of nodes should contain some node which has sufficient neighbors from outside and brings in useful information to at least one of the sets. It was shown in [14] that if the network is sufficiently robust (i.e., $(2F+1)$-robust), the W-MSR algorithm facilitates resilient consensus under the 'worst-case' fault model (i.e., the F-local Byzantine model).

With the same set of initial values as described above (i.e., the top and bottom sets have initial values 1 and 0, respectively), it is not difficult to see that MCA also fails on the network shown in Fig. 1. However, the reason why MCA fails is different from that of the W-MSR algorithm. Under the MCA algorithm, each node is filtering away almost all of its neighbors values, and adopting only the value in the middle. Thus, for a node in some subset $\mathcal{S}$ to adopt some value from outside $\mathcal{S}$, it must have at least as many neighbors outside $\mathcal{S}$ as inside. In other words, for a set of nodes, whether a node will adopt outside information depends on the *relative* number of its neighbors that are outside the set versus inside.[5] To capture this idea, we extend the concepts of reachable sets and robustness as follows.

Definition 7 (r-excess reachable set): Given a graph $\mathcal{G}$ and a nonempty subset $\mathcal{S}$ of nodes of $\mathcal{G}$, we say $\mathcal{S}$ is an *r-excess reachable set* if $\exists i \in \mathcal{S}$ such that $|\mathcal{V}_i \setminus \mathcal{S}| - |\mathcal{V}_i \cap \mathcal{S}| \geq r$, where $r \in \mathbb{Z}_{\geq 0}$. When the context is clear, we will also say that node i (with regard to set $\mathcal{S}$) is r-excess reachable. □

Definition 8 (r-excess robustness): A graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ is *r-excess 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-excess reachable. □

The following result provides an important property for r-excess robust graphs.

Lemma 1: Given an r-excess robust graph $\mathcal{G}$, where $r \geq 2$, the minimum in-degree of $\mathcal{G}$ is r.

Proof: We will prove by contradiction. Assume that there exists an r-excess robust graph $\mathcal{G}$ with some node i such that $d_i < r$. We consider the pair of subsets $\{i\}$ and $\mathcal{V} \setminus \{i\}$. When $r \geq 2$, it is clear that neither of them is r-excess reachable (since $d_i < r$) which contradicts the assumption that $\mathcal{G}$ is r-excess robust. ■

Note that the above result does not apply to the case when $r = 1$, i.e., there exist graphs which are 1-excess robust but with minimum in-degree 0 (e.g., the directed chain). Excess robustness represents a type of information redundancy in the network. Inspired by [15], we also define another type of redundancy by identifying the *number* of nodes that have more neighbors outside their set than inside.

Definition 9 ((r, s)-excess reachable set): Given a graph $\mathcal{G}$ and a nonempty subset of nodes $\mathcal{S}$, we say that $\mathcal{S}$ is an *(r, s)-excess reachable set* if given $\mathcal{X}_{\mathcal{S}}^r = \{i \in \mathcal{S} \colon |\mathcal{V}_i \setminus \mathcal{S}| - |\mathcal{V}_i \cap \mathcal{S}| \geq r\}$, then $|\mathcal{X}_{\mathcal{S}}^r| \geq s$, where $r, s \in \mathbb{Z}_{\geq 0}$. □

Definition 10 ((r, s)-excess robustness): Consider $r \in \mathbb{Z}_{\geq 0}$ and $s \in \{1, \ldots, n\}$. A graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ is *(r, s)-excess robust* if for every pair of nonempty, disjoint subsets $\mathcal{S}_1$ and $\mathcal{S}_2$ of $\mathcal{V}$, at least one of the following holds (recall that $\mathcal{X}_{\mathcal{S}}^r = \{i \in \mathcal{S} \colon |\mathcal{V}_i \setminus \mathcal{S}| - |\mathcal{V}_i \cap \mathcal{S}| \geq r\}$):

1) $|\mathcal{X}_{\mathcal{S}_1}^r| + |\mathcal{X}_{\mathcal{S}_2}^r| \geq s$;
2) $|\mathcal{X}_{\mathcal{S}_1}^r| = |\mathcal{S}_1|$;
3) $|\mathcal{X}_{\mathcal{S}_2}^r| = |\mathcal{S}_2|$.

□

The concept of (r, s)-excess robustness captures a form of redundancy such that for a pair of sets, there are at least s nodes (in the union of the two sets) each of which has r more neighbors from outside its own set (i.e., condition (i)). Conditions (ii) and (iii) capture the cases when all of the nodes in one of the sets have at least r more neighbors outside that set than inside.

In the next section, we will show that the notion of excess robustness is key to characterizing the performance of MCA, i.e., MCA succeeds if the network is sufficiently excess robust.

V. Consensus Results

Recall that in Definition 4, there are two conditions for resilient consensus: the safety condition and the agreement condition. In this section, we first focus on the safety condition and show that the degrees of the nodes play a key role in this condition.

Lemma 2: Suppose each normal node updates its value according to MCA. Under any of the four fault models (F-total or F-local, Byzantine or malicious), the safety condition is guaranteed *if and only if* $d_i[t] = 0$ or $d_i[t] \geq 2F + 1, \forall i \in \mathcal{V}, \forall t \in \mathbb{Z}_{\geq 0}$.

Proof: *(Necessity)* If there is a node i such that $1 \leq d_i[t] \leq 2F$ at some time-step t, then we can always choose a set $\mathcal{A}$ of misbehaving (Byzantine or malicious) nodes such that $|\mathcal{A} \cap \mathcal{V}_i[t]| \geq |\mathcal{N} \cap \mathcal{V}_i[t]|$. If the misbehaving nodes choose their values to be bigger than $\frac{2}{\alpha}(M[0] - m[0]) + m[0]$ (or smaller than $\frac{2}{\alpha}(m[0] - M[0]) + M[0]$), then the set $\mathcal{L}_i$ will have at least half of its values equal to this and the value $\tilde{x}_i[t]$ in (2) will be bigger than $\frac{1}{\alpha}(M[0] - m[0]) + m[0]$ (or smaller than $\frac{1}{\alpha}(m[0] - M[0]) + M[0]$). Thus, we get $x_i[t+1] > M[0]$ (or $x_i[t+1] < m[0]$) and the invariance of the normal nodes' values is disrupted.

(Sufficiency) First note that if a normal node has zero in-degree at time-step t, then its value will not change at the next time-step and thus we just need to focus on the normal nodes which have nonzero in-degree. Recall that there are at most F misbehaving nodes in each normal

[5] This is in contrast to the notion of reachable sets in Definition 5, where only the absolute number of neighbors outside the set mattered.

node's neighborhood. Thus each normal node which has nonzero in-degree will receive at most F values outside the interval $[m[t], M[t]]$ at each time-step t. If $\min\{d_i[t] : i \in \mathcal{V}, d_i[t] > 0\} \geq 2F + 1, \forall t$, then such normal node will receive at least $F + 1$ values inside the interval $[m[t], M[t]]$ and by using MCA, no normal nodes will use a value outside $[m[t], M[t]]$ at each time-step. Since the update rule in (3) is a convex combination of each normal node's own value and the median value in its neighborhood, both $M[t]$ and $m[t]$ are monotone and bounded functions of t. Thus, the safety condition is guaranteed. ■

Note that the W-MSR algorithm is 'naturally' safe since each normal node will not use any value if it does not have sufficient neighbors (i.e., if a node has $2F$ or fewer neighbors and the parameter of the algorithm is F, then all of its neighbors' values will be removed). However, by using MCA, each normal node which has nonzero in-degree will always adopt some value from its neighbors; this may be helpful for convergence but not for safety. Thus, the above minimum degree condition is required.

In the rest of this section, we will show that (r, s)-excess robustness together with the minimum degree of the network captures the *necessary and sufficient* conditions for MCA to succeed under the F-total malicious model (in time-invariant networks). Then we will give a sufficient condition for MCA to succeed under the F-local Byzantine model, which also applies to the other fault models.

A. F-Total Malicious Model

Using the concept of (r, s)-excess robustness, we obtain the main result in this paper.

Theorem 1: Consider a time-invariant network modeled by a directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ where each normal node updates its value according to MCA. Under the F-total malicious model with $F \geq 1$, resilient asymptotic consensus is achieved *if and only if* the minimum in-degree of the network is at least $2F + 1$ and the network is $(0, F + 1)$-excess robust. □

Proof: (Necessity) First note that if there is more than one node with zero in-degree, then consensus will not be reached if these nodes are normal and have different initial values. If there is only one node which has zero in-degree, we can choose this node to be normal and choose its initial value to be the maximum (or minimum) value in the network. Then when using MCA, no other normal nodes will use its value (even when there are no malicious nodes); in other words, there is some choice of initial values such that consensus cannot be achieved. Thus, by Lemma 2, we know that it is necessary for the minimum in-degree of the network to be $2F + 1$.

If $\mathcal{G}$ is not $(0, F + 1)$-excess robust, then there are nonempty, disjoint subsets $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{V}$ such that none of the conditions $(i) - (iii)$ in Definition 10 hold. Suppose the initial value of each node in $\mathcal{S}_1$ is a and each node in $\mathcal{S}_2$ is b, with $a < b$. Let all other nodes have initial values taken from the interval (a, b). Since $|\mathcal{X}^0_{\mathcal{S}_1}| + |\mathcal{X}^0_{\mathcal{S}_2}| \leq F$, suppose all nodes in $\mathcal{X}^0_{\mathcal{S}_1}$ and $\mathcal{X}^0_{\mathcal{S}_2}$ 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 $|\mathcal{X}^0_{\mathcal{S}_1}| < |\mathcal{S}_1|$ and $|\mathcal{X}^0_{\mathcal{S}_2}| < |\mathcal{S}_2|$, respectively. Since these normal nodes are not 0-excess reachable, none of them will use values from outside and no consensus among normal nodes is reached.

(Sufficiency) By Lemma 2, the safety condition is guaranteed and we focus on the agreement condition. 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 the proof of Lemma 2, we know that if the minimum in-degree of the network is at least $2F+1$, then 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 $\epsilon_i > 0$, let $\mathcal{X}_M(t, \epsilon_i) = \{j \in \mathcal{V} : x_j[t] > A_M - \epsilon_i\}$, which includes all normal and malicious nodes that have values larger than $A_M - \epsilon_i$, and let $\mathcal{X}_m(t, \epsilon_i) = \{j \in \mathcal{V} : x_j[t] < A_m + \epsilon_i\}$, which includes all normal and malicious nodes that have values smaller than $A_m + \epsilon_i$. Note that $\mathcal{X}_M(t, \epsilon_0)$ and $\mathcal{X}_m(t, \epsilon_0)$ are disjoint, by the definition of ϵ_0.

Fix $\epsilon < \frac{(\frac{\alpha}{2})^N}{1-(\frac{\alpha}{2})^N}\epsilon_0$, which satisfies $\epsilon_0 > \epsilon > 0$. Let t_ϵ be such that $M[t] < A_M + \epsilon$ and $m[t] > A_m - \epsilon$, $\forall t \geq t_\epsilon$ (we know that such a t_ϵ exists by the definition of convergence). Consider the nonempty and disjoint sets $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ and $\mathcal{X}_m(t_\epsilon, \epsilon_0)$. Since the network is $(0, F+1)$-excess robust and there are no more than F malicious nodes in the network (F-total model), there is at least one normal node in the union that is 0-excess reachable.

Without loss of generality, suppose normal node $i \in \mathcal{X}_M(t_\epsilon, \epsilon_0)$ is 0-excess reachable. Note that the minimum in-degree of the network is $2F + 1$ and thus each normal node always adopts a value from its neighbors, Further note that since node i is 0-excess reachable, the number of node i's neighbors from outside $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ is no less than the number of its neighbors from inside, and thus $\tilde{x}_i[t_\epsilon]$ from (2) will incorporate (or be equal to) a value from outside. By definition, the neighbors outside $\mathcal{X}_M(t_\epsilon, \epsilon_0)$ have values at most equal to $A_M - \epsilon_0$. Since the safety condition is guaranteed, the biggest value that node i will use from inside its own set is $M[t_\epsilon]$. Thus, the median value that node i will use is at most $\tilde{x}_i[t_\epsilon] = \frac{M[t_\epsilon]+A_M-\epsilon_0}{2}$.

Note that at each time-step, every normal node's value is a convex combination of its own value and the median value it chooses 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 $M[t_\epsilon]$ in (3) produces

$$\begin{aligned} x_i[t_\epsilon+1] &\le (1-\alpha)M[t_\epsilon] + \alpha\frac{M[t_\epsilon]+A_M-\epsilon_0}{2} \\ &< (1-\alpha)(A_M+\epsilon) + \alpha\frac{A_M+\epsilon+A_M-\epsilon_0}{2} \\ &= A_M - \frac{\alpha}{2}\epsilon_0 + (1-\frac{\alpha}{2})\epsilon. \end{aligned}$$

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 and will not use a value bigger than $M[t_\epsilon]$. Similarly, if $j \in \mathcal{X}_m(t_\epsilon,\epsilon_0)$ is 0-excess reachable, then $x_j[t_\epsilon+1] > A_m + \frac{\alpha}{2}\epsilon_0 - (1-\frac{\alpha}{2})\epsilon$. Again, any normal node that is not in $\mathcal{X}_m(t_\epsilon,\epsilon_0)$ will have the same lower bound.

Define $\epsilon_1 = \frac{\alpha}{2}\epsilon_0 - (1-\frac{\alpha}{2})\epsilon$, which satisfies $0 < \epsilon < \epsilon_1 < \epsilon_0$. Consider the sets $\mathcal{X}_M(t_\epsilon+1,\epsilon_1)$ and $\mathcal{X}_m(t_\epsilon+1,\epsilon_1)$. Since at least one of the normal nodes in $\mathcal{X}_M(t_\epsilon,\epsilon_0)$ decreases at least to $A_M-\epsilon_1$ (or below), or one of the nodes in $\mathcal{X}_m(t_\epsilon,\epsilon_0)$ increases at least to $A_m+\epsilon_1$ (or above), it must be that either $|\mathcal{X}_M(t_\epsilon+1,\epsilon_1)| < |\mathcal{X}_M(t_\epsilon,\epsilon_0)|$ or $|\mathcal{X}_m(t_\epsilon+1,\epsilon_1)| < |\mathcal{X}_m(t_\epsilon,\epsilon_0)|$, or both. Since $\epsilon_1 < \epsilon_0$, $\mathcal{X}_M(t_\epsilon+1,\epsilon_1)$ and $\mathcal{X}_m(t_\epsilon+1,\epsilon_1)$ are still disjoint.

For $j \ge 1$, define ϵ_j recursively as $\epsilon_j = \frac{\alpha}{2}\epsilon_{j-1} - (1-\frac{\alpha}{2})\epsilon$, and observe that $\epsilon_j < \epsilon_{j-1}$. As long as there are still normal nodes in both $\mathcal{X}_M(t_\epsilon+j,\epsilon_j)$ and $\mathcal{X}_m(t_\epsilon+j,\epsilon_j)$, then we can repeat the above analysis for time-steps $t_\epsilon+j$. Furthermore, at time-step $t_\epsilon+j$, either $|\mathcal{X}_M(t_\epsilon+j,\epsilon_j)| < |\mathcal{X}_M(t_\epsilon+j-1,\epsilon_{j-1})|$ or $|\mathcal{X}_m(t_\epsilon+j,\epsilon_j)| < |\mathcal{X}_m(t_\epsilon+j-1,\epsilon_{j-1})|$, or both.

Since $|\mathcal{X}_M(t_\epsilon,\epsilon_0)|+|\mathcal{X}_m(t_\epsilon,\epsilon_0)| \le N$, there must be some time-step $t_\epsilon+T$ (where $T \le N$) 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 normal nodes in the network at time-step $t_\epsilon+T$ have value at most $A_M-\epsilon_T$, and in the latter case all normal nodes in the network at time-step $t_\epsilon+T$ have value no less than $A_m+\epsilon_T$. We will show that $\epsilon_T > 0$, which will contradict the fact that the largest value monotonically converges to A_M (in the former case) or that the smallest value monotonically converges to A_m (in the latter case). To do this, note that

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

Since $\epsilon < \frac{(\frac{\alpha}{2})^N}{1-(\frac{\alpha}{2})^N}\epsilon_0$, we obtain $\epsilon_T > 0$, providing the desired contradiction. It must thus be the case that $\epsilon_0 = 0$, proving that $A_M = A_m$. ■

Remark 1: Note that since the other fault models are stronger than the F-total malicious model, the necessary condition of $(0, F+1)$-excess robustness also applies to these models. Furthermore, by a similar proof, we can show that 0-excess robustness is necessary for MCA to succeed even without misbehaving nodes. □

B. F-Local Byzantine Model

We now provide a sufficient condition for MCA to succeed under the F-local Byzantine model; furthermore, we extend this condition to time-changing networks. Since the other fault models are special cases of the F-local Byzantine model, these results also apply to those models.

Theorem 2: Consider a time-invariant network modeled by a directed graph $\mathcal{G} = \{\mathcal{V},\mathcal{E}\}$ where each normal node updates its value according to MCA. Under the F-local Byzantine model, resilient asymptotic consensus is achieved if the topology of the network is $(2F+1)$-excess robust. □

Proof: The proof is similar to the proof of sufficiency of Theorem 1. Note that by Lemma 1 and 2, the safety condition is guaranteed if the network is $(2F+1)$-excess robust. ■

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

VI. Construction of Excess Robust Graphs

Having shown that excess robustness is the key concept to capture the dynamics of MCA, we now provide a construction method for excess robust graphs. Note that the notion of (r,s)-excess robustness in Definition 10 is a strict generalization of excess robustness (Definition 6). Thus, the following result (adapted from a construction for robust graphs provided in [14], [15]) also applies to excess robust graphs.

Theorem 3: Let $\mathcal{G} = \{\mathcal{V},\mathcal{E}\}$ be an (r,s)-excess robust graph (with $s \ge 1$). Then the directed graph $\mathcal{G}' = \{\mathcal{V}',\mathcal{E}'\} = \{\mathcal{V}\cup\{v_{\text{new}}\},\mathcal{E}\cup\mathcal{E}_{\text{new}}\}$, where v_{new} is a new node added to $\mathcal{G}$ and $\mathcal{E}_{\text{new}}$ is the directed edge set related to v_{new}, is (r,s)-excess robust if $d_{v_{\text{new}}} \ge r+2s-2$ and $d^{\text{out}}_{v_{\text{new}}} = 0$. □

Proof: First note that since $d^{\text{out}}_{v_{\text{new}}} = 0$, if node $i \in \mathcal{V}$ is r-excess reachable in $\mathcal{G}$, then i is also r-excess reachable in $\mathcal{G}'$. For a pair of nonempty, disjoint sets $\mathcal{S}_1$ and $\mathcal{S}_2$ in $\mathcal{V}'$, there are three cases to check: (i) $v_{\text{new}} \notin \mathcal{S}_i$, (ii) $\{v_{\text{new}}\} = \mathcal{S}_i$, and (iii) $v_{\text{new}} \in \mathcal{S}_i$, $i \in \{1,2\}$. In the first case, since $\mathcal{G}$ is (r,s)-excess robust, the conditions in Definition 10 must hold. In the second case, $\mathcal{X}^r_{\mathcal{S}_i} = \mathcal{S}_i$, and we are done. In the third case, suppose without loss of generality that $\mathcal{S}_2 = \mathcal{S}'_2\cup\{v_{\text{new}}\}$. Since $\mathcal{G}$ is (r,s)-excess robust, at least one of the following conditions hold: $|\mathcal{X}^r_{\mathcal{S}_1}|+|\mathcal{X}^r_{\mathcal{S}'_2}| \ge s$, $|\mathcal{X}^r_{\mathcal{S}_1}| = |\mathcal{S}_1|$, or $|\mathcal{X}^r_{\mathcal{S}'_2}| = |\mathcal{S}'_2|$. If either of the first two hold, then the corresponding conditions hold for the pair $\mathcal{S}_1,\mathcal{S}_2$ in $\mathcal{G}'$. So assume only $|\mathcal{X}^r_{\mathcal{S}'_2}| = |\mathcal{S}'_2|$ holds. Then, the negation of the first condition $|\mathcal{X}^r_{\mathcal{S}_1}|+|\mathcal{X}^r_{\mathcal{S}'_2}| \ge s$ implies $|\mathcal{X}^r_{\mathcal{S}'_2}| = |\mathcal{S}'_2| \le s-1$ and $|\mathcal{V}_{v_{\text{new}}}\cap\mathcal{S}_2| \le s-1$. Hence, $|\mathcal{V}_{v_{\text{new}}}\setminus\mathcal{S}_2| \ge r+s-1$, and $|\mathcal{X}^r_{\mathcal{S}_2}| = |\mathcal{S}_2|$, completing the proof. ■

The above result indicates that to construct an (r, s)-excess robust graph with n nodes, we can start with an (r, s)-excess robust graph with relatively smaller order, and continually add new nodes with incoming edges from at least $r + 2s - 2$ nodes in the existing graph (and with no outgoing edges).

Note that unlike robustness, excess robustness and (r, s)-excess robustness are not monotonic properties (i.e., the properties are not invariant under the addition of edges). Thus, it is more difficult to analyze these properties and to find specific topologies where they exist. Here, we give one example of excess robust graphs.

Proposition 1: Consider a complete graph $\mathcal{G}$ with n nodes ($n > 3$). If n is even, $\mathcal{G}$ is $(1, n)$-excess robust; if n is odd, $\mathcal{G}$ is $(2, n)$-excess robust.

Proof: Suppose we choose a subset $\mathcal{S}$ of m nodes. If n is even (odd), all the nodes in $\mathcal{S}$ are at least 1-excess reachable (2-excess reachable) if $m \leq \frac{n}{2}$ ($\leq \lfloor \frac{n}{2} \rfloor$). When choosing a pair of subsets, at least one of the sets will have size at most $\frac{n}{2}$ ($\lfloor \frac{n}{2} \rfloor$). Thus, by Definition 10, when we choose a pair of sets $\mathcal{S}_1$ and $\mathcal{S}_2$, either condition $|\mathcal{X}^1_{\mathcal{S}_1}| = |\mathcal{S}_1|$ ($|\mathcal{X}^2_{\mathcal{S}_1}| = |\mathcal{S}_1|$) or condition $|\mathcal{X}^1_{\mathcal{S}_2}| = |\mathcal{S}_2|$ ($|\mathcal{X}^2_{\mathcal{S}_2}| = |\mathcal{S}_2|$) (or both) will hold and the result follows. ■

Note that by Theorem 1 and Proposition 1, a complete graph with n nodes can tolerate up to $\lfloor \frac{n}{2} \rfloor - 1$ malicious nodes using MCA since the minimum in-degree of the network is $n - 1$. However, in the following proposition, we show that complete graphs cannot tolerate even three Byzantine nodes. This implies that MCA is very sensitive under the Byzantine model but is robust under the malicious model.

Proposition 2: Consider a complete graph with n nodes ($n > 3$). When the normal nodes use MCA, 2 Byzantine nodes can prevent consensus if n is even and 3 Byzantine nodes can prevent consensus if n is odd.

Proof: We will prove this result by providing a strategy for the Byzantine nodes to prevent consensus. Assume $\lfloor \frac{n}{2} \rfloor - 1$ normal nodes have initial value a and the other $\lfloor \frac{n}{2} \rfloor - 1$ normal nodes have initial value b, where $a \neq b$. The Byzantine nodes will send a and b to those normal nodes with initial values a and b, respectively. For a normal node with initial value a (or b), by using MCA, the median value it will adopt is a (or b). Thus, consensus cannot be reached. ■

VII. Relationship to Contagion Models and Graphical Games

It is of interest to relate our results to those obtained in the literature on contagion and graphical games [24], [25], [18]. In those settings, one is given a network where each node can be in one of several states, and nodes adopt a new state if a certain fraction (or number) of their neighbors have adopted that state. One then asks the question: if a specific subset of nodes adopts a certain state, what properties of the network will allow this state to spread to all other nodes?

To be more precise, consider a graph $\mathcal{G}$ under the fractional adoption model (where a node adopts the state if a fraction q of its neighbors have adopted it). Suppose that a set of nodes $\mathcal{S}$ start in state A, and all other nodes start in state B. A set $\mathcal{S}' \subset \mathcal{V}$ is said to be r-cohesive if every node in $\mathcal{S}'$ has at least a fraction r of its neighbors inside $\mathcal{S}'$ [24], [25], [18]. The following result characterizes the spread of state A.

Theorem 4: ([24], [25], [18]) Consider a network $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where each node can be in either state A or state B, and a node switches from state B to state A if a fraction q of its neighbors are in state A. Let $\mathcal{S} \subset \mathcal{V}$ be a set of nodes that start in state A, and let all other nodes start in state B. Then A will spread to all nodes if and only if there is no subset $\mathcal{S}' \subseteq \mathcal{V} \setminus \mathcal{S}$ that is more than $(1 - q)$-cohesive. □

The notion of a cohesive set is similar to that of an excess reachable set developed in this work. Indeed, a 0-excess reachable set is equivalent to a set that is at most $\frac{1}{2}$-cohesive. These similarities arise because under both the contagion and MCA dynamics, nodes are adopting new information only when the number of neighbors with that information exceeds a certain threshold. Our work shows that such graph properties are useful outside the contagion scenario; indeed, our setting also includes misbehaving nodes and allows the state-space of the nodes to be infinite, leading to more complicated dynamics. A study of further similarities between the two scenarios is a rich avenue for future research.

VIII. Summary

In this paper, we proposed the Median Consensus Algorithm (MCA) to achieve resilient consensus. MCA is computationally lightweight and does not require the normal nodes to obtain or estimate the number F of misbehaving nodes (as in the W-MSR algorithm), or to know anything about the network topology (other then its own neighbors). We showed that previously studied metrics (such as connectivity) fail to characterize the performance of MCA. The reason for this failure is that by using MCA, a node in a given subset of nodes will use a value from outside that subset only if at least half of its neighbors are outside the set. This motivated our development of the concept of *excess robustness*. We showed that this concept is very useful to capture the dynamics of MCA; we provided necessary and sufficient conditions for MCA to succeed under the F-total malicious model, and provided separate necessary and sufficient conditions for other fault models. Our results showed that unlike previously developed filtering algorithms, MCA in complete networks is very sensitive to Byzantine nodes, while being highly robust to malicious nodes. We provided a construction for excess robust graphs, and made connections to similar concepts from the literature on contagion and graphical games.

Acknowledgements

The authors thank Nitin Vaidya for the initial discussions that prompted the study of the Median Consensus Algorithm. The authors also thank Heath LeBlanc for his careful reading and for providing feedback.

References

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

[2] J. Tsitsiklis, D. Bertsekas, and M. Athans, “Distributed asynchronous deterministic and stochastic gradient optimization algorithms,” *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803–812, 1986.

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

[4] N. A. Lynch, *Distributed Algorithms*. Elsevier (imprint: Morgan Kaufmann), 1996.

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

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

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

[8] L. Lamport, R. Shostak, and M. Pease, “The Byzantine Generals problem,” *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, July 1982.

[9] N. Agmon and D. Peleg, “Fault-tolerant gathering algorithms for autonomous mobile robots,” *SIAM J. Comput.*, vol. 36, no. 1, pp. 56–82, 2006.

[10] X. Défago, M. Gradinariu, S. Messika, and P. Raipin-Parvédy, “Fault-tolerant and self-stabilizing mobile robots gathering,” *Distributed Computing*, vol. 4167, pp. 46–60, 2006.

[11] D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl, “Reaching approximate agreement in the presence of faults,” *Journal of the ACM*, vol. 33, pp. 499–516, May 1986.

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

[13] H. LeBlanc and X. Koutsoukos, “Low complexity resilient consensus in networked multi-agent systems with adversaries,” in *Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control*, ser. HSCC ’12, 2012, pp. 5–14.

[14] H. Zhang and S. Sundaram, “Robustness of information diffusion algorithms to locally bounded adversaries,” in *Proceedings of the 31st American Control Conference*, 2012, pp. 5855–5861.

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

[16] 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*, 2012, to appear.

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

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

[19] R. Olfati-Saber, J. A. Fax, and R. M. Murray, “Consensus and cooperation in networked multi-agent systems,” *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.

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

[21] J. N. Tsitsiklis, “Problems in decentralized decision making and computation,” Ph.D. dissertation, Massachusetts Institute of Technology, 1984.

[22] 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*, 2006, pp. 3437–3478.

[23] D. Wagner, “Resilient aggregation in sensor networks,” in *Proc. 2nd ACM Workshop on Security of Ad Hoc and Sensor Networks*, 2004, pp. 78–87.

[24] S. Morris, “Contagion,” *The Review of Economic Studies*, vol. 67, no. 1, pp. 57–78, 2000.

[25] M. O. Jackson, *Social and Economic Networks*. Princeton University Press, Princeton, New Jersey, 2008.