

# Broadcast Using Certified Propagation Algorithm in Presence of Byzantine Faults<sup>☆</sup>

Lewis Tseng<sup>1,\*</sup>, Nitin H. Vaidya<sup>2</sup>, Vartika Bhandari<sup>3</sup>

---

## Abstract

We explore the correctness of the Certified Propagation Algorithm (CPA) [6, 1, 8, 5] in solving broadcast with locally bounded Byzantine faults. CPA allows the nodes to use only local information regarding the network topology. We provide a *tight* necessary and sufficient condition on the network topology for the correctness of CPA. To the best of our knowledge, this work is the first to solve the open problem in [8]. We also present some simple extensions of this result.

*Keywords:*

Algorithms, Distributed computing, Fault tolerance, Broadcast, Byzantine faults, Tight condition

---

---

<sup>☆</sup>This research is supported in part by Army Research Office grant W-911-NF-0710287. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the funding agencies or the U.S. government.

\*Corresponding author.

*Email addresses:* [lt seng3@illinois.edu](mailto:lt seng3@illinois.edu) (Lewis Tseng), [nhv@illinois.edu](mailto:nhv@illinois.edu) (Nitin H. Vaidya), [vartika@google.com](mailto:vartika@google.com) (Vartika Bhandari)

<sup>1</sup>Department of Computer Science and Coordinated Science Laboratory, University of Illinois at Urbana-Champaign

<sup>2</sup>Department of Electrical and Computer Engineering and Coordinated Science Laboratory, University of Illinois at Urbana-Champaign

<sup>3</sup>Google, Inc.

## 1. Introduction

In this work, we explore fault-tolerant broadcast with locally bounded Byzantine faults in synchronous point-to-point networks. We assume a *f*-locally bounded model, in which at most *f* Byzantine faults occur in the neighborhood of every *fault-free* node [6]. In particular, we are interested in the necessary and sufficient condition on the underlying communication network topology for the correctness of the Certified Propagation Algorithm (CPA) – the CPA algorithm has been analyzed in prior work [6, 1, 8, 5].

*Problem Formulation.* Consider an arbitrary directed network of *n* nodes. One node in the network, called the *source* (*s*), is given an initial input, which the source node needs to transmit to all the other nodes. The source *s* is assumed to be *fault-free*. We say that CPA is *correct*, if it satisfies the following properties, where  $x_s$  denotes the input at source node *s*:

- **Termination:** every fault-free node *i* eventually decides on an output value  $y_i$ .
- **Validity:** for every fault-free node *i*, its output value  $y_i$  equals the source's input, i.e.,  $y_i = x_s$ . As stated above, the source node is assumed to be fault-free.

In this paper, we study the condition on the network topology for the correctness of CPA.

*Related Work.* Several researchers have addressed CPA problem. [6] studied the problem in an infinite grid. [1] developed a sufficient condition in the context of arbitrary network topologies, but the sufficient condition proposed

is not tight. [8] provided necessary and sufficient conditions, but the two conditions are not identical (not tight). [5] provided another condition that can approximate (within a factor of 2) the largest  $f$  for which CPA is correct in a given graph.

Independently, conditions similar to our condition are also discovered by other researchers [9, 3] under other contexts. [9] proved a similar condition to be sufficient (but not tight) to solve Shamir's  $(n, k)$  threshold secret sharing problem, where the source wants to transmit shares of secret to all the other nodes, and all nodes are assumed to be honest-but-curious. In the context of cascading behavior in the network, [3] showed that a similar condition is necessary and sufficient to achieve a complete cascade, i.e., all nodes have learned the value transmitted by a cluster of sources with same input values using only local information. In their model, all nodes are assumed to be fault-free. Due to our assumption of existence of Byzantine failures, the proofs in this paper are different from the ones in [9, 3].

## 2. System Model

The system is assumed to be *synchronous*. The synchronous communication network consisting of  $n$  nodes including source node  $s$  is modeled as a simple *directed* graph  $G(\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  is the set of  $n$  nodes, and  $\mathcal{E}$  is the set of directed edges between the nodes in  $\mathcal{V}$ . We assume that  $n \geq 2$ , since the problem for  $n = 1$  is trivial. Node  $i$  can transmit messages to another node  $j$  if and only if the directed edge  $(i, j)$  is in  $\mathcal{E}$ . Each node can transmit messages to itself as well; however, for convenience, we exclude self-loops from set  $\mathcal{E}$ . That is,  $(i, i) \notin \mathcal{E}$  for  $i \in \mathcal{V}$ . All the links (i.e., communication channels) are

assumed to be reliable, FIFO (first-in first-out) and deliver each transmitted message exactly once. With a slight abuse of terminology, we will use the terms *edge* and *link* interchangeably.

For each node  $i$ , let  $N_i^-$  be the set of nodes from which  $i$  has incoming edges. That is,  $N_i^- = \{j \mid (j, i) \in \mathcal{E}\}$ . Similarly, define  $N_i^+$  as the set of nodes to which node  $i$  has outgoing edges. That is,  $N_i^+ = \{j \mid (i, j) \in \mathcal{E}\}$ . Nodes in  $N_i^-$  and  $N_i^+$  are, respectively, said to be incoming and outgoing neighbors of node  $i$ . Since we exclude self-loops from  $\mathcal{E}$ ,  $i \notin N_i^-$  and  $i \notin N_i^+$ . However, we note again that each node can indeed transmit messages to itself.

We consider the  $f$ -local fault model, with at most  $f$  incoming neighbors of any fault-free node becoming Byzantine faulty. [6, 1, 8, 5] also studied CPA problem under this fault model. Yet, to the best of our knowledge, the tight necessary and sufficient conditions for the correctness of CPA in synchronous arbitrary point-to-point networks under  $f$ -local fault model have not been developed previously.

### 3. Feasibility of CPA under $f$ -local fault model

#### 3.1. Certified Propagation Algorithm (CPA)

In this subsection, we describe the Certified Propagation Algorithm (CPA) from [6] formally. Note that the faulty nodes may deviate from this specification arbitrarily. Possible misbehavior includes sending incorrect and mismatching (or inconsistent) messages to different outgoing neighbors.

Source node  $s$  commits to its input  $x_s$  at the start of the algorithm, i.e., sets its output equal to  $x_s$ . The source node is said to have committed to  $x_s$

in round 0. The algorithm for each round  $r$  ( $r > 0$ ), is as follows:

1. Each node that commits in round  $r - 1$  to some value  $x$ , transmits message  $x$  to all its outgoing neighbors, and then terminates.
2. If any node receives message  $x$  directly from source  $s$ , it commits to output  $x$ .
3. Through round  $r$ , if a node has received messages containing value  $x$  from at least  $f + 1$  distinct incoming neighbors, then it commits to output  $x$ .

### 3.2. The Necessary Condition

For CPA to be correct, the network graph  $G(\mathcal{V}, \mathcal{E})$  must satisfy the necessary condition proved in this section. We borrow two relations  $\Rightarrow$  and  $\not\Rightarrow$  from our previous paper [11].

**Definition 1.** For non-empty disjoint sets of nodes  $A$  and  $B$ ,

- $A \Rightarrow B$  iff there exists a node  $v \in B$  that has at least  $f + 1$  distinct incoming neighbors in  $A$ , i.e.,  $|N_v^- \cap A| > f$ .
- $A \not\Rightarrow B$  iff  $A \Rightarrow B$  is not true.

**Definition 2.** Set  $F \subseteq \mathcal{V}$  is said to be a feasible  $f$ -local fault set, if for each node  $v \notin F$ ,  $F$  contains at most  $f$  incoming neighbors of node  $v$ . That is, for every  $v \in \mathcal{V} - F$ ,  $|N_v^- \cap F| \leq f$ .

We now derive the necessary condition on the network topology.

**Theorem 1.** *Suppose that CPA is correct in graph  $G(\mathcal{V}, \mathcal{E})$  under the  $f$ -local fault model. Let sets  $F, L, R$  form a partition<sup>4</sup> of  $\mathcal{V}$ , such that (i) source  $s \in L$ , (ii)  $R$  is non-empty, and (iii)  $F$  is a feasible  $f$ -local fault set. Then*

- $L \Rightarrow R$ , or
- $R$  contains an outgoing neighbor of  $s$ , i.e.,  $N_s^+ \cap R \neq \emptyset$ .

*Proof.* Consider any partition  $F, L, R$  such that  $s \in L$ ,  $R$  is non-empty, and  $F$  is a feasible  $f$ -local fault set. Suppose that the input at  $s$  is  $x_s$ . Consider any single execution of the CPA algorithm such that the nodes in  $F$  behave as if they have crashed.

By assumption, CPA is correct in the given network under such a behavior by the faulty nodes. Thus, all the fault-free nodes eventually commit their output to  $x_s$ . Let round  $r$  ( $r > 0$ ), be the earliest round in which at least one of the nodes in  $R$  commits to  $x_s$ . Let  $v$  be one of the node in  $R$  that commits in round  $r$ . Such a node  $v$  must exist since  $R$  is non-empty, and it does not contain source node  $s$ . For node  $v$  to be able to commit, as per specification of the CPA algorithm, either node  $v$  should receive the message  $x_s$  directly from the source  $s$ , or node  $v$  must have  $f + 1$  distinct incoming neighbors that have already committed to  $x_s$ . By definition of node  $v$ , nodes that have committed to  $x_s$  prior to  $v$  must be outside  $R$ ; since nodes in  $F$  behave as crashed, these  $f + 1$  nodes must be in  $L$ . Thus, either  $(s, v) \in \mathcal{E}$ , or node  $v$  has at least  $f + 1$  distinct incoming neighbors in set  $L$ .

□

---

<sup>4</sup>Sets  $X_1, X_2, X_3, \dots, X_p$  are said to form a partition of set  $X$  provided that (i)  $\cup_{1 \leq i \leq p} X_i = X$ , and (ii)  $X_i \cap X_j = \emptyset$  if  $i \neq j$ .

### 3.3. Sufficiency

We now show that the condition in Theorem 1 is also sufficient.

**Theorem 2.** *If  $G(\mathcal{V}, \mathcal{E})$  satisfies the condition in Theorem 1, then CPA is correct in  $G(\mathcal{V}, \mathcal{E})$  under the  $f$ -local fault model.*

*Proof.* Suppose that  $G(\mathcal{V}, \mathcal{E})$  satisfies the condition in Theorem 1. Let  $F'$  be the set of faulty nodes. By assumption,  $F'$  is a feasible local fault set. Let  $x_s$  be the input at source node  $s$ . We will show that, (i) fault-free nodes do not commit to any value other than  $x_s$  (Validity), and, (ii) until all the fault-free nodes have committed, in each round of CPA, at least one additional fault-free node commits to value  $x_s$  (Termination). The proof is by induction.

*Induction basis:* Source node  $s$  commits in round 0 to output equal to its input  $x_s$ . No other fault-free nodes commit in round 0.

*Induction:* Suppose that  $L$  is the set of fault-free nodes that have committed to  $x_s$  through round  $r$ ,  $r \geq 0$ . Thus,  $s \in L$ . Define  $R = \mathcal{V} - L - F'$ . If  $R = \emptyset$ , then the proof is complete. Let us now assume that  $R \neq \emptyset$ .

Now consider round  $r + 1$ .

- Validity:

Consider any fault-free node  $u$  that has not committed prior to round  $r + 1$  (i.e.,  $u \in R$ ). All the nodes in  $L$  have committed to  $x_s$  by the end of round  $r$ . Thus, in round  $r + 1$  or earlier, node  $u$  may receive messages containing values different from  $x_s$  only from nodes in  $F'$ . Since there are at most  $f$  incoming neighbors of  $u$  in  $F'$ , node  $u$  cannot commit to any value different from  $x_s$  in round  $r + 1$ .

- Termination:

By the condition in Theorem 1, there exists a node  $w$  in  $R$  such that (i) node  $w$  has an incoming link from  $s$ , or (ii) node  $w$  has incoming links from  $f + 1$  nodes in  $L$ . In case (i), node  $w$  will commit to  $x_s$  on receiving  $x_s$  from node  $s$  in round  $r + 1$  (in fact,  $r + 1$  in this case must be 1). In case (ii), since all the nodes in  $L$  from whom node  $w$  has incoming links have committed to  $x_s$  (by definition of  $L$ ), node  $w$  will be able to commit to  $x_s$  after receiving messages from at least  $f + 1$  incoming neighbors in  $L$ , since all nodes in  $L$  have committed to  $x_s$  by the end of round  $r$  by the definition of  $L$ .<sup>5</sup> Thus, node  $w$  will commit to  $x_s$  in round  $r + 1$ .

This completes the proof. □

#### 4. CPA without prior knowledge of $f$

In this section, we present a parameter-independent algorithm CPA-P that does not require prior knowledge of  $f$ , and each node only needs to know  $n$ , the number of nodes in the system. That is, given a graph  $G$  that can tolerate  $f$ -local faults (where  $f$  is unknown), the algorithm CPA-P presented below solves the broadcast problem in  $G$  without usage of  $f$ .

The core idea of CPA-P is for each node to exhaustively test all possible parameters by running  $n + 1$  instances of CPA algorithm in parallel. Each instance of CPA algorithm corresponds to a tested parameter ranging from

---

<sup>5</sup>Since node  $w$  did not commit prior to round  $r + 1$ , it follows that at least one node in  $L$  must have committed in round  $r$ .

0 to  $n$ . That is, each instance assumes that the tested parameter is the real bound ( $f$ ) on the local faults at each node.<sup>6</sup> The correctness of CPA-P is based on the following observation: For each fault-free node, when the tested parameter is larger than or equal to the real parameter  $f$ , then there are only two outcomes: (i) it cannot commit, since it did not receive enough identical messages (violating Step 3 in CPA as specified in 3.1), or (ii) it commits to a correct value, i.e., the input value of the source. Thus, in the end of the CPA-P,<sup>7</sup> each node can simply commit to the non-null value corresponding to the largest tested parameter. Now, we describe CPA-P formally.

Throughout the execution, each node  $i$  (excluding  $s$  and outgoing neighbors of  $s$ ) maintains an  $(n+1)$ -entry vector  $v_i$ , where  $v_i[t]$  ( $0 \leq t \leq n$ ) is the estimate of output corresponding to the tested parameter  $t$ . In the beginning of the algorithm, every entry of vector  $v_i$  is initialized to be a null value  $\perp$ , where  $\perp$  is distinguished from all possible values of  $x_s$ .

Source node  $s$  commits to its input  $x_s$  at the start of the algorithm (round 0), and transmits message  $x_s$  to all its outgoing neighbors in round 1. For the other nodes, the algorithm is as follows.

- For outgoing neighbor of the source  $s$ :
  1. In round 1, it receives message  $x$  directly from source  $s$ , and commits to output  $x$ .

---

<sup>6</sup>For simplicity of presentation, we assume that every node keeps track of  $n+1$  instances (of the CPA algorithm) at the same time, even if the node already knows that some instances cannot terminate, since it may never receive enough identical messages if the tested parameter is too large. In a real implementation, each node  $i$  only needs to keep track of  $\lceil \frac{d_i}{2} \rceil - 1$  instances of CPA algorithm, where  $d_i$  is the number of incoming neighbors at node  $i$ .

<sup>7</sup>Note that CPA is guaranteed to terminate in  $n$  steps, and so is CPA-P.

2. In round 2, it transmits messages  $\langle x, 0 \rangle, \langle x, 1 \rangle, \dots, \langle x, n \rangle$  to all its outgoing neighbors, and terminates.
- For node that is not an outgoing neighbor of  $s$ , in each round  $r$  ( $r > 0$ ):
    1. For  $0 \leq t \leq n$ , each node  $i$  that sets  $v_i[t]$  in round  $r - 1$  to some value  $x$ , transmits message  $\langle x, t \rangle$  to all its outgoing neighbors.
    2. For  $0 \leq t \leq n$ , through round  $r$ , if a node  $i$  has received messages containing value  $\langle x, t \rangle$  from at least  $t + 1$  distinct incoming neighbors, then it sets  $v_i[t] = x$ .
    3. In round  $n$ , each node  $i$  commits to value  $v_i[t']$ , where  $t'$  is the largest value in range  $[0, n]$  such that  $v_i[t'] \neq \perp$ .

Note that the algorithm performs  $n$  rounds.

Now, we show that CPA-P is correct.

**Theorem 3.** *Given a graph  $G$  that can tolerate  $f$ -local faults, CPA-P achieves both validity and termination.*

*Proof.* Denote by CPA-P- $t$  ( $0 \leq t \leq n$ ) the instance of CPA-P corresponding to the tested parameter  $t$ . Then by assumption of  $G$ , CPA-P- $f$  is correct. Thus, for each fault-free node  $i$ ,  $v_i[f] = x_s$ , the input value at source  $s$ . Now, we prove the following claim:

**Claim 1.** *For  $t > f$ , in CPA-P- $t$ , fault-free nodes never decide on an invalid value, i.e., for each fault-free node  $i$ , either  $v_i[t] = x_s$  or  $v_i[t] = \perp$ .*

*Proof.* The proof is by induction.

*Induction basis:* Source node  $s$  and its outgoing neighbors commit to output equal to the source's input  $x_s$  in round 0 and 1, respectively. No other fault-free nodes commit in round 0 and 1.

*Induction:* Suppose that  $L$  is the set of fault-free nodes that have committed to  $x_s$  through round  $r$  ( $r > 0$ ). Thus,  $s \in L$ . Let  $F'$  be the set of faulty nodes, and  $|F'| = f$ . Define  $R = \mathcal{V} - L - F'$ . If  $R = \emptyset$ , then the proof is complete. Let us now assume that  $R \neq \emptyset$ .

Now consider round  $r + 1$ .

Consider any fault-free node  $u$  that has not committed prior to round  $r + 1$  (i.e.,  $u \in R$ ). All the nodes in  $L$  have committed to  $x_s$  by the end of round  $r$ . Thus, in round  $r + 1$  or earlier, node  $u$  may receive messages containing values different from  $x_s$  only from nodes in  $F'$ . Therefore, node  $u$  cannot commit to any value different from  $x_s$  in round  $r + 1$ , since by assumption  $|N_u^- \cap F'| \leq f < t$ .

Unlike the proof in Theorem 2, node  $u$  may never gather enough (i.e., at least  $t + 1$ ) identical messages from its incoming neighbors, since  $t > f$ . Thus, for CPA-P- $t$ , node  $u$  may never terminate. In this case,  $v_u[t] = \perp$ .  $\square$

The source node  $s$  and fault-free outgoing neighbors of  $s$  commit to  $x_s$  in round 0 and 1, respectively. By Claim 1 and the fact that CPA-P- $f$  satisfies both validity and termination, each fault-free node  $i$  (excluding  $s$  and outgoing neighbors of  $s$ ) commits to  $x_s$ . Thus, CPA-P is correct.  $\square$

## 5. Discussion

This section discusses some extensions on the result presented above.

### 5.1. Generalized Fault Model

In this subsection, we briefly discuss how to extend the above results under a generalized fault model. The generalized fault model [10] is characterized using *fault domain*  $\mathcal{F} \subseteq 2^{\mathcal{V}}$  as follows: Nodes in set  $F$  may fail during an execution of the algorithm only if there exists set  $F^* \in \mathcal{F}$  such that  $F \subseteq F^*$ . Set  $F$  is then said to be a *feasible* fault set.

**Definition 3.** *Set  $F \subseteq \mathcal{V}$  is said to be a feasible fault set, if there exists  $F^* \in \mathcal{F}$  such that  $F \subseteq F^*$ .*

For a set of nodes  $B$ , define  $N^-(B) = \{i \mid (i, j) \in \mathcal{E}, i \notin B, j \in B\}$ , the set of incoming neighbors of  $B$ .

**Definition 4.** *Given  $\mathcal{F}$ , for disjoint sets of nodes  $A$  and  $B$ , where  $B$  is non-empty.*

- $A \stackrel{g}{\Rightarrow} B$  iff for every  $F^* \in \mathcal{F}$ ,  $N^-(B) \cap A \not\subseteq F^*$ .
- $A \not\stackrel{g}{\Rightarrow} B$  iff  $A \stackrel{g}{\Rightarrow} B$  is not true.

Under the generalized fault model, step 3 of CPA needs to be modified as follows. Let us call the modified algorithm CPA-G.

3. Through round  $r$ , if a node has received messages containing value  $x$  from a set  $M$ , where  $M$  is not a feasible fault set, then the node commits to value  $x$ .

It is easy to show that a modified version of Theorem 1 stated below holds for the generalized fault model.

**Theorem 4.** *Suppose that CPA-G is correct in graph  $G(\mathcal{V}, \mathcal{E})$  under the generalized fault model. Let sets  $F, L, R$  form a partition of  $\mathcal{V}$ , such that source (i)  $s \in L$ , (ii)  $R$  is non-empty, and (iii)  $F$  is a feasible fault set, then*

- $L \Rightarrow R$ , or
- $R$  contains an outgoing neighbor of  $s$ , i.e.,  $N_s^+ \cap R \neq \emptyset$ .

## 5.2. Broadcast Channel

We have so far assumed that the underlying network is a point-to-point network. The results, however, can be easily extended to the *broadcast* or *radio model* [6, 1] as well. In the *broadcast model*, when a node transmits a value, all of its outgoing neighbors receive this value identically. Thus, no node can transmit mismatching values to different outgoing neighbors. Then, it is easy to see that the same condition as the point-to-point network can be shown to be necessary and sufficient for CPA under the broadcast model as well.

Now consider the following variation of the CPA algorithm: if the outgoing neighbors of source  $s$  do not receive a message from  $s$  in round 1, the message value is assumed to be some default value. With this modification, the condition in Theorem 1 can also be shown to be necessary and sufficient to perform Byzantine Broadcast [7] under the broadcast model, while satisfying the following three conditions (allowing  $s$  to be faulty):

- **Termination:** every fault-free node  $i$  eventually decides on an output value  $y_i$ .

- **Agreement:** the output values of all the fault-free nodes are equal, i.e., there exists  $y$  such that, for every fault-free node  $i$ ,  $y_i = y$ .
- **Validity:** if the source node is fault-free, then for every fault-free node  $i$ , the output value equals the source's input, i.e.,  $y = x_s$ .

The proof follows from the proof of Theorem 1 and the observation that if  $s$  transmits a value, then all the outgoing neighbors of  $s$  receive identical value from  $s$ , which equals its input  $x_s$  when  $s$  is fault-free.

### 5.3. Asynchronous Network

In our analysis so far, we have assumed that the system is synchronous. For a point-to-point network with fault-free source  $s$ , it should be easy to see that the condition in Theorem 1 is also necessary and sufficient to achieve agreement using a CPA-like under the asynchronous model [2] as well. In this case, the algorithm may not proceed in rounds, but a node still commits to value  $x$  either on receiving the value directly from  $s$ , or from  $f + 1$  nodes.

This claim may seem to contradict the FLP result [4]. However, our claim assumes that the source node is fault-free, unlike [4].

## 6. Conclusion

In this paper, we explore broadcast in arbitrary network using the CPA algorithm in  $f$ -local fault model. In particular, we provide a *tight* necessary and sufficient condition on the underlying network for the correctness of CPA.

## References

## References

- [1] V. Bhandari and N. H. Vaidya. On reliable broadcast in a radio network: A simplified characterization. Technical report, University of Illinois at Urbana-Champaign, 2005.
- [2] D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl. Reaching approximate agreement in the presence of faults. *J. ACM*, 33:499–516, May 1986.
- [3] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge, 2010.
- [4] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32:374–382, April 1985.
- [5] A. Ichimura and M. Shigeno. A new parameter for a broadcast algorithm with locally bounded byzantine faults. *Inf. Process. Lett.*, 110(12-13):514–517, June 2010.
- [6] C.-Y. Koo. Broadcast in radio networks tolerating byzantine adversarial behavior. In *Proc. 23rd Annual ACM Symp. on Principles of Distributed Computing (PODC' 04)*, 2004.
- [7] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, Apr. 1980.

- [8] A. Pelc and D. Peleg. Broadcasting with locally bounded byzantine faults. *Inf. Process. Lett.*, 93(3):109–115, Feb. 2005.
- [9] N. B. Shah, K. V. Rashmi, and K. Ramchandran. Efficient and distributed secret sharing in general networks. CoRR abs/1207.0120, 2012.
- [10] L. Tseng and N. H. Vaidya. Iterative approximate byzantine consensus under a generalized fault model. International Conference on Distributed Computing and Networking, 2013. [http://www.crhc.illinois.edu/wireless/papers/ICDCN\\_general.pdf](http://www.crhc.illinois.edu/wireless/papers/ICDCN_general.pdf)
- [11] N. H. Vaidya, L. Tseng, and G. Liang. Iterative approximate byzantine consensus in arbitrary directed graphs. In *Proceedings of the thirty-first annual ACM symposium on Principles of distributed computing*, PODC '12. ACM, 2012.