



HAL
open science

On Byzantine Broadcast in Planar Graphs

Alexandre Maurer, Sébastien Tixeuil

► **To cite this version:**

Alexandre Maurer, Sébastien Tixeuil. On Byzantine Broadcast in Planar Graphs. 2013. hal-00773343v1

HAL Id: hal-00773343

<https://hal.sorbonne-universite.fr/hal-00773343v1>

Submitted on 13 Jan 2013 (v1), last revised 7 Dec 2013 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Byzantine Broadcast in Planar Graphs

Alexandre Maurer and Sébastien Tixeuil
UPMC Sorbonne Universités, Paris, France
E-mail: Alexandre.Maurer@lip6.fr, Sebastien.Tixeuil@lip6.fr
Phone: +33 1 44 27 87 62, +33 1 44 27 87 75

January 13, 2013

Abstract

We consider the problem of reliably broadcasting information in a multihop asynchronous network in the presence of Byzantine failures: some nodes have an unpredictable malicious behavior. We focus on non-cryptographic solutions. Very few Byzantine-robust algorithms exist for loosely connected networks. A recent algorithm guarantees reliable broadcast on a lattice when $D > 4$, D being the minimal distance between two Byzantine nodes.

In this paper, we generalize this result to planar graphs where edges delimit convex polygons. We show that reliable broadcast is guaranteed when $D > Z$, Z being the maximal number of edges per polygon. We also show that no algorithm can improve this bound. In a second part, we assume that the delay between two activations of a correct process has an upper and lower bound. We show that the good information is delivered within a linear time. We also show that reliable broadcast is still guaranteed with a finite memory (exhaustion safety).

1 Introduction

As modern networks grow larger, they become more likely to fail. Indeed, some nodes can be subject to crashes, attacks, transient bit flips, etc. The most general model of failures is the *Byzantine* model [11], where the failing nodes have a totally arbitrary behavior. This model encompasses all other possible types of failures, and has important security applications.

In this paper, we study the problem of reliably broadcasting information in a multihop network. In the ideal case, the source node transmits the information to its neighbors, that transmit it to their neighbors, and so forth. However, a Byzantine node can broadcast a false information and lie to the entire network. Our goal is to design a more elaborated solution, that guarantees reliable broadcast in the presence of Byzantine failures.

Related works. Many Byzantine-robust protocols are based on *cryptography* [3, 5]: the nodes use digital signatures to authenticate the sender across multiple hops. However, as the malicious nodes are supposed to ignore some cryptographic secrets, their behavior is not completely arbitrary. Besides, manipulating cryptographic operations requires important resources, which may not always be available. At last, cryptography requires a trusted infrastructure that initially distributes public and private keys: if this initial infrastructure fails, the whole network fails. Yet, we want to design a solution where *any* element can fail without destabilizing the whole system. For these reasons, we focus on cryptography-free solutions.

Cryptography-free solutions have first been studied in completely connected networks [11, 1, 12, 13, 19]: a node can directly communicate with any other node, which implies the presence of a channel between each pair of nodes. Therefore, these approaches are hardly scalable, as the number of channels per node can be physically limited. We thus study solutions in multihop networks, where a node must rely on other nodes to broadcast informations.

A notable class of algorithms tolerates Byzantine failures with either space [15, 20, 23] or time [14, 9, 8, 7, 6] locality. Yet, the emphasis of space local algorithms is on containing the fault as close to its source as possible. This is only applicable to the problems where the information from remote nodes is unimportant: vertex coloring, link coloring, dining philosophers, etc. Also, time local algorithms presented so far can hold at most one Byzantine node and are not able to mask the effect of Byzantine actions. Thus, this approach is not applicable to reliable broadcast.

In [4], it was shown that, for agreement in the presence of up to k Byzantine nodes, it is necessary and sufficient that the network is $(2k + 1)$ -connected, and that the number of nodes in the system is at least $3k + 1$. Also, this solution assumes that the topology is known to every node, and that nodes are scheduled according to the synchronous execution model. Both requirements have been relaxed in [21]: the topology is unknown and the scheduling is asynchronous. Yet, this solution retains $2k + 1$ connectivity for reliable broadcast and $k + 1$ connectivity for detection.

Another existing approach is based on the fraction of Byzantine neighbors per node. Solutions have been proposed for nodes organized on a grid [10, 2]. Both approaches are based on a local voting system, and perform correctly if every node has strictly less than a $1/4$ fraction of Byzantine neighbors. This result was later generalized to other topologies [22], assuming that each node knows the global topology.

All aforementioned approaches are hardly applicable to sparse network, such as planar graph networks. For instance, in a lattice network (see Figure 1), no existing solution can tolerate more than one Byzantine node. New solutions have been proposed for such networks [16, 18], but only give probabilistic guarantees, and require the nodes to know their position on the network. This requirement was finally relaxed in [17]: reliable broadcast is guaranteed on a lattice when $D > 4$, D being the minimal number of hops between two Byzantine nodes.




Figure 1: Lattice (left) and planar graph considered in the paper (right)

Our contribution. In this paper, we generalize the result of [17] to planar graphs where edges delimit convex polygons (see Figure 1). We show that reliable broadcast is guaranteed when $D > Z$, Z being the maximal number of edges per polygon. We also show the tightness of this bound: if we only have $D \geq Z$, no algorithm can guarantee reliable broadcast.

In a second part, we make a weak time assumption: the delay between two activations of a correct process has an upper and lower bound. We show that the good information is delivered within a linear time, like in a simple broadcast protocol. We also show that, unlike the previous solutions [10, 2, 22, 16, 18, 17], reliable broadcast is still guaranteed with a finite memory in both nodes and channels.

The hypotheses and the protocol are described in Section 2. The reliability properties are established in Section 3, and the time properties are established in Section 4.

2 Setting

In this section, we present our hypotheses and describe the broadcast protocol.

2.1 Hypotheses

Topology Let $\mathcal{G} = (G, E)$ be a graph representing the topology of the network. G denotes the *nodes*, and E denotes the *edges* connecting two nodes. The graph \mathcal{G} is *planar*: there exists a bidimensional representation of this graph where edges do not cross each other, and therefore delimit polygons. The edges delimit strictly convex polygons, and each edge connects 2 polygons (see Figure 1). This second condition implies that the network has no border, and therefore is infinite. However, our results can easily be extended to finite bidimensional spaces, like the sphere or the torus.

The “strictly convex” hypothesis is a sufficient condition that is only used in the proof of Lemma 3, to obtain a node-disjoint circular path. We could use a tighter (yet more complex) condition, but we chose this one to keep the claims simple.

Let $Z \geq 3$ be the maximal number of edges per polygon, and let Y be the maximal number of edges per node. We assume that Z is known by the correct nodes.

Network Two nodes (or *processes*) connected by an edge (or *channel*) are called *neighbors*. A node can only send messages to its neighbors. Some nodes are *correct* and follow the protocol described thereafter. The other nodes are *Byzantine*, and have a totally unpredictable behavior. The correct nodes do not know which nodes are Byzantine.

We consider an *asynchronous* network: any message sent is eventually received, but it can be at any time. We assume that, in an infinite execution, any process is activated infinitely often; however, we make no hypothesis on the order of activation of the processes. Finally, we assume *authenticated channels* (or “oral” model): each node has a unique identifier, and knows the identifier of its neighbors. Therefore, when a node receives a message from a neighbor p , it knows that p is the author of the message.

2.2 Protocol

Preliminaries An arbitrary correct node, called the *source*, wants to broadcast an information m_0 in the network. We say that a correct node *multicasts* a message when it sends it to all its neighbors, and *delivers* m when it permanently considers that m was broadcast by the source. We say that we achieve *reliable broadcast* if all correct nodes eventually deliver m_0 .

Principle of the protocol We use the same principle that in [17]: to deliver an information, a node must receive it from a direct neighbor q , but also from another node located at at most $Z - 2$ hops. The intuitive idea is that, if two Byzantine nodes are distant from more than Z hops, they can never cooperate to make a correct node deliver a false information.

In this paper, we use a modified version of the protocol of [17]: instead of storing all the received message in a set Rec , a correct node uses a variable $Rec(q)$ for each neighbor q , storing the last message received from q . This modification makes the correct nodes *exhaustion-safe* (see 4.2).

The messages exchanged in the protocol are tuples of the form (m, S) , where m is the information broadcast by the source (or pretending to be it), and S is a set containing the identifiers of the nodes already visited by the message.

Description of the protocol

- The source multicasts an arbitrary information m_0 .
- The correct nodes that are neighbors of the source wait until they receive an information m from the source, then deliver m and multicast (m, \emptyset) .
- The other correct nodes have the following behavior:
 - When (m, S) is received from a neighbor q , with $q \notin S$ and $\text{card}(S) \leq Z - 3$: assign the value (m, S) to $\text{Rec}(q)$ and multicast $(m, S \cup \{q\})$.
 - When there exists m, p, q and S such that $q \neq p$, $q \notin S$, $\text{Rec}(q) = (m, \emptyset)$ and $\text{Rec}(p) = (m, S)$: deliver m , multicast (m, \emptyset) and stop.

3 Reliability properties

In this section, we prove the main result of the paper: if $D > Z$, we achieve reliable broadcast. We also show that this bound is tight: if we only have $D \geq Z$, no algorithm can guarantee reliable broadcast.

3.1 Definitions

Definition 1 (Path and circular path). *A path is a sequence of nodes (u_1, \dots, u_n) such that u_i and u_{i+1} are neighbors. This path is circular if $u_1 = u_n$. Notice that, in our definition, we do not require that the nodes are distinct.*

Definition 2 (Neighbor and adjacent polygons). *Two polygons are neighbors if they share at least one node, and adjacent if they share an edge.*

Definition 3 (Polygonal path). *A polygonal path is a sequence of polygons (P_1, \dots, P_n) such that P_i and P_{i+1} are adjacent.*

Definition 4 (Connected polygons). *A set S of polygons is connected if, for each pair of polygons (P, Q) of S , there exists a polygonal path (P, P_1, \dots, P_n, Q) in S .*

Definition 5 (Correct and Byzantine polygons). *A polygon is correct if all its nodes are correct. Otherwise, it is Byzantine.*

3.2 Main theorem

Let us show that, if $D > Z$, we achieve reliable broadcast (Theorem 1).

Lemma 1. *Let us suppose that $D > Z$. Then, if two polygons are neighbors, the set of their nodes contains at most one Byzantine node.*

Proof. The proof is by contradiction. Let us suppose the opposite: there exists two neighbor polygons P and Q , and the set of their nodes contains two distinct Byzantine nodes b_1 and b_2 .

As P and Q are neighbors, let u be a node shared by P and Q . Let (u, p_1, \dots, p_n, u) be a circular path on P , and let (u, q_1, \dots, q_m, u) be a circular path on Q . Therefore, $(u, p_1, \dots, p_n, u, q_1, \dots, q_m, u)$ is a circular path containing all the nodes of P and Q .

As this circular path contains at most $2Z$ hops, two nodes of this path are distant of at most Z hops. In particular, b_1 and b_2 are distant of at most Z hops, which contradicts $D > Z$. Thus, the result. \square

Lemma 2. *Each correct node belongs to at least one correct polygon.*

Proof. Let us suppose the opposite: there exists a correct node v that does not belong to any correct polygon. Let (e_1, \dots, e_n) be the edges connected to v , clockwise, with $e_1 = e_n$. Let P_i be the polygon containing e_i and e_{i+1} . As P_i and P_{i+1} are Byzantine, according to Lemma 1, they share the same Byzantine node b . Therefore, by induction, $\{P_1, \dots, P_{n-1}\}$ share the same Byzantine node b . As v is the only node shared by all the polygons $\{P_1, \dots, P_{n-1}\}$, $b = v$, and v is Byzantine: contradiction. Thus, the result. \square

Lemma 3. *Let v be a node, and let V be the set of polygons containing v . Let S be the set of polygons that are not in V , but are neighbors with a polygon of V . Then, S is connected.*

Proof. Let (e_1, \dots, e_n) be the edges connected to v , clockwise, with $e_1 = e_n$. Let u_i be the node connected to v by e_i . Let P_i be the polygon containing the edges e_i and e_{i+1} . At last, let $(u_i, p_1^i, p_2^i, \dots, u_{i+1})$ be a path on P_i , defined clockwise, that does not contain v . We define a circular path $(q_1, \dots, q_m) = (u_1, p_1^1, p_2^1, \dots, u_2, p_1^2, p_2^2, \dots, u_n)$, containing all the nodes of V except v . As the polygons are strictly convex, the angle (q_1, v, q_i) strictly increases with i . Therefore, the nodes $\{q_1, \dots, q_m\}$ are distinct.

Let us consider a given index $i \in \{1, \dots, m\}$. Let $(e_1^i, \dots, e_{n(i)}^i)$ be the edges connected to q_i , clockwise, such that e_1^i is connected to q_{i-1} and $e_{n(i)}^i$ is connected to q_{i+1} . Let P_k^i be the polygon containing e_k^i and e_{k+1}^i . As the circular path is defined clockwise, P_k^i does not belong to V . Therefore, as P_k^i shares q_i with a polygon of V , P_k^i belongs to S . As P_k^i and P_{k+1}^i are adjacent, $S_i = \{P_1^i, \dots, P_{n(i)-1}^i\}$ is connected.

As S_i and S_{i+1} have the polygon $P_{n(i)-1}^i = P_1^{i+1}$ in common, $S_i \cup S_{i+1}$ is connected. Therefore, $S = S_1 \cup \dots \cup S_{m-1}$ is connected. \square

Lemma 4. *Let us suppose that $D > Z$. Let (P, P_1, \dots, P_n, Q) be a polygonal path such that P and Q are correct, and $\{P_1, \dots, P_n\}$ are Byzantine. Then, there exists a polygonal path (P, Q_1, \dots, Q_m, Q) such that $\{Q_1, \dots, Q_m\}$ are correct.*

Proof. According to Lemma 1, P_i and P_{i+1} share the same Byzantine node b . Therefore, by induction, the polygons $\{P_1, \dots, P_n\}$ share the same Byzantine node b .

Let V be the set of polygons containing b , and let S be the set of polygons that are not in V , but are neighbor with a polygon of V . As V contains P_1 and P_n , by definition, S contains P and Q . According to Lemma 3, S is connected: there exists a polygonal path (P, Q_1, \dots, Q_m, Q) in S . To complete the proof, let us show that the polygons of S are correct.

Let us suppose the opposite: a polygon P^* of S is Byzantine. Let b^* be the Byzantine node contained by P^* . Then, as P^* has a neighbor polygon in V , according to Lemma 1, $b^* = b$. It implies that P^* belongs to V : contradiction. Thus, the polygons of S are correct. Thus, the result. \square

Lemma 5. *If $D > Z$, the set of correct polygons is connected.*

Proof. Let P and Q be two correct polygons, and let (P, P_1, \dots, P_n, Q) be a polygonal path. If $\{P_1, \dots, P_n\}$ are correct, the result is trivial. Otherwise, let us consider the following process.

Let N be the smallest integer such that P_N is Byzantine, and let M be the smallest integer greater than N such that P_{M+1} is correct. Then, according to Lemma 4, there exists a polygonal path $(P_{N-1}, Q_1, \dots, Q_m, P_{M+1})$ such that the polygons $\{Q_1, \dots, Q_m\}$ are correct. Therefore, we can replace the sequence (P_N, \dots, P_M) by (Q_1, \dots, Q_m) . We repeat this process until all the polygons of the path are correct. \square

Lemma 6. *Let us suppose that $D \geq Z$. Then, if a correct node delivers an information, it is necessarily m_0 .*

Proof. The proof is by contradiction. Let us suppose the opposite: $D \geq Z$, yet at least one correct node delivers $m' \neq m_0$. Let u be the first correct node to deliver m' . It implies that there exists p , q and S such that $q \neq p$, $q \notin S$, $Rec(q) = (m', \emptyset)$ and $Rec(p) = (m', S)$.

$Rec(q) = (m', \emptyset)$ implies that u received (m', \emptyset) from a neighbor q . Let us suppose that q is correct. Then, as q sent (m', \emptyset) , it implies that q delivered m' . This is impossible, as u is the first correct node to deliver m' . So q is necessarily Byzantine. Besides, according to the protocol, $Rec(p) = (m', S)$ implies that $card(S) \leq Z - 3$.

Let us prove the following property \mathcal{P}_i by induction, for $0 \leq i \leq card(S)$: a correct node p_i , located at $i + 2$ hops or less from q , sent (m', S_i) with $card(S_i) = card(S) - i$.

- First, let us show that \mathcal{P}_0 is true. $Rec(p) = (m', S)$ implies that p sent (m', S) . Let us suppose that p is Byzantine. Then, as q is also Byzantine, $D \leq 2$, which is impossible as $D \geq Z \geq 3$. So p is necessarily correct, and \mathcal{P}_0 is true if we take $p_0 = p$ and $S_0 = S$. If $Z = 3$, ignore the following step.
- Let us suppose that \mathcal{P}_i is true, with $i < card(S)$. As $card(S_i) = card(S) - i \geq 1$, p_i necessarily received (m', S_{i+1}) from a node p_{i+1} located at $i+3$ hops or less from q , with $S_i = S_{i+1} \cup \{p_{i+1}\}$ and $p_{i+1} \notin S_{i+1}$. Thus, we have $card(S_{i+1}) = card(S_i) - 1 = card(S) - i - 1$. Let us suppose that p_{i+1} is Byzantine. Then, as q is also Byzantine, $D \leq i + 3 \leq card(S) + 2 < Z$, which is impossible as $D \geq Z$. So p_{i+1} is necessarily correct, and \mathcal{P}_{i+1} is true.

Therefore, $\mathcal{P}_{card(S)}$ is true, and $p_{card(S)}$ sent (m', \emptyset) , as $card(S_{card(S)}) = card(S) - card(S) = 0$. According to the protocol, it implies that $p_{card(S)}$ delivered m' before u , which contradicts our initial hypothesis. Thus, the result. \square

Lemma 7. *Let us suppose that $D \geq Z$. Let (u_1, \dots, u_n) be a path of distinct correct nodes, with $3 \leq n \leq Z$, such that u_1 and u_n deliver m_0 . Then, at least one of the nodes u_2 and u_{n-1} delivers m_0 .*

Proof. As u_1 and u_n deliver m_0 , and therefore multicast (m_0, \emptyset) , let E_1 and E_2 be the two following events: (E_1) u_2 receives (m_0, \emptyset) from u_1 and (E_2) u_{n-1} receives (m_0, \emptyset) from u_n . Let us suppose that E_2 is the first event to occur. As u_n delivers m_0 , according to the protocol, u_n stops. Therefore, for the node u_{n-1} , $Rec(u_n) = (m_0, \emptyset)$ until the end of the execution.

Let us prove the following property \mathcal{P}_i by induction, for $1 \leq i \leq n - 2$: u_i multicasts (m_0, S_i) , with $S_i \subseteq \{u_1, \dots, u_{n-2}\}$ and $card(S_i) \leq i - 1$.

- As u_1 delivers m_0 , u_1 multicasts (m_0, \emptyset) . Therefore, \mathcal{P}_1 is true if we take $S_0 = \emptyset$
- Let us suppose that \mathcal{P}_i is true, for $i < n - 2$. Then, u_{i+1} receives (m_0, S_i) from u_i , with $card(S_i) \leq i - 1 < n - 3 \leq Z - 3$. When it does, two possibilities:
 - If u_{i+1} has stopped, u_{i+1} has necessarily delivered an information. As $D \geq Z$, according to Lemma 6, this information was m_0 . Thus, according to the protocol, u_{i+1} has already multicast (m_0, \emptyset) , and \mathcal{P}_{i+1} is true if we take $S_{i+1} = \emptyset$.
 - Otherwise, as $card(S_i) \leq Z - 3$, u_{i+1} multicasts $(m_0, S_i \cup \{u_i\})$. Thus, \mathcal{P}_{i+1} is true if we take $S_{i+1} = S_i \cup \{u_i\}$.

Therefore, \mathcal{P}_{n-2} is true, and u_{n-1} receives (m_0, S_{n-2}) from u_{n-2} , with $S_{n-2} \subseteq \{u_1, \dots, u_{n-2}\}$ and $card(S_{n-2}) \leq n - 3 \leq Z - 3$. Thus, for the node u_{n-1} , $Rec(u_{n-2}) = (m_0, S_{n-2})$, with $u_n \notin S_{n-2}$. Thus, as we already have $Rec(u_n) = (m_0, \emptyset)$, according to the protocol, u_{n-1} delivers m_0 .

If E_1 is the first event to occur, by a perfectly symmetric reasoning, we show that u_2 delivers m_0 . Thus, the result. \square

Lemma 8. *Let us suppose that $D \geq Z$. Let P be a correct polygon, and let p_1 and p_2 be two neighbor nodes of P that deliver m_0 . Then, all the nodes of P deliver m_0 .*

Proof. Let $z \leq Z$ be the number of nodes of P . Let us prove the following property \mathcal{P}_i by induction, for $1 \leq i \leq z - 1$: there exists a path of $i + 1$ nodes of P that deliver m_0 .

- \mathcal{P}_1 is true, as (p_1, p_2) is a path of 2 nodes that deliver m_0 .
- Let us suppose that \mathcal{P}_i is true for $i < z - 1$. Let (u_1, \dots, u_{i+1}) be a path of $i + 1$ nodes that deliver m_0 . Let $\{q_1, \dots, q_n\}$ be n nodes such that $(u_1, \dots, u_{i+1}, q_1, \dots, q_n, u_1)$ is a circular path on P . Then, $(u_{i+1}, q_1, \dots, q_n, u_1)$ is a path of correct nodes where u_{i+1} and u_1 deliver m_0 . Therefore, according to Lemma 7, at least one of the nodes q_1 and q_n deliver m_0 . Thus, at least one of the paths $(q_n, u_1, \dots, u_{i+1})$ and $(u_1, \dots, u_{i+1}, q_1)$ contains $i + 2$ nodes of P that deliver m_0 , and \mathcal{P}_{i+1} is true.

Therefore, \mathcal{P}_{z-1} is true, and the z nodes of P deliver m_0 . □

Theorem 1. *If $D > Z$, we achieve reliable broadcast.*

Proof. Let s be the source and let p be a correct node. According to Lemma 2, s belongs to a correct polygon P and p belongs to a correct polygon P' . According to Lemma 5, there exists a correct polygonal path (Q_1, \dots, Q_n) such that $Q_1 = P$ and $Q_n = P'$.

Let us prove the following property \mathcal{P}_i by induction, for $1 \leq i \leq n$: all the nodes of Q_i deliver m_0 .

- First, let us show that \mathcal{P}_1 is true. Let q be a neighbor of s on Q_1 . As Q_1 is correct, according to the protocol, q delivers m_0 . Then, according to Lemma 8, \mathcal{P}_1 is true.
- Let us suppose that \mathcal{P}_i is true, for $i < n$. Let u_1 and u_2 be the two nodes shared by Q_i and Q_{i+1} . As \mathcal{P}_i is true, u_1 and u_2 deliver m_0 . Then, according to Lemma 8, \mathcal{P}_{i+1} is true.

Thus, \mathcal{P}_n is true, and p delivers m_0 . Thus, the result. □

3.3 Bounds tightness

Let us show that the bound of Theorem 1 is tight.

Theorem 2. *If we only have $D \geq Z$, no algorithm can guarantee reliable broadcast.*

Proof. Let us suppose the opposite: there exists an algorithm guaranteeing reliable broadcast when $D \geq Z$. Let us consider the network of Figure 2.




Figure 2: Undecidable case for $D \geq Z$

In this network, $D = Z = 4$, thus $D \geq Z$ is satisfied. The 4 bold nodes (2 correct, 2 Byzantine) form a node-cut isolating the grey part of the network, which contains the source.

As there is a perfect symmetry between the 2 correct nodes and the 2 Byzantine nodes, the outer nodes can never determine m_0 with certitude, and reliable broadcast is impossible. This contradiction achieves the proof. \square

4 Time properties

In this section, we assume that the delay between two activations of a correct process has an upper bound T and a lower bound T' . We assume that, when activated, a process instantly receives and treats the messages stored in its channels.

First, we show that the good information is delivered within a linear time. Then, we show that reliable broadcast is still guaranteed with a limited memory (*exhaustion safety*).

4.1 Delivery time

Let us show that the good information is delivered within a linear time. More precisely, let us consider a correct node p located at $L \geq 1$ hops from the source. We show that p delivers m_0 within $o(LT)$ time units (Theorem 3). This is the same time complexity as a basic broadcast protocol, where any information received is immediately delivered and retransmitted.

Let P be a correct polygon containing the source s , and let P' be a correct polygon containing p . Such polygons exist, according to Lemma 2. In the following proofs, we use crude bounds in order to keep formulas simple.

Lemma 9. *There exists a polygonal path of at most $2LY$ polygons connecting P and P' .*

Proof. As p is located at L hops from the source, let (u_0, \dots, u_L) be a path, with $u_0 = s$ and $u_L = p$. Let (e_1, \dots, e_m) be the edges connected to u_L , clockwise, such that $e_m = e_1$ is connected to u_{L-1} . Let k be such that e_k and e_{k+1} belong to P' . Let u_{L+1} be the node connected to e_{k+1} .

Let us prove the following property \mathcal{P}_i by induction, for $0 \leq i \leq L + 1$: there exists a polygonal path $(Q_0, \dots, Q_{n(i)})$, with $n(i) \leq iY$, such that $Q_{n(i)}$ contains u_i .

- \mathcal{P}_0 is true if we take $Q_0 = P$.

- Let us suppose that \mathcal{P}_i is true, with $i \leq L$. Let (e_1^i, \dots, e_m^i) be the edges connected to u_i , clockwise, such that e_1^i and e_2^i belong to Q_i . Let $k(i)$ be such that the edge $e_{k(i)}$ is connected to u_{i+1} . For each $j < k(i)$, let $Q_{n(i)+j}$ be the polygon containing e_j^i and e_{j+1}^i . Then, $(Q_{n(i)}, \dots, Q_{n(i)+k(i)-1})$ is a polygonal path of $k(i)$ polygons, such that $Q_{n(i)+k(i)-1}$ contains u_{i+1} . Thus, \mathcal{P}_{i+1} is true if we take $n(i+1) = n(i) + k(i) - 1 \leq (i+1)Y$.

Therefore, \mathcal{P}_{L+1} is true, and $(Q_0, \dots, Q_{n(L+1)})$ is a polygonal path, with $n(L+1) \leq (L+1)Y \leq 2LY$. By definition of u_{L+1} , $Q_{n(L+1)} = P'$. Thus, the result. \square

Lemma 10. *There exists a correct polygonal path of at most Y^3ZL polygons connecting P and P' .*

Proof. According to Lemma 9, there exists a polygonal path (P_1, \dots, P_n) with $P_1 = P$, $P_n = P'$ and $n \leq 2LY$. If this path is correct, the result is trivial. Otherwise, let (P_N, \dots, P_M) be a sequence of Byzantine nodes, as defined in Lemma 5.

Let us consider the proof of Lemma 3. The circular path (q_1, \dots, q_m) contains at most YZ nodes, and each set S_i contains at most Y polygons. Thus, the set $S = S_1 \cup \dots \cup S_m$ contains at most Y^2Z polygons.

Therefore, according to the proof of Lemma 5, (P_N, \dots, P_M) can be replaced by a sequence of at most Y^2Z polygons. As the number of Byzantine sequences in (P_1, \dots, P_n) is strictly inferior to $n/2$, the correct path thus obtained contains at most $Y^2Zn/2 \leq Y^2Z(2LY)/2 \leq Y^3ZL$ polygons. \square

Theorem 3. *If s multicasts m_0 at t_0 , p delivers m_0 before $t_0 + o(LT)$.*

Proof. Let Q be a correct polygon, and let us suppose that two nodes of Q deliver m_0 before t_1 . Then, according to the proof of Lemma 8, a third node delivers m_0 before $t_1 + ZT$, and so forth. Thus, all the nodes of Q deliver m_0 before $t_1 + Z^2T$. Similarly, all the nodes of P deliver m_0 before $t_0 + Z^2T$.

According to Lemma 10, there exists a correct polygonal path of Y^3ZL polygons connecting P and P' . Therefore, according to the proof of Theorem 1, p delivers m_0 before $t_0 + Y^3Z^3LT = t_0 + o(LT)$. \square

4.2 Exhaustion safety

In the existing protocols [10, 2, 22, 16, 18, 17], the correct nodes can store as many messages as necessary. Therefore, if the memory is finite, the Byzantine nodes can adopt a very simple strategy: they can send a lot of false messages to their neighbors, very quickly, and overflow their memory before they deliver the good information. The problem of resisting to such attacks is referred to as *exhaustion safety* [24, 25].

Let us consider the memory of both nodes and channels.

- In our protocol, the memory of correct nodes is limited to one variable Rec per neighbor. More precisely, let M be the maximal number of bits of an information m , and let X be the maximal number of bits of a node identifier. As the largest tuple (m, S) that a correct node can accept verifies $card(S) \leq Z - 3$, each variable Rec requires at most $M + ZX$ bits. Thus, each correct node requires at most $Y(M + ZX)$ bits of memory.
- Let N be the smallest integer such that $N > T/T'$. Then, as a node receives all the messages of its channels when activated, a channel connecting two correct nodes contains at most N tuples (m, S) . If a channel is connected to a Byzantine node, it can be overflowed without

consequences: it is unimportant that a Byzantine node receives messages, and the messages received from a Byzantine node are already unpredictable. Thus, the channels require at most $N(M + XZ)$ bits of memory.

Therefore, our protocol still works with a finite quantity of memory in both nodes and channels.

5 Conclusion

In this paper, we have generalized the condition on the distance between Byzantine nodes to a class of planar graphs, and shown its tightness. Our solution has the same time complexity as a basic broadcast, and is exhaustion safe with reasonable time assumptions.

An open problem is, of course, to extend this condition to more general graphs. Also, even if we already have a linear time complexity, some optimizations could be made to reduce the time to deliver the good information.

References

- [1] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. McGraw-Hill Publishing Company, New York, May 1998. 6.
- [2] Vartika Bhandari and Nitin H. Vaidya. On reliable broadcast in a radio network. In Marcos Kawazoe Aguilera and James Aspnes, editors, *PODC*, pages 138–147. ACM, 2005.
- [3] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *OSDI*, pages 173–186, 1999.
- [4] D. Dolev. The Byzantine generals strike again. *Journal of Algorithms*, 3(1):14–30, 1982.
- [5] Vadim Drabkin, Roy Friedman, and Marc Segal. Efficient byzantine broadcast in wireless ad-hoc networks. In *DSN*, pages 160–169. IEEE Computer Society, 2005.
- [6] Swan Dubois, Toshimitsu Masuzawa, and Sébastien Tixeuil. The impact of topology on byzantine containment in stabilization. In *Proceedings of DISC 2010*, Lecture Notes in Computer Science, Boston, Massachusetts, USA, September 2010. Springer Berlin / Heidelberg.
- [7] Swan Dubois, Toshimitsu Masuzawa, and Sébastien Tixeuil. On byzantine containment properties of the min+1 protocol. In *Proceedings of SSS 2010*, Lecture Notes in Computer Science, New York, NY, USA, September 2010. Springer Berlin / Heidelberg.
- [8] Swan Dubois, Toshimitsu Masuzawa, and Sébastien Tixeuil. Bounding the impact of unbounded attacks in stabilization. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2011.
- [9] Swan Dubois, Toshimitsu Masuzawa, and Sébastien Tixeuil. Maximum metric spanning tree made byzantine tolerant. In David Peleg, editor, *Proceedings of DISC 2011*, Lecture Notes in Computer Science (LNCS), Rome, Italy, September 2011. Springer Berlin / Heidelberg.
- [10] Chiu-Yuen Koo. Broadcast in radio networks tolerating byzantine adversarial behavior. In Soma Chaudhuri and Shay Kutten, editors, *PODC*, pages 275–282. ACM, 2004.
- [11] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [12] D. Malkhi, Y. Mansour, and M.K. Reiter. Diffusion without false rumors: on propagating updates in a Byzantine environment. *Theoretical Computer Science*, 299(1–3):289–306, April 2003.

- [13] D. Malkhi, M. Reiter, O. Rodeh, and Y. Sella. Efficient update diffusion in byzantine environments. In *The 20th IEEE Symposium on Reliable Distributed Systems (SRDS '01)*, pages 90–98, Washington - Brussels - Tokyo, October 2001. IEEE.
- [14] Toshimitsu Masuzawa and Sébastien Tixeuil. Bounding the impact of unbounded attacks in stabilization. In Ajoy Kumar Datta and Maria Gradinariu, editors, *SSS*, volume 4280 of *Lecture Notes in Computer Science*, pages 440–453. Springer, 2006.
- [15] Toshimitsu Masuzawa and Sébastien Tixeuil. Stabilizing link-coloration of arbitrary networks with unbounded byzantine faults. *International Journal of Principles and Applications of Information Science and Technology (PAIST)*, 1(1):1–13, December 2007.
- [16] Alexandre Maurer and Sébastien Tixeuil. Limiting byzantine influence in multihop asynchronous networks. *IEEE International Conference on Distributed Computing Systems (ICDCS 2012)*.
- [17] Alexandre Maurer and Sébastien Tixeuil. On byzantine broadcast in loosely connected networks. *International Symposium on Distributed Computing (DISC 2012)*.
- [18] Alexandre Maurer and Sébastien Tixeuil. A scalable byzantine grid. *International Conference on Distributed Computing and Networking (ICDCN 2013)*.
- [19] Y. Minsky and F.B. Schneider. Tolerating malicious gossip. *Distributed Computing*, 16(1):49–68, 2003.
- [20] Mikhail Nesterenko and Anish Arora. Tolerance to unbounded byzantine faults. In *21st Symposium on Reliable Distributed Systems (SRDS 2002)*, pages 22–29. IEEE Computer Society, 2002.
- [21] Mikhail Nesterenko and Sébastien Tixeuil. Discovering network topology in the presence of byzantine nodes. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 20(12):1777–1789, December 2009.
- [22] Andrzej Pelc and David Peleg. Broadcasting with locally bounded byzantine faults. *Inf. Process. Lett.*, 93(3):109–115, 2005.
- [23] Yusuke Sakurai, Fukuhito Ooshita, and Toshimitsu Masuzawa. A self-stabilizing link-coloring protocol resilient to byzantine faults in tree networks. In *Principles of Distributed Systems, 8th International Conference, OPODIS 2004*, volume 3544 of *Lecture Notes in Computer Science*, pages 283–298. Springer, 2005.
- [24] P. Sousa, N. F. Neves, and P. Veríssimo. How resilient are distributed f fault/intrusion-tolerant systems? In *Proceedings of the 2005 International Conference on Dependable Systems and Networks (DSN'05)*, pages 98–107, June 2005.
- [25] Paulo Sousa, Nuno Ferreira Neves, Paulo Veríssimo, and William H. Sanders. Proactive resilience revisited: The delicate balance between resisting intrusions and remaining available. In *Proceedings of the 25th IEEE Symposium on Reliable Distributed Systems (SRDS)*, pages 71–80, October 2006.