

On Byzantine Broadcast in Planar Graphs

Alexandre Maurer¹ and Sébastien Tixeuil^{1,2}

¹ UPMC Sorbonne Universités, Paris, France

² Institut Universitaire de France

E-mail: Alexandre.Maurer@lip6.fr, Sebastien.Tixeuil@lip6.fr

Phone: +33 1 44 27 87 62, +33 1 44 27 87 75

December 7, 2013

Abstract

We consider the problem of reliably broadcasting information in a multihop asynchronous network in the presence of Byzantine failures: some nodes may exhibit unpredictable malicious behavior. We focus on completely decentralized solutions. Few Byzantine-robust algorithms exist for loosely connected networks. A recent solution guarantees reliable broadcast on a torus when $D > 4$, D being the minimal distance between two Byzantine nodes.

In this paper, we generalize this result to 4-connected planar graphs. We show that reliable broadcast can be guaranteed when $D > Z$, Z being the maximal number of edges per polygon. We also show that this bound on D is a lower bound for this class of graphs. Our solution has the same time complexity as a simple broadcast. This is also the first solution where the memory required increases linearly (instead of exponentially) with the size of transmitted information.

Important disclaimer These results have NOT yet been published in an international conference or journal. This is just a technical report presenting intermediary and incomplete results. A generalized version of these results may be under submission.

1 Introduction

As modern networks grow larger, they become more likely to fail, as nodes may be subject to crashes, attacks, transient bit flips, etc. To encompass all possible cases, we consider the most general model of failure: the *Byzantine* model [11], where the failing nodes can exhibit arbitrary malicious behavior. In other words, tolerating Byzantine nodes implies guaranteeing they are not able to cause problems in the correct part of the network.

In this paper, we study the problem of reliably broadcasting information in a multihop network. In the ideal case, the source node sends the information to its neighbors, that in turn send it to their own neighbors, and so forth (this is denoted in the sequel as a “simple broadcast”). However, a single Byzantine node can forward a false information and lie to the entire network. Our goal is to design a solution that guarantees reliable broadcast in the presence of Byzantine retransmitters.

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 cannot be considered as *entirely* arbitrary. Besides, manipulating asymmetric cryptography requires important resources, which may not always be available. The most important point is that cryptography requires some

degree of trusted infrastructure to initially distributes public and private keys: therefore, if this initial infrastructure fails, the whole network fails. Yet, we want to design a totally decentralized solution, where any element can fail independently without compromising the whole system. For these reasons, we focus on non-cryptographic 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. Space local algorithms try to contain the fault as close to its source as possible. This is only applicable to the problems where the information from distant 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 $3k + 1$. However, this solution assumes that the topology is known to every node, and that the network is synchronous. 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 failure detection.

Another existing approach is based on the fraction of Byzantine neighbors per node. Solutions have been proposed for nodes organized on a lattice [10, 2]. Reliable broadcast was shown possible 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 loosely connected networks, where each node has a limited (possibly upper bounded by a constant) number of neighbors. For instance, on a torus topology (see Figure 1), no existing solution can tolerate more than one Byzantine node. Efficient solutions have been proposed for such networks [16, 18], but only give probabilistic guarantees, and require the nodes to know their position in the network. This last requirement was relaxed in [17]: reliable broadcast is guaranteed on a torus when $D > 4$, D being the minimal number of hops between two Byzantine nodes.




Figure 1: Torus (left) and 4-connected planar graph (right)

Our contribution. In this paper, we generalize the result of [17] to 4-connected planar graphs (see Figure 1). We show that reliable broadcast can be guaranteed when $D > Z$, Z being the maximal number of edges per polygon. We also show that this bound is tight: if we only have $D \geq Z$, no algorithm can guarantee reliable broadcast for this class of graphs.

Then, if we assume that the delay between two activations of a same process is bounded, we show that reliable broadcast can be achieved in $O(d)$ time units, d being the network diameter. So, tolerating Byzantine failures yields the same time complexity as a simple broadcast.

Finally, we show that, unlike previous solutions [10, 2, 22, 16, 18, 17], the local memory required for broadcasting is $O(M)$ (instead of $O(2^M)$), M being the maximal size of an information message.

Organization of the paper In Section 2, we present the hypotheses and describe the broadcast protocol. In Section 3, we prove the condition for reliable broadcast, and show its tightness. In Section 4, we establish the time complexity. Finally, in Section 5, we discuss about the memory requirements.

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 bi-dimensional representation of this graph where edges do not cross. Besides, we assume that the graph is 4-connected: to disconnect the graph, at least 4 nodes must be removed (see Definition 2). From this hypothesis, each node connects at least 4 edges.

As the graph is planar, the edges delimit *polygons* (see Figure 1 and Definition 3). Let $Z \geq 3$ be the maximal number of edges per polygon, and let $Y \geq 4$ be the maximal number of edges per node. Z is a parameter of the algorithm.

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 actual 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 underlying principle as in [17]: to actually deliver an information message, a node must receive it from a direct neighbor q , but also (indirectly) 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.

Besides generalizing the aforementioned protocol to planar graphs, our new protocol improves memory efficiency. Indeed, instead of storing all received messages in a set Rec , a correct node uses a variable $Rec(q)$ for each neighbor q , storing only the last message received from q . This modification enables to reduce the memory required by the nodes (see Section 5).

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 $card(S) \leq Z - 3$: assign the value (m, S) to $Rec(q)$ and multicast $(m, S \cup \{q\})$.
 - When there exists m, p, q and S such that $q \neq p$, $q \notin S$, $Rec(q) = (m, \emptyset)$ and $Rec(p) = (m, S)$: deliver m , multicast (m, \emptyset) and stop.

3 Condition for reliable broadcast

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

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 and u_n are also neighbors. Unless we mention it, we do not require that these nodes are distinct.*

Definition 2 (Node-cut and k -connected network). *As set S of nodes is a node-cut if the graph $G - S$ is disconnected, that is: there exists a pair of nodes $\{p, q\} \notin S$ such that no path connects p and q in $G - S$. The network is k -connected if no node-cut contains less than k nodes.*

Definition 3 (Polygon). *A polygon is a circular path that does not surround any node in the bidimensionnal representation of the planar graph.*

Definition 4 (Neighbor and adjacent polygons). *Two polygons are neighbors if they share at least one node, and adjacent if they share an edge.*

Definition 5 (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 6 (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 7 (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 exist 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) be a circular path on P , and let (u, q_1, \dots, q_m) be a circular path on Q . Therefore, $(u, p_1, \dots, p_n, u, q_1, \dots, q_m)$ 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 $D > Z$. Hence, the result. \square

Lemma 2. *Let v be a node, and let V be the set of polygons containing v . Then, v is the only node common to these polygons.*

Proof. Let us suppose the opposite: there exists a node $w \neq v$ common to these polygons. Let P be a polygon containing v . Let q_1 and q_2 be the two neighbors of v contained by P . Let Q_1 (resp. Q_2) be the polygon adjacent to P containing v and q_1 (resp. q_2). Let S be the set of nodes contained by P . As a polygon contains at least 3 nodes, $S - \{v, w\}$ contains at least one node. Then, as w is also common to P , Q_1 and Q_2 , $\{v, w\}$ is a node-cut isolating $S - \{v, w\}$ from the rest of the network. This is impossible, as the network is 4-connected. Hence, the result. \square

Lemma 3. *If $D > Z$, 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 V be the set of polygons containing v . Let P_1 and P_2 be two polygons of V . As P_1 and P_2 are Byzantine, according to Lemma 1, they share the same Byzantine node b . Therefore, by induction, all the polygons of V share the same Byzantine node b . But according to Lemma 2, v is the only node shared by the polygons of V . Therefore, $b = v$, and v is Byzantine: contradiction. Hence, the result. \square

Lemma 4. *Let v be a node, and let V be the set of polygons containing v . Let X be the set of nodes contained by the polygons of V . Then, there exists a circular path (q_1, \dots, q_m) such that nodes $\{q_1, \dots, q_m\}$ are distinct and that contains all nodes of $X - \{v\}$, and only contains nodes of X .*

Proof. Let (e_1, \dots, e_n) be the edges connected to v , ordered clockwise, and let $e_{n+1} = e_1$. Let u_i be the node connected to v by e_i . If, $\forall i \in \{1, \dots, n\}$, there exists a polygon containing the edges e_i and e_{i+1} , go to paragraph 1. Else, go to paragraph 2.

1. Let P_i be the polygon containing the edges e_i and e_{i+1} . Let $(v, u_i, p_1^i, p_2^i, \dots, u_{i+1})$ be a circular path on P_i , ordered clockwise. We define a path $(u_1, p_1^1, p_2^1, \dots, u_2, p_1^2, p_2^2, \dots, u_{n+1}) = (q_1, \dots, q_{m+1})$, containing all the nodes of $X - \{v\}$. Let us show that the nodes $\{q_1, \dots, q_m\}$ are distinct. Let us suppose the opposite: there exists k and $k' > k$ such that $u_k = u_{k'}$. Then, $\{u_k, v\}$ is a node-cut disconnecting $\{u_{k+1}, \dots, u_{k'-1}\}$ from the rest of the network, which is impossible as the network is 4-connected. Thus, the nodes are distinct. Hence, the result.

2. Let k be the first integer such that e_k and e_{k+1} do not belong to any polygon. Let us notice that there is no other integer $k' > k$ satisfying this property – otherwise, $\{v\}$ would be a node-cut isolating u_k from $u_{k'}$. Then, let (e'_1, \dots, e'_n) be the edges connected to v , clockwise, such that $e'_1 = e_{k+1}$.

Let P_i be the polygon containing the edges e'_i and e'_{i+1} . Let $(v, u_i, p_1^i, p_2^i, \dots, u_{i+1})$ be a circular path on P_i , ordered clockwise. We define a path $(u_1, p_1^1, p_2^1, \dots, u_2, p_1^2, p_2^2, \dots, u_n) = (q_1, \dots, q_{m-1})$, containing all nodes of $X - \{v\}$. For the same reasons as in paragraph 1, the nodes $\{q_1, \dots, q_{m-1}\}$ are distinct. Hence, the result, if we take $q_m = v$.

□

Lemma 5. *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 (q_1, \dots, q_m) be the circular path of Lemma 4. Then, $S = S_1 \cup \dots \cup S_m$, where S_i is the set of polygons containing q_i . If each set S_i is connected, as S_i and S_{i+1} share a polygon containing q_i and q_{i+1} , S is connected. Now, let us suppose that there exists a k such that S_k is not connected.

S_k contains only two disconnected parts, otherwise $\{v\}$ would be a node-cut. Let (q'_1, \dots, q'_m) be a circular path containing nodes $\{q_1, \dots, q_m\}$, ordered clockwise, such that $q'_1 = q_k$. Let S'_1 (resp. S'_{m+1}) be the part of S_k containing the node q_2 (resp. q_m). $\forall i \in \{2, \dots, m\}$, let S'_i be the set of polygons containing q'_i . Then, $S = S_1 \cup \dots \cup S_m = S'_1 \cup \dots \cup S'_{m+1}$. Let us prove the following property \mathcal{P}_i by induction, $\forall i \in \{1, \dots, m+1\}$: $S'_1 \cup \dots \cup S'_i$ is connected.

- \mathcal{P}_1 is true, as S'_1 is connected.
- Let us suppose that \mathcal{P}_i is true, for $i \in \{1, \dots, m\}$. Let us suppose that $S'_1 \cup \dots \cup S'_{i+1}$ is not connected. It implies that S'_{i+1} is not connected. S'_{i+1} contains only two disconnected parts, otherwise $\{q'_{i+1}\}$ would be a node-cut. Let S'^A_{i+1} be the part containing the node q'_i , and let S'^B_{i+1} be the other part. Then, $\{q'_1, v, q'_{i+1}\}$ is a node-cut isolating $S'_1 \cup \dots \cup S'_i \cup S'^A_{i+1}$ from S'^B_{i+1} , which is impossible as the network is 4-connected. Thus, \mathcal{P}_{i+1} is true.

Therefore, \mathcal{P}_{m+1} is true, and S is connected. □

Lemma 6. *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 neighbors to a polygon of V . As V contains P_1 and P_n , by definition, S contains P and Q . According to Lemma 5, 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: there exists a polygon P' of S that 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. Hence, the result. □

Lemma 7. *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 6, 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 8. *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. Hence, the result. \square

Lemma 9. *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 8, 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 $\text{card}(S_i) \leq Z - 3$, 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 $\text{card}(S_{n-2}) \leq n - 3 \leq Z - 3$. Thus, for the node u_{n-1} , $\text{Rec}(u_{n-2}) = (m_0, S_{n-2})$, with $u_n \notin S_{n-2}$. Thus, as we already have $\text{Rec}(u_n) = (m_0, \emptyset)$, 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 . Hence, the result. \square

Lemma 10. *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 9, 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 . \square

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 3, s belongs to a correct polygon P and p belongs to a correct polygon P' . According to Lemma 7, 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 10, \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 10, \mathcal{P}_{i+1} is true.

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

3.3 Bounds tightness

Let us show that the bound on D (Theorem 1) cannot be improved.

Theorem 2. *If $D \geq Z$, no algorithm can guarantee reliable broadcast on 4-connected planar graphs.*

Proof. Let us suppose the opposite: there exists an algorithm guaranteeing reliable broadcast on 4-connected planar graphs for $D \geq Z$. Let us consider the network of Figure 2.




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

In this network, $D = Z = 4$, thus $D \geq Z$ is satisfied. Here, we have 4 nodes (2 correct, 2 Byzantine) forming a node-cut that isolates 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

Nevertheless, notice that it does not make the condition $D > Z$ necessary for all graphs: the necessary and sufficient condition to achieve byzantine resilient broadcast may be more complex than the distance between Byzantine failures. We leave this as an open question.

4 Time complexity

In this section, we assume that the delay between two activations of the same process has an upper bound T . Then, we show that reliable broadcast is achieved in $O(d)$ time units, d being the diameter of the network. This is the same time complexity as a simple broadcast, where any information received is retransmitted without verification.

Lemma 11. *Let p be a node located a $L \geq 1$ hops from the source. Then, there exists a correct polygonal path of at most $Y^3 Z L$ polygons connecting p to the source.*

Proof. 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 3. Let (u_1, \dots, u_{L+1}) be a path connecting s and p , and let U_i be the set of polygons containing u_i . Each set U_i is connected, otherwise $\{u_i\}$ would be a node-cut. Therefore, $U = U_1 \cup \dots \cup U_{L+1}$ is connected. As each set U_i contains at most Y polygons, U contains at most $Y(L+1)$ polygons.

Therefore, there exists a polygonal path (P_1, \dots, P_n) of at most $Y(L+1)$ polygons, with $P_1 = P$ and $P_n = P'$. 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 7.

Let us consider the proof of Lemma 5. 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 $Y^2 Z$ polygons.

Therefore, according to the proof of Lemma 7, (P_N, \dots, P_M) can be replaced by a sequence of at most $Y^2 Z$ 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 $Y^2Zn/2 \leq Y^2ZY(L+1)/2 \leq Y^3ZL$ polygons. \square

Theorem 3. *Reliable broadcast is achieved in $O(d)$ time units.*

Proof. Let us suppose that the source broadcasts m_0 at a date t_0 .

Let Q be a correct polygon, and let us suppose that two nodes of Q have delivered m_0 at a date t . Then, according to the proof of Lemma 10, a third node delivers m_0 before $t + ZT$, and so forth. Thus, all nodes of Q deliver m_0 before $t + Z^2T$. Similarly, all nodes of P deliver m_0 before $t_0 + Z^2T$.

According to Lemma 11, for any node p located at $L \geq 1$ hops from the source, there exists a correct polygonal path of Y^3ZL polygons connecting this node to the source. Thus, according to the proof of Theorem 1, p delivers m_0 before $t_0 + Y^3Z^3LT$.

Therefore, as $L \leq d$, reliable broadcast is achieved in Y^3Z^3Td time units. Thus, as Y , Z and T are bounded, reliable broadcast is achieved in a $O(d)$ time. \square

5 Required memory

In this section, we show that our solution is the first Byzantine resilient broadcast in sparse multi-hop networks where the used memory increases linearly with the size of informations, and not exponentially.

Indeed, the existing solutions [10, 2, 21, 16, 17, 18], the nodes are supposed to store as many information messages m as necessary. However, the Byzantine nodes can potentially broadcast all possible false informations $m' \neq m_0$. This strategy is referred to as *exhaustion* in the literature [24, 25]. Therefore, the correct nodes implicitly require $O(2^M)$ bits of memory to ensure reliable broadcast, M being the maximal number of bits of an information m .

In our protocol, we made the following modification : instead of storing all the messages received, we only store the last message received from a neighbor q in the variable $Rec(q)$. Thus, the nodes only require $O(M)$ bits of memory. More precisely, let us consider a finite network, 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 $card(S) \leq Z$, each variable Rec requires at most $M + ZX$ bits. Thus, each correct node requires at most $Y(M + ZX)$ bits of memory.

Concerning the memory required in channels, the problem is the same for all solutions: we must assume that the delay between two activations of a same process belongs to an interval $[T_1, T_2]$, $T_1 > 0$ – otherwise, the memory is impossible to bound. Indeed, let N be the smallest integer such that $N > T_2/T_1$. 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) . Besides, 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, each channel requires at most $N(M + XZ)$ bits of memory.

Therefore, the local memory required is now $O(M)$ instead of $O(2^M)$.

6 Conclusion

We 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 requires less memory than the previous solutions.

An open problem is to find more involved criteria for the placement of Byzantine failures, and to extend it to more general graphs. Also, even if we already have a linear time complexity, some optimizations could be made to further reduce the time to deliver genuine 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.