



HAL
open science

Reliable Broadcast in Dynamic Networks with Locally Bounded Byzantine Failures

Silvia Bonomi, Giovanni Farina, Sébastien Tixeuil

► **To cite this version:**

Silvia Bonomi, Giovanni Farina, Sébastien Tixeuil. Reliable Broadcast in Dynamic Networks with Locally Bounded Byzantine Failures. [Technical Report] Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6, F-75005 Paris, France; Dipartimento di Ingegneria Informatica Automatica e Gestionale "Antonio Ruberti", Università degli Studi di Roma La Sapienza, Rome, Italy. 2018. hal-01712277v3

HAL Id: hal-01712277

<https://hal.science/hal-01712277v3>

Submitted on 31 Oct 2018

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.

Reliable Broadcast in Dynamic Networks with Locally Bounded Byzantine Failures [★]

Silvia Bonomi¹, Giovanni Farina^{2,1} ✉, and Sébastien Tixeuil²

¹ Dipartimento di Ingegneria Informatica Automatica e Gestionale Antonio Ruberti,
Sapienza Università di Roma, Rome, Italy

`bonomi@diag.uniroma1.it`

² Sorbonne Université,

CNRS, Laboratoire d'Informatique de Paris 6, LIP6, F-75005 Paris, France
{`giovanni.farina,sebastien.tixeuil`}@lip6.fr

Abstract. Ensuring reliable communication despite possibly malicious participants is a primary objective in any distributed system or network. In this paper, we investigate the possibility of reliable broadcast in a dynamic network whose topology may evolve while the broadcast is in progress. In particular, we adapt the Certified Propagation Algorithm (CPA) to make it work on dynamic networks and we present conditions (on the underlying dynamic graph) to enable safety and liveness properties of the reliable broadcast. We furthermore explore the complexity of assessing these conditions for various classes of dynamic networks.

Keywords: Byzantine Reliable Broadcast · Locally bounded failures · Dynamic Networks.

1 Introduction

Designing dependable and secure systems and networks that are able to cope with various types of adversaries, ranging from simple errors to internal or external attackers, requires to integrate those risks from the very early design stages. The most general attack model in a distributed setting is the Byzantine model, where a subset of nodes participating in the system may behave arbitrarily (including in a malicious manner), while the rest of processes remain correct. Also, reliable communication primitives are a core building block of any distributed software. Finally, as current applications are run for extended periods of time with expected high availability, it becomes mandatory to integrate dynamic changes in the underlying network while the application is running. In this paper, we address the reliable *broadcast* problem (where a *source* node must

[★] This work was performed within Project ESTATE (Ref. ANR-16-CE25-0009-03), supported by French state funds managed by the ANR (Agence Nationale de la Recherche), and it has been partially supported by the INOCS Sapienza Ateneo 2017 Project (protocol number RM11715C816CE4CB). Giovanni Farina thanks the *Université Franco-Italienne/Università Italo-Francese* (UFI/UIF) for supporting his mobility through the Vinci grant 2018.

send data to every other node) in the context of dynamic networks (whose topology may change while the broadcast is in progress) that are subject to Byzantine failures (a subset of the nodes may act arbitrarily). The reliable broadcast primitive is expected to provide two guarantees: *(i) safety*, namely if a message m is delivered by a correct process, then m was sent by the source and *(ii) liveness*, namely if a message m is sent by the source, it is eventually delivered by every correct process.

Related Works. In static multi-hop networks (in which the topology remains fixed during the entire execution of the protocol) the necessary and sufficient condition enabling reliable broadcast while the maximum number of Byzantine failure is bounded by f has been identified by Dolev [5], stating that this problem can be solved if and only if the network is $2f + 1$ -connected. Subsequently, the reliable broadcast problem has been analyzed assuming a local condition on the number of Byzantine neighbors a node may have [10, 16]. All aforementioned works require high network connectivity. Indeed, extending a reliable broadcast service to sparse networks required to weaken the achieved guarantees [12–14]: *(i)* accepting that a small minority of correct nodes may accept invalid messages (thus compromising safety), or accepting that a small minority of correct nodes may not deliver genuine messages (thus compromising liveness).

Adapting to dynamic networks proved difficult, as the topology assumptions made by the mentioned proposals may no longer hold: the network changes during the execution. Some core problems of distributed computing have been considered in the context of dynamic networks subject to Byzantine failures [1, 8] but, to the best of our knowledge, there exists a single contribution for the reliable communication problem, due to Maurer *et al.* [15]. Their work can be seen as the dynamic network extension of the Dolev [5] solution for static networks, and assumes that no more than f Byzantine processes are present in the network. Also, the protocol to be executed spreads an exponential number of messages with respect to the size of the network and requires each node to compute the minimal cut over the set of paths traversed by each received message, making the protocol unpractical for real applications.

The Byzantine tolerant reliable broadcast can also be solved by employing cryptography (e.g., digital signatures) [4, 6] that enable all nodes to exchange messages guaranteeing authentication and integrity. The main advantage of cryptographic protocols is that they allow solving the problem with simpler solutions and weaker conditions (in terms of connectivity requirements). However, on the negative side, the safety of the protocols is bounded to the crypto-system.

Contributions. In this paper, we investigate the possibility of reliable broadcast in a dynamic network that is subject to Byzantine faults. More precisely, we address the possibility of a local criterion on the number of Byzantine (as opposed to a global criterion as in Maurer *et al.* [15]) in the hope that a practically efficient protocol can be derived in case the criterion is satisfied. Our starting point is the CPA protocol [2, 10, 16, 18], that was originally designed for static networks. In particular, our contributions can be summarized as follows: *(i)* we extend the CPA algorithm to make it work in dynamic networks; *(ii)* we prove that the

original safety property of CPA naturally extends to dynamic networks and we define new liveness conditions specifically suited for the dynamic networks and (iii) we investigate the impact of nodes awareness about the dynamic network on reliable broadcast possibility and efficiency.

2 System Model & Problem Statement

We consider a distributed system composed by a set of n processes $\Pi = \{p_1, p_2, \dots, p_n\}$, each one having a unique integer identifier. The passage of time is measured according to a fictional global clock spanning over natural numbers \mathbb{N} . The processes are arranged in a multi-hop communication network. The network can be seen as an undirected graph where each node represents a process $p_i \in \Pi$ and each edge represents a communication channel between two elements $p_i, p_j \in \Pi$ such that p_i and p_j can communicate.

Dynamic Network Model. The communication network is *dynamic* i.e., the set of edges (or available communication channels) changes over time. More formally, we model the network as a *Time Varying Graph* (TVG) [3] i.e., a graph $\mathcal{G} = (V, E, \rho, \zeta)$ where:

- V is the set of processes (in our case $V = \Pi$);
- $E \subseteq V \times V$ is the set of edges (i.e., communication channels).
- $\rho : E \times \mathbb{N} \rightarrow \{0, 1\}$ is the *presence* function. Given an edge $e_{i,j}$ between two nodes p_i and p_j , $\rho(e_{i,j}, t) = 1$ indicates that edge $e_{i,j}$ is present at time t ;
- $\zeta : E \times \mathbb{N} \rightarrow \mathbb{N}$ is the *latency* function that indicates how much time is needed to cross an edge starting from a given time t . In particular, $\zeta(e_{i,j}, t) = \delta_{i,j,t}$ indicates that a message m sent at time t from p_i to p_j takes $\delta_{i,j,t}$ time units to cross edge $e_{i,j}$.

The evolution of \mathcal{G} can also be described as a sequence of static graphs $\mathcal{S}_{\mathcal{G}} = G_0, G_1, \dots, G_T$ where G_i corresponds to the *snapshot* of \mathcal{G} at time t_i (i.e. $G_i = (V, E_i)$ where $E_i = \{e \in E \mid \rho(e, t_i) = 1\}$). No further assumptions on the evolution of the dynamic network are made. The static graph $G = (V, E)$ that considers all the processes and all the possible existing edges is called *underlying graph* of \mathcal{G} and it flattens the time dimension indicating only the pairs of nodes that have been connected at some time t' . In the following, we interchangeably use terms *process* and *node* and we will refer to *edges* and *communication channels* interchangeably. Let us note that the TVG model is one among the most general available and it is able to abstract and characterize several real dynamic networks [3].

Communication model and Timing assumption. Processes communicate through message exchanges. Every message has (i) a *source*, which is the id of the process that has created the message and (ii) a *sender*, that is the id of the process that is relaying the message. The source and the sender may coincide. The sender is always a neighbor in the communication network. The ID of the source is included inside the message, i.e. any message is composed by its content and the source ID. We refer with m_s to a message m with p_s as source.

We assume *authenticated* and *reliable* point-to-point channels where (a) *authenticated* ensures that the identity of the sender cannot be forged; (b) *reliable* guarantees that the channel delivers a message m if and only if (i) m was previously sent by its sender and (ii) the channel has been up long enough to allow the reception (i.e. given a message m sent at time t from p_i to p_j and having latency $\delta_{i,j,t}$, we will have reliable delivery if $\rho(e_{i,j}, \tau) = 1$ for each $\tau \in [t, t + \delta_{i,j,t}]$). Notice that these channel assumptions are implicitly made also on analysis of CPA on static networks and that they are both essential to guarantee the reliable broadcast properties.

At every time unit t each process takes the following actions: (i) *send* where processes send all the messages for the current time unit (potentially none), (ii) *receive* where processes receive and store all the messages for the current time unit (potentially none) and (iii) *computation* where processes process the buffer of received messages and compute the messages to be sent during the next time unit according to the deterministic distributed protocol \mathcal{P} that they are executing. Thus, the system is assumed to be synchronous in the sense that (i) every channel has a latency function that is bounded and the overall message delivery time is bounded by the maximum channel latency and (ii) computation steps are bounded by a constant that is negligible with respect to the overall message delivery time and we consider it equal to 0. We discuss the implications and consequences of lack of synchrony inside the full version paper.

Failure model. We assume an omniscient adversary able to control several processes of the network allowing them to behave arbitrarily (including corrupting/dropping messages or simply crashing). We call them *Byzantine* processes. Processes that are not Byzantine faulty are said to be *correct*. Correct processes do not a priori know which processes are Byzantine. Specifically to reliable broadcast protocols, a Byzantine process can spread messages carrying a fake source ID and/or content or it can drop any received message preventing its propagation.

We considered the *f-locally bounded* failure model [10] as all CPA related works, i.e., along time every process p_i can be connected with at most f Byzantine processes. In other words, given the underlying static graph $G = (V, E)$, every process $p_i \in V$ has at most f Byzantine neighbors in G .

Problem Statement. In this paper, we consider the problem of *Reliable Broadcast over dynamic networks* assuming a *f-locally bounded* Byzantine failure model from a given *correct*³ source p_s . We say that a protocol \mathcal{P} satisfies *reliable broadcast*, if a message m broadcast by a correct process $p_s \in \Pi$ (also called *source* or *author*) is eventually delivered (i.e., accepted as a valid message) by every correct process $p_j \in \Pi$. Said differently, a protocol \mathcal{P} satisfies *reliable broadcast*, if the following conditions are met:

- **Safety** if a message m is delivered by a correct process, then such message has been sent by the source p_s ;

³ note the assumption of a possibly faulty source leads to a more general problem, the *Byzantine Agreement* [5]

- **Liveness:** if a message m is broadcast by the source p_s , it is eventually delivered by every correct process.

In other words, a reliable broadcast protocol extends the guarantees provided by the communication channels to the message exchanges between a node and any correct process not directly connected to it.

3 The Certified Propagation Algorithm (CPA)

The Certified Propagation Algorithm (CPA) [10, 16] is a protocol enforcing reliable broadcast, from a correct source p_s , in *static* multi-hop networks with a f -locally bounded Byzantine adversary model, where nodes have no knowledge on the global network topology. Given a message m to be broadcast, CPA starts the propagation of m_s from p_s and applies three acceptance policies (denoted by AC) to decide if m_s should be accepted and forwarded (*i.e.*, transmitted also by nodes different from the source) by a process p_j . Specifically:

- p_s delivers m_s (**AC1**), forwards it to all of its neighbors, and stops;
- when receiving m_s from p_i , if p_i is the source then p_j delivers m_s (**AC2**), forwards m_s to all of its neighbors and stops; otherwise the message is buffered.
- upon receiving $f + 1$ copies of m_s from distinct neighbors, p_j delivers m_s (**AC3**), then forwards it to all its neighbors and stops.

The correctness of CPA on static networks has been proved to be dependent on the network topology. In particular, Litsas *et al.* [11] provided topological conditions based on the concept of *k-level ordering*. Informally, given a graph $G = (V, E)$ and considering a node p_s as the source, we can define a k -level ordering as a partition of nodes into *ordered levels* such that: (i) p_s belongs to level L_0 , (ii) all the neighbors of p_s belong to level L_1 , and (iii) each node in a level L_i has at least k neighbors over levels L_j , with $j < i$. A k -level ordering is *minimum* if every node appears in the minimum level possible.

Definition 1 (MKLO). Let $G = (V, E)$ be a graph and let p_s be a node of G called source. The minimum k -level ordering (MKLO) of G from p_s is the partition P_k of nodes into disjoint subsets called levels L_i defined as follows:

$$\begin{cases} p \in L_0 & \text{if } p = p_s \\ p \in L_1 & \text{if } p \in N_s \\ p \in L_{i>1} & \text{if } p \in V \setminus \left(\bigcup_{j=0}^{i-1} L_j \right) \text{ and } |N_p \cap \left(\bigcup_{j=0}^{i-1} L_j \right)| \geq k \end{cases}$$

For CPA to ensure reliable broadcast from p_s , a sufficient condition is that a k -level ordering from p_s exists, with $k \geq 2f + 1$. Conversely, the necessary condition demands a k -level ordering from p_s with $k \geq f + 1$ (see [11]). Those conditions can be verified with an algorithm whose time complexity is polynomial in the size of the network, specifically with a modified Breadth-First-Search. In the case that a graph $G = (V, E)$ satisfies the necessary condition from p_s but

not the sufficient one, then further analysis must be carried out. In particular, in order to verify whether G enables reliable broadcast from p_s , one should check whether a k -level ordering from p_s exists (with $k = f + 1$) in every sub-graph G' obtained from G by removing all nodes corresponding to possible Byzantine placement in the f -locally bounded assumption. The verification of the strict condition has been proven to be NP-Hard [9].

4 The Certified Propagation Algorithm on Dynamic Networks

In this section, we analyze how CPA behaves on dynamic networks, *i.e.* networks whose topology may evolve over time, and how it needs to be extended to work in such settings.




Fig. 1: Example of a simple TVG and its underlying static graph.

Let us consider the TVG shown in Figure 1 and suppose process p_2 is Byzantine. If we consider the static underlying graph $G = (V, E)$ shown in Figure 1b, it is easy to verify that running CPA from the source node p_s is possible to achieve reliable broadcast in a 1-locally bounded adversary. However, if we consider snapshots of the TVG at different times⁴ as shown in Figure 1a, one can verify that nodes p_3 and p_4 remain unable to deliver the message forever. In fact, p_3 is not a neighbor of the source p_s when the message is broadcast by p_s (*i.e.*, at time t_0), and even if it had happened ($e_{s,3}$ at time t_0) the edge connecting p_4 with its correct neighbor p_3 appears only before the message would have been delivered and accepted by p_3 , and thus it is not available for the retransmission. From this simple example its easy to see that the temporal dimension plays a fundamental role in the definition of topological constraints that a TVG must satisfy to enable reliable broadcast.

⁴ For the sake of simplicity, we consider the channel delay always equal to 1 in the example.

4.1 CPA Safety in Dynamic Networks

In the following, we show that the authenticated and reliable channels are necessary to ensure the reliable broadcast through CPA.

Lemma 1. *The CPA algorithm does not ensure safety of reliable broadcast when channels are not both authenticated and reliable (even on static graphs).*

Proof. An authenticated channel guarantees that the identity of the sender of a message cannot be forged. Without this assumption a Byzantine process can impersonate an arbitrary number of processes and invalidate the f -locally bounded assumption.

A reliable channel guarantees that a message is received as it was sent by its sender. Without this assumption, an unreliable channel can potentially simulate a Byzantine process (namely the channel can deliver a message different from the one that was sent by the sender). \square

The same channel assumptions are sufficient for ensuring safety also on dynamic networks.

Theorem 1. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be the TVG of a network with f -locally bounded Byzantine adversary. If every correct process p_i runs CPA on top of reliable authenticated channels, then if a message m_s is delivered by p_i , m_s was previously sent by the correct source p_s .*

Proof. The proof trivially follows from CPA correctness in static networks with f -locally bounded adversary, considering that in the underlying graph $G = (V, E)$, we still have a f -locally bounded adversary. \square

4.2 CPA Liveness in Dynamic Networks

The CPA liveness in static networks is based on the availability of a certain topology that supports the message propagation. Indeed every edge is always up so, once the communication network satisfies the topological constraints imposed by the protocol, the assumption that channels do not lose messages is sufficient to guarantee their propagation. In dynamic networks, this is no longer true. Let us recall that each edge e in a TVG is up according to its presence function $\rho(e, t)$. At the same time, the message delivery latency are determined by the edge latency function $\zeta(e, t)$. As a consequence, in order to ensure that a message m sent at time t from p_i to p_j is delivered, we need that (p_i, p_j) remains up until time $t + \zeta(e, t)$. Contrarily, there could exist a communication channel where every message sent has no guarantee to be delivered as the edge disappears while the message is still traveling. Thus, in addition to topological constraints, moving to dynamic networks we need to set up other constraints on when edges appear and for how long they remain up. Considering that processes have no information about the network evolution, they do not know if and when a given transmitted message will reach its receiver. Hence, without assuming extra knowledge, a correct process must re-send messages infinitely often.

As a consequence, CPA must be extended to the dynamic context incorporating the following additional steps:

- if process p_i delivers a message m , it forwards m to all of its neighbors infinitely often, at every time unit.

As a consequence, each time that the neighbors of p_i changes, p_i attempts to propagate the message. Let us notice that such an infinite retransmission can be avoided/stopped only if a process get the acknowledgments about the delivery of the communication channels. This issue has been analyzed by considering further assumptions on the dynamic network [7, 17]. To ease of explanation, we will refer to this extended version of CPA as Dynamic CPA (DCPA).

We now characterize the conditions enabling a communication channel to deliver messages in order to argue about liveness. For this purpose, we define a boolean predicate whose value is `true` if and only if the TVG allows the reliable delivery of a message m sent from p_i to p_j at time t .

Definition 2. Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG. We define the predicate Reliable Channel Delivery at time t' , $\text{RCD}(p_i, p_j, t')$ as follows:

$$\text{RCD}(p_i, p_j, t') = \begin{cases} \text{true} & \text{if } \rho(\langle p_i, p_j \rangle, \tau) = 1, \forall \tau \in [t', t' + \zeta(e_{i,j}, t')]. \\ \text{false} & \text{otherwise.} \end{cases}$$

The communication channels do not usually have memory, thus we consider any message sent while the $\text{RCD}()$ predicate is `false` as dropped.

Now that we are able to express constraints on each edge through the $\text{RCD}()$ predicate, we need to define those $\text{RCD}()$ that enable liveness of reliable broadcast. Let us define the *k-acceptance function*, that encapsulates temporal aspects for the three acceptance conditions of CPA.

Definition 3. Let $p_s \in \Pi$ be a process that starts a reliable broadcast at time t_{br} . The *k-acceptance function* $\mathcal{A}_k(p, t)$ over the time $t \in \mathbb{N}$ is defined as follows:

$$\mathcal{A}_k(p_j, t) = \begin{cases} 1 & \text{if } p_j = p_s \text{ with } t \geq t_{br} & (AK1) \\ 1 & \text{if } \exists t' \geq t_{br} : \text{RCD}(p_s, p_j, t') = \text{true with } t \geq t' + \zeta(e_{s,j}, t') & (AK2) \\ 1 & \text{if } \exists p_1, \dots, p_k : \forall i \in [1, k], \mathcal{A}_k(p_i, t_i) = 1 \text{ and} \\ & \quad \exists t'_i \geq t_i : \text{RCD}(p_j, p_i, t'_i) = \text{true with } t \geq t'_i + \zeta(e_{i,j}, t'_i) & (AK3) \\ 0 & \text{otherwise} \end{cases}$$

Definition 4. Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG, and let p_s be a node called source. A temporal minimum *k-level ordering* of \mathcal{G} (TMKLO) from p_s is a partition of the nodes in levels L_{t_i} defined as follows:

$$p \in L_{t_i} \text{ iff } t_i = \min t \in \mathbb{N} \text{ such that } \mathcal{A}_k(p, t_i) = 1$$

Let us denote as P_k the partition identifying the temporal minimum *k-level ordering*.




Fig. 2: TVG example.

As an example, let us consider the TVG presented in Figure 2: it evolves in five discrete time instants (*i.e.*, t_0, t_1, \dots, t_4), its latency function $\zeta(e, t)$ is equal to 1 for every edge e at any time t . Now, let us consider process p_s as a source node that broadcasts m at time $t_{br} = 0$, and let us assume that $k = 2$. Such a TVG admits a temporal minimum 2-level ordering $P_2 = \{L_{t_0} = \{p_s\}, L_{t_1} = \{p_1\}, L_{t_2} = \{p_3\}, L_{t_4} = \{p_2, p_4\}\}$. Indeed:

- The 2-acceptance function $\mathcal{A}_2(p_s, t)$ is equal to 1 for $t \geq t_{br} = t_0$ according to *AK1*.
- The acceptance function evaluated on process p_1 is equal to 1 for $t \geq 1$ according to *AK2* (*i.e.*, $t' = 0$ and $RCD(p_s, p_1, 0) = true$ due to the presence function $\rho(\langle p_s, p_1 \rangle, \tau) = 1, \forall \tau \in [0, 1]$).
- On processes p_3 and p_2 , the acceptance function evaluates to 1 respectively for $t \geq 2$ and for $t \geq 4$, for the same reasons as p_1 .
- The acceptance function on p_4 evaluates to 1 for $t \geq 4$ according to *AK3* (*i.e.*, $RCD(p_i, p_4, t'_i) = true$ for $p_i = p_1, t'_i = 1$, and for $p_i = p_3, t'_i = 3$).

We now present a sufficient condition (Theorem 2) and a necessary condition (Theorem 3) for the liveness of reliable broadcast based on the TMKLO.

Theorem 2 (DCPA liveness sufficient condition). *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG, let p_s be the source which broadcasts m at time t_{br} , and let us assume f -locally bounded Byzantine failures. If there exists a partition $P_k = \{L_{t_{br}}, L_{t_1} \dots L_{t_x}\}$ of the nodes in V representing a TMKLO of \mathcal{G} associated to m with $k > 2f$, then the message m spread using DCPA is eventually delivered by every correct process in \mathcal{G} .*

Proof. We need to prove that if there exist a TMKLO with $k > 2f$ associated to message m , then any correct process eventually satisfies one of the CPA acceptance policies. A TMKLO with $k > 2f$ implies that there exist a time t such that the $2f + 1$ -acceptance function $\mathcal{A}_k(p, t)$ is equal to 1 for every node of the network.

The process p_s belongs to any TMKLO due to *AK1*: as the source of the broadcast, p_s delivers the message according to *AC1*. Remind that the correct processes running DCPA spread the delivered messages over their neighborhood infinitely often. Then, the other nodes belong to the TMKLO due to the occurrence of *AK2* or *AK3*.

If *AK2* is satisfied by a node p_j from time t_j , then m : (i) can be delivered by the channel interconnecting p_s with p_j by definition of $\text{RCD}()$, and (ii) it is transmitted by p_s , because t_j is greater than t_{br} . It follows that p_j delivers m according to *AC2*: indeed, p_j has received m directly from the source.

If *AK3* is satisfied on a node p_j , it is possible to identify two scenarios:

- **Case 1:** $\text{RCD}()$ is satisfied between p_j and $2f + 1$ nodes p_i where *AK2* is already satisfied. We have shown that the processes satisfying *AK2* accept m , and so they retransmit m . Assuming the f -locally bounded failure model, at most f nodes among the neighbors of p_i can be Byzantine and may not propagate m . Thus, p_j receives at least $f + 1$ copies of m from distinct neighbors. According to *AC3* of DCPA p_j delivers m .
- **Case 2:** $\text{RCD}()$ is satisfied between p_j and $2f + 1$ nodes p_i where *AK2* or *AK3* is already satisfied. Inductively, as the nodes considered in **Case 1** deliver m , it follows that the nodes p_j satisfying *AK3* due to at least $2f + 1$ nodes p_i where *AK2* or *AK3* already holds also deliver m .

□

Theorem 3 (DCPA liveness necessary condition). *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG, let p_s be the source that starts to broadcast m at time t_{br} , and let us assume f -locally bounded Byzantine failures. The message m can be delivered by every correct process in \mathcal{G} only if a partition $P_k = \{L_{t_{br}}, L_{t_1} \dots L_{t_x}\}$ of nodes in V representing a TMKLO of \mathcal{G} associated to m with $k > f$ exists.*

Proof. Let us assume for the purpose of contradiction that: (i) every correct process in \mathcal{G} delivers m , (ii) the Byzantine failures are f -locally bounded, and (iii) there does not exist a TMKLO associated to m with $k > f$. The latter implies that the TMKLO with $k = f + 1$ does not include all the nodes, i.e. $\exists p \in \Pi \mid \forall t \in \mathbb{N}, \mathcal{A}_{f+1}(p, t) = 0$.

The process p_s is always included in a TMKLO of any k . Thus, p_s is included in P_{f+1} . The nodes that deliver m according to *AC2* have received m from p_s . Thus, the $\text{RCD}()$ predicate evaluated between p_s and p_i was true at least once after the delivery of m by p_s . It follows that the condition defined in *AK2* is eventually satisfied, and that those nodes are included in P_{f+1} .

The remaining nodes that deliver according to *AC3* have received the message from $f + 1$ distinct neighbors. Let us initially assume that such neighbors have delivered the message by *AC2*. Again, the RCD predicate evaluated between the receiving node p_j and the distinct $f + 1$ neighbors p_i has been true at least once after the respective deliveries of m . We already proved that such neighbors of p_i are included in P_{f+1} , therefore the condition defined in *AK2* is satisfied by those p_j and they are included in P_{f+1} .

It naturally follows that the remaining nodes (the ones that have received the message from neighbors satisfying *AC2* or *AC3*) are included in P_{f+1} . This is in contradiction with the assumptions we made, because eventually every process satisfies one of the conditions *AK1*, *AK2* or *AK3*, and the claim follows. □

5 On the Detection of DCPA Liveness

In Section 4, we proved that DCPA always ensure the reliable broadcast safety, and we provided the necessary and sufficient conditions about the dynamic network to enforce the reliable broadcast liveness. In this section, we are investigating the ability of individual processes to detect whether the reliable broadcast liveness is actually achieved in the current network. In more detail, we seek answers to the following questions:

- **(Conscious Termination):** Given a message m_s sent by a source p_s on TVG \mathcal{G} , is p_s able to detect if m_s will eventually be delivered by every correct process?
- **(Bounded Broadcast Latency):** Given a message m_s sent by a source p_s on TVG \mathcal{G} , is p_s able to compute upper and lower bounds for reliable broadcast completion?

Obviously, if p_s has no knowledge about \mathcal{G} , nothing about termination can be detected. As a consequence, some knowledge about \mathcal{G} is required to enable Conscious Termination and Bounded Broadcast Latency. We now formalize the notion of *Broadcast Latency*, and introduce oracles that abstract the knowledge a process may have about \mathcal{G} .

Definition 5 (Broadcast Latency (BL)). *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG and let p_s be a node called source that broadcasts a message m at time t_{br} . We define as Broadcast Latency BL the period between t_{br} and the time of the last delivery of m by a correct process.*

We define the following knowledge oracles (from more powerful to least powerful):

- **Full knowledge Oracle (FKO):** FKO provides full knowledge about the TVG, *i.e.*, it provides $\mathcal{G} = (V, E, \rho, \zeta)$;
- **Partial knowledge Oracle (PKO):** given a TVG $\mathcal{G} = (V, E, \rho, \zeta)$, PKO provides the underlying static graph $G = (V, E)$ of \mathcal{G} ;
- **Size knowledge Oracle (SKO):** given a TVG $\mathcal{G} = (V, E, \rho, \zeta)$, SKO provides the size of \mathcal{G} , that is $|V|$.

5.1 Detecting DCPA Liveness on Generic TVGs

In Section 4 we showed that the conditions guaranteeing the liveness property of reliable broadcast are strictly bounded to the network evolution. It follows that the knowledge provided by an FKO, in particular about the network evolution starting from the broadcast time t_{br} , is necessary to argue on liveness, unless further assumptions are taken into account. In the following, we clarify how a process can employ an FKO to detect Conscious Termination and Bounded Broadcast Latency.

Lemma 2. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG, let p_s be a node called source that broadcasts a message m at time t_{br} and let us assume f -locally bounded Byzantine failures. If p_s has access to an FKO then it is able to verify if there exists a TMKLO for the current broadcast on \mathcal{G} .*

Proof. In order to prove the claim it is enough to show an algorithm that verifies if a TMKLO exists, given the full knowledge of the TVG provided by FKO. Such algorithm works as follow: initially, the source p_s is placed in level $L_{t_{br}}$ of the TMKLO. Then, the snapshots characterizing the TVG have to be analyzed, starting from $G_{t_{br}}$ and following their order. In particular, for each snapshot G_{t_i} , $t_i \geq t_{br}$, we need to verify that:

1. edges with only one endpoint already included in some level of the TMKLO are up enough to satisfy $RCD()$ and
2. whenever $RCD()$ is satisfied for a given edge $e_{i,j}$, we need to check if it allows p_j to be part of the TMKLO as it satisfies one condition among AK2 and AK3.

The algorithm ends when a TMKLO is found or when all the snapshots have been analyzed (and in the latter case we can infer that no TMKLO exists for the considered message on the given TVG). Assuming that \mathcal{G} spans over T time instants, the complexity of this algorithm is:

$$O(|T||E|) + O(|V| + |E|) = O(|V| + |T||E|)$$

□

Theorem 4. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG, let p_s be a node called source that broadcasts a message m at time t_{br} and let us assume f -locally bounded Byzantine failures. If p_s has access to an FKO then it is able to detect if eventually every correct process will deliver m .*

Proof. The claim follows by considering that in order to assess the Conscious Termination of DCPA, the source process p_s needs to compute a TMKLO (i.e., it needs to check that eventually each correct process will be placed in a level) and due to Lemma 2 this can be done by accessing FKO. In particular, to detect Conscious Termination, a process p_i can first verify if the necessary condition holds and this can be done by computing a TMKLO with $k \geq f + 1$. If not, p_i can simply infer that m will not be delivered by every correct process. Contrarily, it can verify if the sufficient condition holds computing a TMKLO with $k \geq 2f + 1$. If it exists, p_i can infer that eventually every correct process will deliver the message otherwise, it needs to verify the necessary condition in every subgraph obtained by \mathcal{G} removing all the possible disposition of Byzantine processes (remind that getting this answer corresponds to solve an NP-Complete problem even considering a static networks, thus the same intractability follows also on dynamic networks). If the necessary condition is always satisfied, it can infer Conscious Termination otherwise not. □

Let us note that if a process has the capability of computing the *TMKLO* for a message m sent at time t_{br} , then it can also establish a lower bound and an upper bound on the time needed by every correct process to deliver m simply evaluating the maximum level of the *TMKLO* that satisfy respectively the necessary and the sufficient condition for DCPA.

Theorem 5. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG and let p_s be a node called source that broadcasts a message m at time t_{br} and let us assume f -locally bounded Byzantine failures. Let $P_{f+1} = \{L_{t_0}, L_{t_1} \dots L_{t_x}\}$ be the *TMKLO* with $k = f + 1$ associated to m and let t_{max}^{f+1} be the time associated to the last level of P_{f+1} . Let assume the existence of the *TMKLO* with $k = 2f + 1$ associated to m , $P_{2f+1} = \{L_{t_0}, L_{t_1} \dots L_{t_x}\}$, and let t_{max}^{2f+1} be the time associated to the last level of P_{2f+1} . The computed *TMKLOs* provide respectively a lower bound and an upper bound for *BL* such that:*

$$t_{max}^{f+1} - t_{br} \leq BL \leq t_{max}^{2f+1} - t_{br}$$

Proof. Lower Bound: Let us assume for the purpose of contradiction that *BL* can be lower than $t_{max}^{f+1} - t_{br}$. It follows that the last process p_i delivering m does it at a time $t_i < t_{max}^{f+1}$. Given the definition of *TMKLO* with $k = f + 1$, a level L_x is created each time that a process not yet inserted in the *TMKLO* delivers a message (due to AK2 or AK3). As a consequence, the last level of the *TMKLO* is created when the last process delivers the message. Thus, considering that p_i is the last process delivering the message, it follows that t_i is the time associated to the last level. Given P_{f+1} , it follows that $t_i = t_{max}^{f+1}$ and we have a contradiction.

Upper Bound: Let us assume for the purpose of contradiction that *BL* can be greater than $t_{max}^{2f+1} - t_{br}$. It follows that the last process p_i delivering m does it at a time $t_i > t_{max}^{2f+1}$. Given the definition of *TMKLO* with $k = 2f + 1$, a level L_x is created each time that a process not yet inserted in the *TMKLO* delivers a message (due to AK2 or AK3). As a consequence, the last level of the *TMKLO* is created when the last process delivers the message. Thus, considering that p_i is the last process delivering the message, it follows that t_i is the time associated to the last level. Given P_{2f+1} , it follows that $t_i = t_{max}^{2f+1}$ and we have a contradiction. \square

Remind that, as the sufficient condition we provided is not strict, a *TMKLO* with $k = 2f + 1$ could not exist even if the reliable broadcast is achievable. It is also possible to provide a stricter upper bound for *BL* as we explained inside the proof of Theorem 4, but is not practical to compute. Finally, let us remark that the knowledge on the underlying topology is not enough on dynamic networks to argue on liveness.

Remark 1. Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG and let p_s be a node called *source* that broadcasts a message m at time t_{br} and let us assume f -locally bounded Byzantine failures. If a process p_s has access only to a *PKO* (and not to an *FKO*) then it is not able to detect either *Conscious Termination* and *Bounded*

Broadcast Latency. Indeed, as we highlighted in section 4.2, moving on dynamic network the knowledge on the underlying graph is not enough, because specific sequences of edge appearances are required in order to guarantee the message propagation (let us take again Figure 1 as clarifying example). Thus, a PKO is not enough in arguing on liveness. The same can be said about Bounded Broadcast Latency as PKO provides no information about the time instants when the edges will appear.

5.2 Detecting DCPA Liveness on Restricted TVGs

Casteigts *et al.* [3] defined a hierarchy of TVG classes based on the strength of the assumptions made about appearance of edges. So far, we considered the most general TVG⁵. In the following, we consider two more specific classes of the hierarchy where we show that liveness can be detected using oracles weaker than FKO. In particular, we consider the following classes that are suited to model recurring networks:

- **Class recurrence of edges, ER** : if an edge e appears once, it appears infinitely often⁶.
- **Class time bounded recurrences, TBER**: if an edge e appears once, it appears infinitely often and there exist an upper bound Δ between two consecutive appearances of e ⁷.

Let us recall that assuming predicate $\text{RCD}(e_{i,j}, t) = \text{true}$ for every edge $e_{i,j}$ at some time t is necessary to guarantee liveness. While considering classes ER and TBER, such condition must be satisfied infinitely often, otherwise it is easy to show that the results presented in the previous section still apply. Let us also note that the conditions we defined in Section 4.2 are related to a single broadcast generated by a specific source p_s i.e., for a source p_s broadcasting a message at time t_{br} the conditions must hold from t_{br} on. Contrarily, exploiting the recurrence of edges it is possible to define different conditions that are valid for every broadcast from the same source p_s , independently from when it starts.

Detecting DCPA Liveness in ER TVG In this section, we prove that considering TVG of class ER, we can get the following results: (i) PKO (an oracle weaker than FKO) is enough to enable Conscious Termination, (ii) despite the more specific TVG considered, FKO is still required to establish upper bounds for *BL*. Intuitively, this results follows from the fact that PKO allows to determine whether a MKLO exists on the static underlying graph, and this is enough to detect if eventually every correct process will be able to deliver the message. However, given the absence of information on when each edge is going to appear, it is impossible to compute an upper bound on the time required to accomplish the broadcast.

⁵ Class 1 TVG according to Casteigts *et al.* [3]

⁶ Class 6 TVG in Casteigts *et al.* [3].

⁷ Class 7 TVG in Casteigts *et al.* [3].

Lemma 3. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG and let $G = (V, E)$ be the associated underlying graph. If p_s has access to a PKO then it can compute a MKLO on G .*

Proof. The PKO provides knowledge on the topology of G . We reminded in Definition 1 that the MKLO is a partition of the nodes on the base of a topological conditions. It follows that it is possible to verify the MKLO on G with PKO through a modified breath-first search [11]. \square

Lemma 4. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class ER that ensures RCD() infinitely often, let $G = (V, E)$ be the static underlying graph of \mathcal{G} , let p_s be a node called source and let us assume f -locally bounded Byzantine failures. If there exists the MKLO of $G = (V, E)$ associated to p_s then there always exists the TMKLO of \mathcal{G} associated to a message m sent by p_s with the same k .*

Proof. We prove the claim showing a mapping from MKLO to TMKLO. The source is placed inside the TMKLO at level t_{br} . Then, given the assumption on the channels and that every node in the MKLO has either (i) an edge connecting it with the source (ii) and/or k neighbors already included in MKLO, it follows that every node eventually satisfies at least one between AK2 and AK3. \square

Theorem 6. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class ER that ensures RCD() infinitely often, and let p_s be a node called source that broadcasts m at time t_{br} , and let us assume f -locally bounded Byzantine failures. If p_s has access to a PKO, then it is able to detect if eventually every correct process delivers m .*

Proof. It follows from Lemma 3 and Lemma 4 \square

Detecting DCPA Liveness in TBER TVG The liveness condition enabling CPA to enforce reliable broadcast relays on the network topology, therefore an oracle weaker than FKO cannot enable Conscious Termination unless further assumptions are made. On the other hand, the weaker oracle SKO allows a process to compute Bounded Broadcast Latency.

Lemma 5. *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class TBER where each edge $e_{i,j}$ reappears in at most Δ time instants satisfying RCD($e_{i,j}, t$). Let $\delta_{max} = \max(\zeta(e, t))$. Let p_s be the source and let us assume f -locally bounded Byzantine failures. The Broadcast Latency BL is upper bounded by*

$$BL \leq |V|(\delta_{max} + \Delta)$$

Proof. Given the assumptions on the TVG, we know that every edge reappears in Δ and satisfies RCD(). The worst case scenario, with respect to the message propagation, is the one in which every node has to wait Δ to forward a message. The worst case scenario, with respect to the network topology, is the one where every process has to wait the last one which has delivered to deliver (in other words, the partitions of the MKLO evaluated over the underlying graph $G(V, E)$, with the exception of the second level, have size equals to 1). \square

Lemma 6. Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class TBER where each edge $e_{i,j}$ reappears in at most Δ time instants satisfying $\text{RCD}(e_{i,j}, t)$. Let $\delta_{max} = \max(\zeta(e, t))$. Let p_s be the source and let us assume f -locally bounded Byzantine failures. Let $P_{2f+1} = \{L_{t_0}, L_{t_1} \dots L_{t_x}\}$ be the MKLO with $k = 2f + 1$ computed on the underlying graph $G = (V, E)$ (if exists) and let S_{2f+1} be size of P_{2f+1} .

An upper bound for BL can be computed from the MKLO with $k = 2f + 1$. In particular:

$$BL \leq S_{2f+1}(\delta_{max} + \Delta)$$

Proof. Given the assumptions on the TVG \mathcal{G} we know that every edge reappears in Δ and guarantees $\text{RCD}()$. The worst case scenario with respect the message propagation is the one where every node as to wait Δ to forward a message.

The bound follows by Theorem 5 and Lemma 4, noting that every node in level L_i delivers in $(\delta_{max} + \Delta)i$ time instants. \square

Theorem 7. Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class TBER where each edge $e_{i,j}$ reappears in at most Δ time instants satisfying $\text{RCD}(e_{i,j}, t)$. Let $\delta_{max} = \max(\zeta(e, t))$. Let p_s be a node called source that broadcasts m at time t_{br} , and let us assume f -locally bounded Byzantine failures. Let $P_{2f+1} = \{L_{t_0}, L_{t_1} \dots L_{t_x}\}$ be the MKLO with $k = 2f + 1$ associated to m and computed on the underlying graph $G = (V, E)$ (if exists) and let S_{2f+1} be size of P_{2f+1} . If p_s uses SKO or PKO, then p_s is able to compute an upper bound for BL. Specifically:

$$BL \leq |V|(\delta_{max} + \Delta) \text{ using SKO}$$

$$BL \leq S_{2f+1}(\delta_{max} + \Delta) \text{ using PKO}$$

Proof. The claim follows from Lemmas 6 and 5. \square

6 Moving to an Asynchronous System

In this work we assumed a synchronous distributed systems. In this section, we briefly discuss consequences of asynchrony on the safety and liveness of DCPA.

In Section 4.1, we showed that a reliable and authenticated channel is necessary and sufficient to enforce safety through CPA in an f -locally bounded failure model. Such channel properties are independent of the latency function. Indeed, they require that if a message m sent by a correct process is eventually received at its destination, it has not been compromised by the channel. As a consequence CPA (and DCPA as well) continues to enforce safety also on asynchronous dynamic networks.

In Section 4.2, we pointed out the need of having channels up long enough to allow the delivery of messages. This imposes constraints on the presence function due to the latency function. The asynchrony affects the latency function $\zeta(e, t)$ that basically is no more bounded. This makes impossible (in asynchronous system) to establish constraints for the liveness due to the fact it is no longer guaranteed the propagation of messages. It follows that we cannot argue on liveness of reliable broadcast on general TVG without making further assumptions.

In Section 5.2 we investigated about liveness in specialised classes of TVG. In particular, we showed in Theorem 6 that assuming recurrent *RCD* and having the knowledge on the underlying static graph it is possible to investigate about. It follows that, although RCDs are not identifiable over the time, if they are satisfied infinitely often, they enable the verification of liveness also in asynchronous systems.

7 Conclusion

We considered the reliable broadcast problem in dynamic networks represented by TVG. We analyzed the porting conditions enabling CPA to be correctly employed on dynamic networks. The analysis of this simple algorithm is important as it works exploiting only local knowledge. This contrasts to the best result so far in the same setting [15], that demands an exponential costs to check when a message can be delivered. Moreover, we presented necessary and sufficient conditions to ensure safety and liveness DCPA. We analyzed how much knowledge of the TVG is needed to detect whether the liveness condition is satisfied, and its cost. Our work is a starting point to identify more general parameters of dynamic networks that guarantees the fulfillment of the conditions we provided, both in a deterministic and probabilistic way. Other interesting points to address in future works are: i) the definition of a more realistic locally bounded failure model that takes also the time dimension into account, ii) the research of conditions on the dynamic network enabling nodes to conscious termination with just local information.

References

1. Augustine, J., Pandurangan, G., Robinson, P.: Fast byzantine agreement in dynamic networks. In: Fatourou, P., Taubenfeld, G. (eds.) ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013. pp. 74–83. ACM (2013)
2. Bhandari, V., Vaidya, N.H.: Reliable broadcast in radio networks with locally bounded failures. *IEEE Trans. Parallel Distrib. Syst.* **21**(6), 801–811 (2010)
3. Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N.: Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems* **27**(5), 387–408 (2012)
4. Castro, M., Liskov, B., et al.: Practical byzantine fault tolerance. In: OSDI. vol. 99, pp. 173–186 (1999)
5. Dolev, D.: Unanimity in an unknown and unreliable environment. In: Foundations of Computer Science, 1981. SFCS'81. 22nd Annual Symposium on. pp. 159–168. IEEE (1981)
6. Drabkin, V., Friedman, R., Segal, M.: Efficient byzantine broadcast in wireless ad-hoc networks. In: Dependable Systems and Networks, 2005. DSN 2005. Proceedings. International Conference on. pp. 160–169. IEEE (2005)
7. Gómez-Calzado, C., Casteigts, A., Lafuente, A., Larrea, M.: A connectivity model for agreement in dynamic systems. In: Euro-Par 2015: Parallel Processing - 21st

- International Conference on Parallel and Distributed Computing, Vienna, Austria, August 24-28, 2015, Proceedings. pp. 333–345 (2015)
8. Guerraoui, R., Huc, F., Kermarrec, A.: Highly dynamic distributed computing with byzantine failures. In: Fatourou, P., Taubenfeld, G. (eds.) ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013. pp. 176–183. ACM (2013)
 9. Ichimura, A., Shigeno, M.: A new parameter for a broadcast algorithm with locally bounded byzantine faults. *Information processing letters* **110**(12-13), 514–517 (2010)
 10. Koo, C.Y.: Broadcast in radio networks tolerating byzantine adversarial behavior. In: Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing. pp. 275–282. ACM (2004)
 11. Litsas, C., Pagourtzis, A., Sakavalas, D.: A graph parameter that matches the resilience of the certified propagation algorithm. In: International Conference on Ad-Hoc Networks and Wireless. pp. 269–280. Springer (2013)
 12. Maurer, A., Tixeuil, S.: Byzantine broadcast with fixed disjoint paths. *J. Parallel Distrib. Comput.* **74**(11), 3153–3160 (2014)
 13. Maurer, A., Tixeuil, S.: Containing byzantine failures with control zones. *IEEE Trans. Parallel Distrib. Syst.* **26**(2), 362–370 (2015)
 14. Maurer, A., Tixeuil, S.: Tolerating random byzantine failures in an unbounded network. *Parallel Processing Letters* **26**(1) (2016)
 15. Maurer, A., Tixeuil, S., Defago, X.: Communicating reliably in multihop dynamic networks despite byzantine failures. In: Reliable Distributed Systems (SRDS), 2015 IEEE 34th Symposium on. pp. 238–245. IEEE (2015)
 16. Pelc, A., Peleg, D.: Broadcasting with locally bounded byzantine faults. *Information Processing Letters* **93**(3), 109–115 (2005)
 17. Raynal, M., Stainer, J., Cao, J., Wu, W.: A simple broadcast algorithm for recurrent dynamic systems. In: 28th IEEE International Conference on Advanced Information Networking and Applications, AINA 2014, Victoria, BC, Canada, May 13-16, 2014. pp. 933–939 (2014)
 18. Tseng, L., Vaidya, N.H., Bhandari, V.: Broadcast using certified propagation algorithm in presence of byzantine faults. *Inf. Process. Lett.* **115**(4), 512–514 (2015)