



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-01712277v1

HAL Id: hal-01712277

<https://hal.science/hal-01712277v1>

Submitted on 19 Feb 2018 (v1), last revised 31 Oct 2018 (v3)

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 Bonomi^{*}, Giovanni Farina^{†*}, Sébastien Tixeuil[†]

^{*}Dipartimento di Ingegneria Informatica Automatica e Gestionale Antonio Ruberti,
Università degli Studi di Roma La Sapienza, Rome, Italy

bonomi@diag.uniroma1.it

[†]Sorbonne Université,

CNRS, Laboratoire d'Informatique de Paris 6, LIP6, F-75005 Paris, France

Giovanni.Farina@lip6.fr, 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.

Submission Type: Regular paper.

This paper can also be considered for brief announcement.

Eligible for best student paper award.

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 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.

This problem can be solved using cryptographic digital signatures [8, 10]. Yet, cryptographic infrastructures make the assumption that some components cannot be compromised (*e.g.* secret key distributor, public key provider, etc.), weakening the generality of the Byzantine model. For this reason, we focus in this paper on non-cryptographic solutions.

Related Works. In static networks (that is, in networks whose topology remains fixed during the entire execution of the protocol), a solution for the reliable broadcast problem has been initially provided for complete networks [14], requiring that no more than one third of the nodes are Byzantine (that is, $n > 3f + 1$, where n denotes the size of the network and f the maximum number of Byzantine nodes). In multi-hop networks, a necessary and sufficient condition is due to Dolev [9], and states that reliable broadcast can be solved if and only if the network is $2k + 1$ -connected. This global condition was replaced by a local condition on the number of Byzantine neighbors a node may have [13, 21]. All aforementioned works require high network connectivity.

Extending reliable broadcast service to sparse networks required to weaken the achieved guarantees [18, 17, 16]: (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 topology assumptions made by the various proposal may not hold anymore as the topology of the network evolves during execution. Without consideration of Byzantine failures, the problem of Broadcasting in dynamic networks was extensively studied [4, 6, 5, 3, 1]. Some core problems of distributed computing have been considered in the context of dynamic networks subject to Byzantine failures [11, 2] but, to the best of our knowledge, there exists a single contribution for the reliable communication problem, due to Maurer *et al.* [19]. This reliable broadcast can be seen as the dynamic network extension of the Dolev [9] solution for static networks, and assume that no more than f Byzantine can exist in the entire network. Also, the protocol to be executed by each node upon reception of a message grows exponentially with the size of the network, as each node computes the minimal cut of a set of paths traversed by each received message, making it unpractical in practice.

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.* [19]) in the hope that a practically efficient protocol can be derived in case the criterion is satisfied. Our starting point is the CPA protocol [13, 21], that

was originally designed for static networks. In particular, our contributions can be summarised 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} . 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 along time. More formally, we model the network as a *Time Varying Graph* (TVG) [7] 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 many 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 be also described as a sequence of static graphs $\mathcal{S}_{\mathcal{G}} = G_0, G_1, \dots, G_k$ 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 assumption on the evolution of the dynamic network are done.

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.

Communication model and Timing assumption. Every process is able to communicate only with its direct neighbours in the graph through message passing. We call *sender of a message m* the process p that transmits m . We assume *authenticated* and *reliable* point-to-point channels where (a) *authenticated* ensures that the identity of the sender of a message cannot be forged; (b) *reliable* guarantees that the channel delivers a message m if and only if (i) m was previously sent by p 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 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.

We assume that the computation time is negligible with respect to communication and we consider it equal to 0. Let us note that the timing assumption we are considering here model a dynamic synchronous system. We discuss the implication and consequences of lack of synchrony in the Appendix A.

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.

We considered the *f*-locally bounded failure model [13] 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 neighbours 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 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 ;
- **Liveness:** if a message m is broadcast by the source p_s , it is eventually delivered by every correct process.

3 The Certified Propagation Algorithm (CPA)

The Certified Propagation Algorithm (CPA) [13, 21] 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 from p_s and applies three acceptance policies (denoted by *AC*) to decide if m should be delivered and forwarded (i.e., transmitted also by nodes different from the source) by a process p_j . Specifically:

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

CPA correctness on static networks has been proved to be dependent on the network topology. In particular, Litsas *et al.* [15] 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 neighbours of p_s belong to level L_1 , and (iii) each node in a level L_i has at least k neighbours in a level L_j , with $j < i$. A *k*-level ordering is *minimum* if every node appears in the minimum level possible, otherwise it is *relaxed*.

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

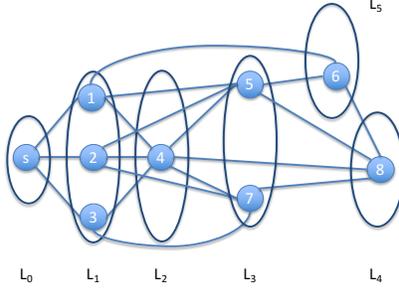


Figure 1: A graph $G = (V, E)$ allowing a k -level ordering with $k = 3$.

follows:

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

An example minimum k -level ordering (with $k = 3$) is shown in Figure 1. A relaxed k -level ordering for the same graph can be obtained by placing p_5 in a level and p_7 in another. Let us note that multiple relaxed k -level ordering may exist while the minimum k -level ordering is unique. Also, once a relaxed k -level ordering is defined, it can be reduced to a minimum k -level ordering.

For CPA to ensure reliable broadcast from p_s , a sufficient condition is that a k -level ordering exists, with $k \geq 2f + 1$. Conversely, a necessary condition is that $k \geq f + 1$ (see [15]). Considering Figure 1 and assuming $f = 1$, CPA ensures reliable broadcast from p_s as the sufficient condition is satisfied. Those conditions can be verified with an algorithm whose time complexity is polynomial in the size of the network through a modified Breadth-first search. When a graph $G = (V, E)$ satisfies only the necessary condition 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 exists (with $k = f + 1$) in every sub-graph G' obtained from G by removing all nodes corresponding to possible Byzantine position in the f -locally bounded assumption. The verification of the strict condition has been proven to be NP-Hard [12].

The graph $G = (V, E)$ depicted in Figure 2a does not match the sufficient condition but it satisfies the necessary one for CPA allowing the definition a k -level ordering with $k = f + 1$. Computing all the possible sub-graphs obtained by G by removing all the possible placements of a 1-locally bounded adversary, it is possible to deduce that a k -level ordering, with $k = f + 1$, still exists on every subgraph. Thus, CPA ensures reliable broadcast on G . Contrarily, the graph $G' = (V, E')$ depicted in Figure 2b has a sub-graph (the one obtained by removing p_4 as a possible faulty process) that does not admit a k -level ordering with $k = f + 1$ anymore. Thus, in this case, CPA does not ensure reliable broadcast.

4 The Certified Propagation Algorithm on Dynamic Networks

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

Let us consider the TVG shown in Figure 3, and suppose process p_2 is Byzantine. If we consider the static underlying graph $G = (V, E)$ shown in Figure 3b, it is easy to verify that running CPA from the source

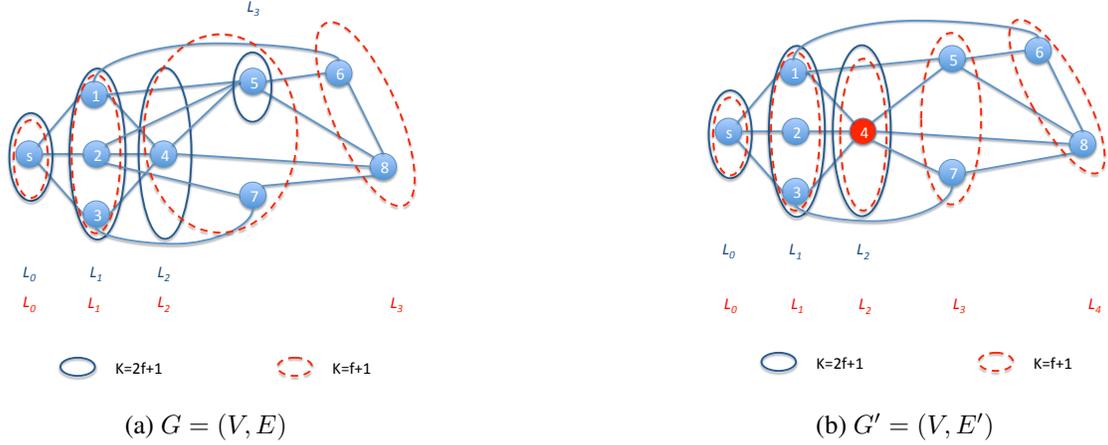


Figure 2: Example of graphs where a k -level ordering exists for $k = f + 1$ but not for $k = 2f + 1$ (with $f = 1$).

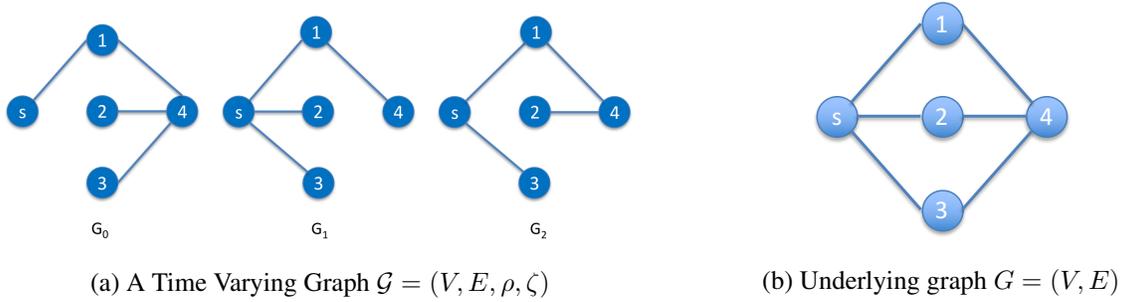


Figure 3: Example of a simple TVG and its underlying static graph.

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 times¹ as shown in Figure 3a, 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.

4.1 CPA Safety in Dynamic Networks

In the following, we show that 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 in static graphs).*

Proof An authenticated channel guarantees that the identity of the sender of a message cannot be forged.

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

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 sent. Without this assumption, an unreliable channel can simulate a Byzantine process that is not bound to authenticated channels, also invalidating the *f*-locally bounded assumption. \square *Lemma 1*

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 is delivered by p_i , m 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 *Theorem 1*

4.2 CPA Liveness in Dynamic Networks

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, assuming that channels do not lose messages is sufficient to guarantee their propagation.

In dynamic networks, this is not anymore 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 times are determined by the edge latency function $\zeta(e, t)$. As a consequence, in order to ensure 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 channel disappears while the message is still traveling.

Thus, in addition to topological constraints 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 neighbours infinitely often.

To ease of explanation, we will refer to this extended version of CPA as Dynamic CPA (DCPA).

We now characterize the conditions enabling a 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:*

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

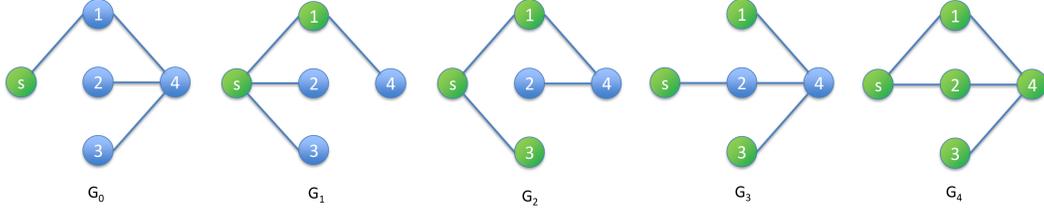


Figure 4: TVG example

Now that we are able to express constraints on each edge through the $RCD()$ predicate, we need to define those 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:

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

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) is a partition the nodes in levels L_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.

Let us consider the TVG presented in Figure 4: 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 $RCD(p_s, p_1, 0) = \text{true}$ due to the presence function $\rho(\langle p_s, p_1 \rangle, \tau) = 1, \forall \tau \in [0, 1]$).
- On process p_3 , the acceptance function evaluates to 1 for $t \geq 2$, 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., $RCD(p_i, p_4, t'_i) = \text{true}$ for $p_i = p_1, t'_i = 1$, and for $p_i = p_3, t'_i = 3$).

In the sequel, we 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 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 instant t such that the $2f + 1$ -acceptance function $\mathcal{A}_k(p, t)$ is equal to 1 for every node of the network.

Now, 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 infinitively 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 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 scenarii:

- **Case 1:** 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:** RCD() is satisfied between p_j and $2f + 1$ nodes p_i where AK2 or AK3 is already satisfied. Inductively, as the last considered nodes 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 2*

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. If there does not exist any 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 $k > f$, then m broadcast using DCPA cannot be delivered by every correct process in \mathcal{G} .

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. $\exists p \in \Pi \mid \forall t \in \mathbb{N}, \mathcal{A}_{f+1}(p, t) = 0$.

Now, 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 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 satisfied, and that 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 because eventually every process satisfies one of the conditions AK1, AK2 or AK3, and the claim follows.

□*Theorem 3*

5 On the Detection of DCPA Liveness

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

- **(Conscious Termination):** Given a message m sent by a source p_s on TVG \mathcal{G} , is p_s able to detect if m is eventually delivered by every correct process?
- **(Bounded Broadcast Latency):** Given a message m 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 formalise 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 the following we prove that assuming a generic TVG the FKO is sufficient to detect Conscious Termination and Bounded Broadcast Latency while PKO is not. This is due to the fact that FKO enables a process to compute a TMKLO.

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 a 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 t_{br} of the TMKLO. Then, the snapshots characterizing the TVG have to be analyzed, starting from $G_{t_{br}+1}$ and following their order. In particular, for each snapshot G_{t_i} 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 $\text{RCD}()$ 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:

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

A more detailed description of the algorithm is delegated to the Appendix B. \square *Lemma 2*

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 a 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 (and in general any process in the system) 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 same intractability follows also on dynamic networks). If the necessary condition is always satisfied, it can infer Conscious Termination otherwise not. \square *Theorem 4*

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 TMKLOs that satisfy respectively the necessary and the sufficient condition for DCPA.

Lemma 3 *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} . The computed TMKLO provides a lower bound for BL such that:*

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

Proof Let us assume by 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. \square *Lemma 3*

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 possible. It is also possible to provide a stricter upper bound for BL as we explained in Theorem 4, but is not practical to compute.

Lemma 4 *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 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 TMKLO provides an upper bound for BL such that:*

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

Proof Let us assume by 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. $\square_{Lemma 4}$

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 , let $P_{2f+1} = \{L_{t_0}, L_{t_1} \dots L_{t_y}\}$ be the TMKLO with $k = 2f + 1$ associated to m (if exists) and let t_{max}^{f+1} and t_{max}^{2f+1} be respectively the times associated to the last level of P_{f+1} and P_{2f+1} . If p_s has access to a FKO then it is able to compute an upper bound and a lower bound on BL such that:*

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

Proof The claim simply follows from Lemma 3 and Lemma 4 $\square_{Theorem 5}$

Theorem 6 *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.*

Proof It has been shown in [15, 20] the necessary and sufficient conditions to argue on the liveness property of CPA. Such conditions are topological. As it is not possible to argue on liveness of CPA on static networks without the knowledge of the topology, it follows the same also on dynamic networks. However, 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 3 as clarifying example). Thus, PKO is not enough arguing on liveness. The same can be said about Bounded Broadcast Latency as PKO provides no information about the time instants where the edge will appear. $\square_{Theorem 6}$

In this section, we shown that FKO is sufficient to detect liveness while PKO is not. Given that the complete knowledge on network topology is required by CPA to argue on liveness on static networks, we conjecture that it is not possible to detect liveness in generic TVG with an oracle weaker than FKO.

5.2 Detecting DCPA Liveness on Restricted TVGs

Casteigts *et al.* [7] defined a hierarchy of TVG classes based on the strength of the assumptions made on appearance of edges. So far, we considered the most general TVG². 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 infinitively often³.
- **Class time bounded recurrences, TBER**: if an edge e appears once, it appears infinitively often and there exist an upper bound Δ between two consecutive appearances of e ⁴.

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 infinitively often otherwise, it is easy to show that the results presented in the previous section still apply.

5.2.1 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 is 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 deliver a message.

Lemma 5 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class ER, 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 $G = (V, E)$ be the underlying graph of \mathcal{G} . 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 have reminded in Definition 1 that the MKLO is a partition of the nodes on the base of topological conditions. It follows that it is possible to verify the MKLO on G with PKO through a modified breath-first search [15]. □_{Lemma 5}

Lemma 6 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG, let p_s be the source which broadcasts a message m at time t_{br} and let us assume f -locally bounded Byzantine failures. If there exists a TMKLO of \mathcal{G} associated to m then there exist also the MKLO of $G = (V, E)$ associated to p_s with the same k .*

Proof We prove the claim showing a mapping from TMKLO to MKLO. Let us enumerate the partitions of *TMKLO* with natural numbers starting from 0, namely for $P_k = \{L_{t_{br}}, L_{t_1} \dots L_{t_x}\} \rightarrow P'_k = \{L_0, L_1 \dots L_{n'}\}$.

Let us analyze P'_K . The source p_s is in level L_0 . The processes $p_i \in L_1$ are neighbors of the source. The processes placed inside the *TMKLO* due to *AK2* are neighbors of the source, thus they can be moved in level L_1 . The remaining nodes could be neighbors or not of the source in G . The former can be moved in L_1 . The

²Class 1 TVG according to Casteigts *et al.* [7]

³Class 6 TVG in Casteigts *et al.* [7].

⁴Class 7 TVG in Casteigts *et al.* [7].

partition P'_k is now a relaxed k -level ordering. We reminded in Section 3 that it is possible to obtain a *MKLO* from a relaxed k -level ordering. Thus, after the shown shiftings, the partition P'_K is a *MKLO*. $\square_{\text{Lemma 6}}$

Lemma 7 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class ER that ensures $\text{RCD}()$ infinitively 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. $\square_{\text{Lemma 7}}$

Theorem 7 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class ER that ensures $\text{RCD}()$ infinitively 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 5 and Lemma 7 $\square_{\text{Theorem 7}}$

Theorem 8 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class ER that ensures $\text{RCD}()$ infinitively often, let $\delta_{min} = \min(\zeta(e, t))$ 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. Let $P_{f+1} = \{L_{t_0}, L_{t_1} \dots L_{t_x}\}$ be the *MKLO* with $k = f + 1$ computed on the underlying graph $G = (V, E)$ and let S_{f+1} be size of P_{f+1} .*

*A lower bound for BL can be computed from the *MKLO* with $k = f + 1$. In particular:*

$$\delta_{min} S_{f+1} - t_{br} \leq BL$$

Proof It follows from Lemma 3 and Lemma 7. $\square_{\text{Theorem 8}}$

Theorem 9 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class ER, 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 a process p_s uses *SKO*, then it is not able to detect either *Conscious Termination* or *Bounded Broadcast Latency*.*

Proof The TVG of class ER guarantees the reappearances of edges without no guarantees about the time between the appearances. Then, the *SKO* provides no information about the topology. Thus, the claim follows considering also Theorem 6. $\square_{\text{Theorem 9}}$

5.2.2 Detecting DCPA Liveness in TBER TVG

In this section, we demonstrate that considering a TVG of class TBER, it is possible: (i) to use *PKO* to enable *Conscious Termination*, and (ii) to use *SKO* to compute *Bounded Broadcast Latency*.

Theorem 10 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class TBER, 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 a process p_s uses *SKO*, then p_s is not able to enable *Conscious Termination*.*

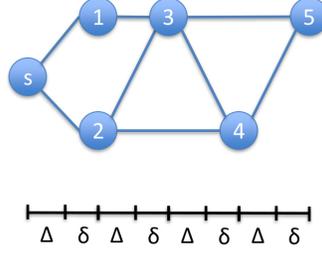


Figure 5: Worst case topology for computing BL.

Proof The TVG of class TBER guarantees timely reappearances of edges but the SKO provides no information about the topology. Thus, the claim follows considering also Theorem 6. $\square_{\text{Theorem 10}}$

Theorem 11 Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class TBER 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 uses PKO, then p_s is able to enable Conscious Termination.

Proof The claim simply follows from Theorem 7 considering that TBER is a class of TVG contained in ER. $\square_{\text{Theorem 11}}$

Lemma 8 Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a Class TBER TVG 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. 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 $\text{RCD}()$. The worst case scenario, with respect the message propagation, is the one in which every node has to wait Δ to forward a message. The worst case scenario, with respect 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). We depicted such a worst case scenario in Figure 5.

$\square_{\text{Lemma 8}}$

Lemma 9 Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a Class TBER TVG 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 Lemma 4 and Lemma 7, noting that every node in level L_i delivers in $(\delta_{max} + \Delta)i$ time instants. $\square_{Lemma 9}$

Theorem 12 *Let $\mathcal{G} = (V, E, \rho, \zeta)$ be a TVG of class TBER 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 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 Lemma 9 and Theorem 8. $\square_{Theorem 12}$

6 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 network. 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 [19], 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 analysed how much knowledge of the TVG is needed to detect whether the liveness condition is satisfied, and its cost. Table 1 summarizes all our results. Our work is a starting point to identify more general parameters of dynamic networks that guarantees the fulfilment of the conditions we provided, both in a deterministic and probabilistic way. Another interesting point to address in future work is the possibility to define a f -locally bounded failure model on a time dimension, considering the failures to be *f-locally bounded* in a temporal interval.

Table 1: Solvability Conscious Termination and Broadcast Latency

		Conscious Termination	Broadcast Latency
FKO	Generic TVG	Th 4	Th 5
PKO	Generic TVG	<i>Impossible</i> - Th 6	<i>Impossible</i> - Th 6
	ER	Th 7	Th 8
	TBER	Th 7	Th 12
SKO	Generic TVG	<i>Impossible</i> - Th 6	<i>Impossible</i> - Th 6
	ER	<i>Impossible</i> - Th 9	<i>Impossible</i> - Th 9
	TBER	<i>Impossible</i> - Th 10	Th 12

References

- [1] Yehuda Afek, Baruch Awerbuch, and Eli Gafni. Applying static network protocols to dynamic networks. In *Foundations of Computer Science, 1987., 28th Annual Symposium on*, pages 358–370. IEEE, 1987.
- [2] John Augustine, Gopal Pandurangan, and Peter Robinson. Fast byzantine agreement in dynamic networks. In Panagiota Fatourou and Gadi Taubenfeld, editors, *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*, pages 74–83. ACM, 2013.
- [3] Baruch Awerbuch. On the effects of feedback in dynamic network protocols. *Journal of Algorithms*, 11(3):342–373, 1990.
- [4] Baruch Awerbuch and Shimon Even. Efficient and reliable broadcast is achievable in an eventually connected network. In *Proceedings of the third annual ACM symposium on Principles of distributed computing*, pages 278–281. ACM, 1984.
- [5] Baruch Awerbuch, Yishay Mansour, and Nir Shavit. Polynomial end-to-end communication. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 358–363. IEEE, 1989.
- [6] Baruch Awerbuch and Michael Sipser. Dynamic networks are as fast as static networks. In *Foundations of Computer Science, 1988., 29th Annual Symposium on*, pages 206–219. IEEE, 1988.
- [7] Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems*, 27(5):387–408, 2012.
- [8] Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.
- [9] Danny Dolev. The byzantine generals strike again. *Journal of algorithms*, 3(1):14–30, 1982.
- [10] Vadim Drabkin, Roy Friedman, and Marc Segal. Efficient byzantine broadcast in wireless ad-hoc networks. In *Dependable Systems and Networks, 2005. DSN 2005. Proceedings. International Conference on*, pages 160–169. IEEE, 2005.
- [11] Rachid Guerraoui, Florian Huc, and Anne-Marie Kermarrec. Highly dynamic distributed computing with byzantine failures. In Panagiota Fatourou and Gadi Taubenfeld, editors, *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*, pages 176–183. ACM, 2013.
- [12] Akira Ichimura and Maiko Shigeno. A new parameter for a broadcast algorithm with locally bounded byzantine faults. *Information processing letters*, 110(12-13):514–517, 2010.
- [13] Chiu-Yuen Koo. Broadcast in radio networks tolerating byzantine adversarial behavior. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 275–282. ACM, 2004.
- [14] Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(3):382–401, 1982.

- [15] Chris Litsas, Aris Pagourtzis, and Dimitris Sakavalas. A graph parameter that matches the resilience of the certified propagation algorithm. In *International Conference on Ad-Hoc Networks and Wireless*, pages 269–280. Springer, 2013.
- [16] Alexandre Maurer and Sébastien Tixeuil. Byzantine broadcast with fixed disjoint paths. *J. Parallel Distrib. Comput.*, 74(11):3153–3160, 2014.
- [17] Alexandre Maurer and Sébastien Tixeuil. Containing byzantine failures with control zones. *IEEE Trans. Parallel Distrib. Syst.*, 26(2):362–370, 2015.
- [18] Alexandre Maurer and Sébastien Tixeuil. Tolerating random byzantine failures in an unbounded network. *Parallel Processing Letters*, 26(1), 2016.
- [19] Alexandre Maurer, Sébastien Tixeuil, and Xavier Defago. Communicating reliably in multihop dynamic networks despite byzantine failures. In *Reliable Distributed Systems (SRDS), 2015 IEEE 34th Symposium on*, pages 238–245. IEEE, 2015.
- [20] Aris Pagourtzis, Giorgos Panagiotakos, and Dimitris Sakavalas. Reliable broadcast with respect to topology knowledge. *Distributed Computing*, 30(2):87–102, 2017.
- [21] Andrzej Pelc and David Peleg. Broadcasting with locally bounded byzantine faults. *Information Processing Letters*, 93(3):109–115, 2005.

Appendix A Moving to an Asynchronous System

In this work we considered a synchronous 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 from the latency function as they just 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 that depend on the latency function. The asynchrony affects the latency function $\zeta(e, t)$ that basically is no more bounded and known by processes. This makes impossible (in asynchronous system) to establish constraints for the liveness due to the fact that the latency function is undefined and thus it is not possible to guarantee the propagation of messages. It follows that it not possible to argue on liveness without making stronger assumptions.

In subsection 5.2 we investigated about liveness in specialised classes of TVG. In particular, we showed in Theorem 7 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 infinitively often, they enable the verification of liveness also in asynchronous systems.

Appendix B An algorithm for verifying if a TMKLO exists

We present an algorithm enabling a process equipped with a FKO to compute TMKLO, if it exist.

The complete pseudo-code is provided in Figure 6. It is composed by:

- the main part, `check_liveness_and_BL_bounds`, where the liveness condition is verified and bounds for BL are computed,
- the code that computes a *TMKLO*, `compute_temporal_minimum_level_ordering`.

Going into details of the main part, it leverages on the *TMKLO* to verify the necessary condition (Theorem 3) on lines 2 – 3 and the sufficient condition (Theorem 2) on lines 5 – 6. It is clear that there exist cases where the necessary condition is verified while the sufficient one is not and, in order to argue on Conscious Termination and Bounded Broadcast Latency, we need to verify the necessary condition considering any possible Byzantine placement over the network (lines 12 – 18) as in [15].

Moving on the part related to the *TMKLO* computation, the corresponding function analyzes the snapshots characterizing the TVG have to be analyzed, starting from $G_{t_{br+1}}$ and following their order, verifying : (i) if the edges with only one endpoint already included in some level of the *TMKLO* are up enough to satisfy `RCD()` and (ii) 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. In the following we provide a summary of variables used in the algorithm:

- `t_level_ordering` is a map $\langle \text{process ID, acceptance time instant} \rangle$, it corresponds to the *TMKLO*;
- `transmitting_edges` is the set of edges which are under analysis for `RCD()`;
- `onGoing_RCDs` is a map $\langle \text{edge, i} \rangle$ which keeps track of the duration of ongoing `RCD()`;

- *succeeded_RCDs* is the set of edges which have succeeded in $RCD()$.

The algorithm takes as input:

- the TVG description;
- the ID of the source node;
- t_{br} , namely the the time instant when the broadcast starts;
- the upper bound for the channel latency δ ; (v) the value of k for *TMKLO*.

It computes the k -*TMKLO* or it returns an incomplete partition. It ends when all processes are included in the *TMKLO* (lines 29), or all network snapshots are analyzed and returns an incomplete *TMKLO* (line 30).

In particular, the k -acceptance function fixes the level of the source (AK1, line 21). Subsequently, the TVG is analyzed through its consecutive snapshots (line 22). Given a time instant t , the edges under analysis are the ones that interconnect processes who are already in the *TMKLO* with others not included (Lines 44 – 48). If an edge does not remain up for δ consecutive time instants, it is not possible for such an edge to ensure $RCD()$. This means that all ongoing transmissions (edges), registered inside the variable *onGoing_RCDs*, have to be compared with the ones included in *transmitting_edges*. In lines 42 – 43 all ongoing $RCD()$ over edges that disappear are dropped. Subsequently, the state of ongoing $RCD()$ is updated considering the transmitting edges (line 26). When one among AK2 or AK3 is verified, the specific nodes are inserted inside the *TMKLO* (lines 36 – 41).

B.1 Algorithm complexity

We analyze the computation complexity of the presented algorithm, focusing on the single functions:

- The `get_transmitting_edges(t)` function analyses the edges present in a specific time instants, filtering out the ones that are not relevant for the *TMKLO*. The filtering process has constant complexity. It follows that the function has complexity $O(|E|)$.
- The `check_ongoing_RCDs()` performs the set difference of collection of size $O(|E|)$. Thus, this function has complexity $O(|E|)$.
- The `check_RCD()` function accesses an entry map and compares the returned value with δ . The *if* condition on line 12 is verified once for every edge and every process is inserted inside the *TMKLO* only once. It follows that the function has cost $O(|V| + |E|)$.

It follows the complexity of computing the *TMKLO* is:

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

We have to notice that even assuming reliable channels with instantaneous delivery (namely, $\delta = 0$ guaranteeing that every message sent over a present channel is delivered) the complexity of the algorithm does not change, because all the edges included inside the snapshot G_{t_i} have still to be analyzed.

We have shown that if the sufficient condition is verified, a process has to compute two *TMKLO* to get knowledge on conscious termination and bounded broadcast latency. Thus, the complexity of the entire algorithm is polynomial with respect our metrics (nodes, edges and time). In the worst case scenario, placed between the gap among the sufficient and necessary conditions, all possible Byzantine placements have to be verified. This, turn our algorithm in an exponential one as shown in [15].

```

check_liveness_and_BL_bounds(source, tbr, δ, f):
(01) graph ← FKO;
(02) TMKLOnec ← compute_temporal_minimum_level_ordering(graph, source, tbr, δ, f + 1);
(03) if TMKLOnec.is_complete(graph.size) then
(04)   BLmin ← TMKLOnec.max_value() - tbr;
(05)   TMKLOsuf ← compute_temporal_minimum_level_ordering(graph, source, tbr, δ, 2f + 1);
(06)   if TMKLOsuf.is_complete(graph.size) then
(07)     liveness ← 1;
(08)     BLmax ← TMKLOsuf.max_value() - tbr;
(09)   else
(10)     BLmax ← 0;
(11)     liveness ← 1;
(12)   for each (subgraph ∈ get_subgraphs_f_local(graph, source, f)) do
(13)     TMKLOi ← compute_temporal_minimum_level_ordering(subgraph, source, tbr, δ, f + 1);
(14)     if TMKLOi.is_complete(subgraph.size) then
(15)       BLmax ← max(BLmax, TMKLOi.max_value());
(16)     else
(17)       liveness ← 0;
(18)       break
(19)   else
(20)     liveness ← 0;

compute_temporal_minimum_level_ordering(graph, source, tbr, δ, k):
(21) t_level_ordering ← t_level_ordering ∪ {< source, tbr >};
(22) for (t ← tbr; t ≤ graph.max_T; t ← t + 1) do
(23)   transmitting_edges ← get_transmitting_edges(t);
(24)   check_ongoing_RCDs();
(25)   for each (e ∈ transmitting_edges) do
(26)     onGoing_RCDs ← onGoing_RCDs ∪ {< e, +1 >};
(27)     check_RCD(e, t);
(28)   if t_level_ordering.keys = graph.nodes then
(29)     return t_level_ordering
(30) return t_level_ordering

check_RCD(e, t):
(31) edge_duration ← onGoing_RCDs.get(e)
(32) if edge_duration = δ then
(33)   out_node ← not_included(e)
(34)   succeeded_RCDs ← succeeded_RCDs ∪ {e}
(35)   onGoing_RCDs ← onGoing_RCDs \ {e}
(36)   if source ∈ e.endpoints then
(37)     t_level_ordering ← t_level_ordering ∪ {< out_node, t >}
(38)   else
(39)     k_acceptance ← k_acceptance{< out_node, +1 >}
(40)     if k_acceptance.get_out_node() = k then
(41)       t_level_ordering ← t_level_ordering ∪ {< out_node, t >}

check_ongoing_RCDs():
(42) interrupted_RCDs ← onGoing_RCDs \ transmitting_edges
(43) onGoing_RCDs ← onGoing_RCDs \ interrupted_channels

get_transmitting_edges(t):
(44) new_transmitting_edges ← ∅
(45) for each (e ∈ graph.getSnapshot(t).edges) do
(46)   if |e.endpoints ∩ t_level_ordering.keys| = 1 AND e ∉ succeeded_RCDs then
(47)     new_transmitting_edges ← new_transmitting_edges ∪ {e}
(48) return new_transmitting_edges

not_included(e):
(49) for each (p ∈ e.endpoints) do
(50)   if t_level_ordering.keys ∩ p ≠ ∅ then
(51)     return p

```

Figure 6: Process code with *FKO*

Lemma 10 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 p_s has access to a FKO then it is able to

- answer about the Conscious Termination of m
- compute upper bound and a lower bound for the broadcast latency of m

with a cost that is polynomial in the best case and exponential in the worst case with respect our metrics (nodes, edges and time).

Proof It follows from Lemmas 2 and Theorems 4 and 5. The complexity cost derives from the algorithm computing a TMKLO: If the sufficient condition is verified (Theorem 2), the conscious termination and bounds for the broadcast latency can be obtained with an algorithm polynomial with respect our metrics (nodes, edges and time), specifically in $O(|V| + |T||E|)$.

Otherwise in the worst case scenario, where the necessary condition (Theorem 3) holds but the sufficient one does not, we have to test any possible Byzantine placement to get our answers: namely the necessary condition has to be verified on every subgraph where a possible set of Byzantine node (f -locally bounded) is removed. This verification has an exponential cost [15]. □*Lemma 10*