Broadcasting Information in Multi-hop Networks prone to Mobile Byzantine Faults *

Silvia Bonomi[1], Giovanni Farina[2,1] ✉, and Sébastien Tixeuil[2]

[1] Sapienza Università di Roma, Rome, Italy
bonomi@diag.uniroma1.it
[2] Sorbonne Université, CNRS, LIP6, F-75005 Paris, France
giovanni.farina@lip6.fr, sebastien.tixeuil@lip6.fr

Abstract. Every non-trivial distributed application needs to exchange information in order accomplish its task, and reliable communication primitives are fundamental in failures prone distributed systems to guarantee correct message exchanges between parties.
Their implementation becomes particularly challenging when considering distributed systems where processes are arranged in a multi-hop network and each of them may temporary and continuously be compromised by an attacker during the execution. Although some fundamental problems (such as the register implementation and the agreement) were investigated considering Mobile Byzantine Faults (MBF), most of the contributions consider a fully connected communication network.
In this paper we analyze the specific difficulty of ensuring reliable communication between parties in a distributed system affected by Mobile Byzantine Faults (compared to the case where the Byzantine failures are static), showing that such a problem is essentially impossible to solve in asynchronous systems with MBF, and we propose a synchronous protocol providing reliable communication both in complete networks and specific multi-hop topologies.

Keywords: Reliable communication · Mobile Byzantine Faults · multi-hop networks.

1 Introduction

Distributed systems are often prone to failures, given the multitude of interconnected components they are composed of, and protocols that are deployed on them are usually designed to guarantee correct execution despite fault occurrences. Besides, distributed systems are more and more frequently subject also to external attackers, who aim to penetrate and compromise them.

* 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 wishes to thank *Université Franco-Italienne/Universitá Italo-Francese* (UFI/UIF) for supporting his mobility through the Vinci grant 2018.

Processes in a distributed system need to communicate in order to achieve non-trivial goals. Indeed, several reliable communication primitives have been defined to guarantee integrity, delivery and authorship of messages exchanged even in case of arbitrary failures. The reliable communication solutions proposed so far mostly put constraints on the spatial distribution of failures or on their duration. Such assumptions capture most of the internal misbehavior that may occurs in a system: data corruptions, link failures, machine faults, etc. On the other side, external malicious attackers commonly start compromising some machines and then they use them to move over the system till reaching their targets, and the research handling such kind of attacks mostly focus on their prevention, detection and reaction.

In this paper, we analyze the specific difficulties of ensuring reliable communication in distributed system affected by Mobile Byzantine Faults (compared to the case where the Byzantine failures are static), showing that reliable communication in asynchronous systems is essentially impossible, and then we propose a synchronous protocol solving reliable communication both in complete networks and specific multi-hop topologies.

2 Related works

The reliable communication problem has been extensively investigated considering *static* Byzantine process failures. Dolev [10] provided the seminal contribution addressing this problem in general networks with a globally bounded number of faulty processes. Subsequently, several failure distributions have been considered, such as neighborhood-bounded [12, 19, 21, 25], probabilistic [16, 20], and the general adversary model [19]. Weaker problem specifications have been proposed to allow solving the reliable communication problem in loosely connected network [13, 15], and dynamic networks have also been considered [2, 17].

In complete communication networks, non-static Byzantine faulty processes were considered by Reischuk [22] who proposed an algorithm solving the Byzantine agreement in the case of f malicious agents that remain stationary on f processes only for a given period of time. Later, Ostrovsky and Yung [18] introduced the notion of an adversary that can inject and distribute faults in the system at a constant rate and they proposed solutions (mixing randomization and self-stabilization) for tolerating the attacks of mobile viruses. Then, Garay [11] considered processes proceedings in synchronous rounds composed by three phases (send, receive, and compute), and Byzantine mobile agents able to move between one process to another during the lifetime of the system. Several subsequent works later specialized his model, making alternative hypothesis on the unawareness of processes of being faulty [24], assuming correct processes sending non-equivocal messages [1], channels delays [23], decoupling the system evolution from the agents movements [5]. All aforementioned works for the mobile attacker model addressed either the Byzantine agreement, the approximate Byzantine agreement [7], or the register abstraction [4] problems in complete

networks. Most related to our work is the solution by Sasaki et al. [24], that is detailed in section 6.

3 System model

Process definition and communication model. We consider a distributed system composed by a set of n processes $\Pi = \{p_1, p_2, \dots p_n\}$, each associated with an unique identifier. Processes communicate by exchanging messages via reliable and authenticated point-to-point links i.e., messages can neither be lost or altered by the links and the identity of the sender of any message cannot be forged. Processes and their links can be abstracted by an undirected graph $G = (\Pi, E)$ where the set of nodes is represented by the processes of the system and the set of edges E contains an element $e_{i,j}$ if and only if there exists a link between processes p_i and p_j. Two processes p_i and p_j can exchange messages only if there is a link between them.

Time assumptions and computational model. Unless differently stated, we consider a *synchronous* system [9]. Specifically, we assume one where the computation evolves in sequential synchronous rounds $r_0, r_1, \dots r_i \dots$ (with $i \in \mathbb{N}$). Every round is divided in three phases: (i) *send* where processes send messages through their links for the current round, (ii) *receive* where processes receive all messages sent at the beginning of the current round, and (iii) *computation* where processes execute a deterministic distributed protocol $\mathcal{P}$ and generate the messages to be sent during the subsequent round. We assume a tamper-proof read-only memory on every process where the code of $\mathcal{P}$ is stored.

Failure model. We assume that the system is affected by *Mobile Byzantine Faults* (MBF) [1, 8, 11, 24]. Informally, in the mobile Byzantine failure model, faults are represented by f computationally unbounded agents that move between processes. When an agent is on a process p_i, it forces p_i to behave as a *Byzantine faulty* process (i.e., it may corrupt its local variables, forces it to execute an arbitrary protocol, to send arbitrary messages, to omit sending messages, etc.). We assume that, at every round r_i, every mobile Byzantine agent is placed on at most one process p_j and that it can move from p_j to another process p_k only if there is a link between the two. The movement of the Byzantine agents is characterized by the *roaming pace* parameter ρ, that is the minimum amount of time between two displacements of an agent. We assume that the Byzantine agents can only move in between the computation and the send phase [11, 24], thus $\rho \geq 1$ round.

We alternatively consider either an *aware* [24] or *unaware* [11] mobile Byzantine failure model: in the former case a process knows about a mobile agent that is moving away from it, in the latter it does not. At every round r_i, a process p_j is either *correct* or *Byzantine faulty*. Precisely, p_j is faulty if a mobile Byzantine agent is on it at r_i, or it is correct otherwise and it executes the distributed protocol $\mathcal{P}$. Notice that every process backs to execute protocol $\mathcal{P}$

right after a Byzantine agent moved away and that the failure state of a process cannot change during a message transmission (send - receive phases). In the aware mobile Byzantine failure model, we refer with *cured* process to a correct one at round r_i that was Byzantine faulty at round r_{i-1}. We assume that every cured process wipes all of its local variables at the beginning of the round.

Link specifications. The point-to-point reliable and authenticated links guarantee the following properties [9]: *Reliable delivery* - if a correct process sends a message m to a correct process p_j, then p_j eventually receives m; *No duplication* - no message is delivered by a link to a process more than once; *Authenticity* - if some correct process p_j receives a message m with sender p_s, then m was previously sent to p_j by p_s.

3.1 Graph metrics

We briefly recall some graph metrics that are employed to characterize reliable communication correctness conditions.

Sasaki et al. [24] defined $G(\alpha, \beta)$ as the class of graphs $G = (V, E)$ such that, for any pair i, j of vertices in V, there are α disjoint paths connecting i and j, whose length (in terms of the number of edges) is at most β.

A *k-clique community* is a graph defined as the union of all k-cliques (i.e., complete subgraphs of size k) that can be reached from each other through a series of adjacent k-cliques (where adjacency means sharing k-1 vertices).

Pelc and Peleg [21] defined the parameter $X(G)$ of a connected graph $G = (V, E)$: for every pair of nodes $i, j \in V$, $X(i, j)$ denotes the number of nodes $x \in \Gamma(i)$ [3] that are closer to j than i; the parameter $X(G)$ is defined as the minimum $X(i, j)$ between any pair of not incident nodes, namely $X(G) := \min\{X(i, j) |\ i, j \in V, (i, j) \notin E\}$. The parameter $X(G)$ allows to arrange nodes of a graph G in disjoint level $L_0, L_1, \dots L_j$ $(j \geq 1)$ with respect their distance to any chosen vertex $s \in V$ such that $L_0 = \{s\}$, $L_1 = \Gamma(s)$ and any node in a level L_i is at distance i from s and it has at least $X(G)$ neighbors in L_{i-1} (i.e. a *level ordering* [12]). A graphical example is provided in Figure 1a.

Litsas et al. [12] defined the parameter $\Psi(G)$ of a graph G. Such a parameter allows to arrange nodes of graphs in disjoint level $L_0, L_1, \dots L_j$ $(j \geq 1)$ with respect to any chosen vertex $s \in V$ such that $L_0 = \{s\}$, $L_1 = \Gamma(s)$ and any node in a level L_i has at least $\Psi(G)$ neighbors in levels $[L_1, L_{i-1}]$ (i.e. a *minimum level ordering* [12]). A graphical example is provided in Figure 1b.

We refer with $\langle k, l \rangle$*-multipartite cycle* to a connected graph G composed by l sets of k not adjacent nodes, such that each set is part of exactly two complete bypartite subgraphs of $2k$ nodes. Figure 1c depicts an example of a $\langle 2, 4 \rangle$-multipartite cycle.

All the graph parameters and topologies we recalled guarantee specific graph topological properties that will be leveraged addressing the reliable communication problem.

[3] $\Gamma(s)$ is the set of nodes in the neighborhood of node s in a graph.

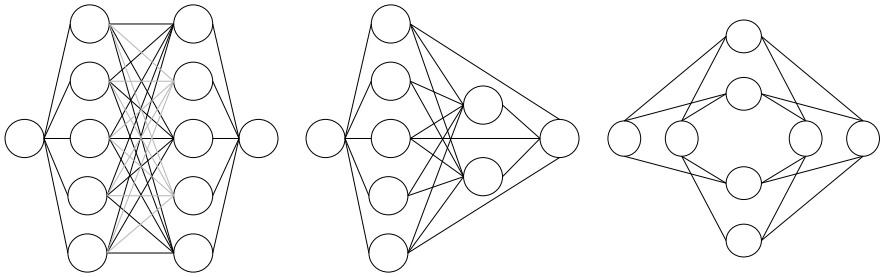

(a) (b) (c)

Fig. 1: (a) Level ordering with $X(G) = 5$. (b) Minimim Level ordering with $\Psi(G) = 5$. (c) $\langle 2, 4 \rangle$-multipartite cycle.

4 Mobile Byzantine reliable communication problem specification

Not all processes in a multi-hop network can directly exchange messages: some of them have to rely on intermediate nodes relaying their messages in order to communicate. Meanwhile, Byzantine faulty processes may diffuse *spurious messages*, i.e. messages that have not been sent by their advertised source. A reliable communication primitive prevent all correct processes from delivering spurious messages while allowing them to communicate.

We aim to define a mobile Byzantine fault tolerant *reliable communication* primitive in a multi-hop network of point-to-point reliable authenticated links, namely to enable message exchanges between every pair of processes extending the guarantees provided by the point-to-point links in a distributed system affected by Mobile Byzantine Faults.

The standard reliable communication (RC) specification [3, 10, 17, 19, 21] between a *source* process p_s and a *target* process p_t requires the following properties to be satisfied: *safety* - if a correct process p_t delivers a message m from p_s, then m has been sent by p_s; *liveness*: if a correct process p_s sends a message m to a correct process p_t, then m is eventually delivered by p_t.

In the system model we are considering, the failure state of processes change over time and no process is permanently correct. Furthermore, processes can be compromised while they are communicating, namely between the computation and send phase. It follows that every process which aims to communicate with a peer must remain correct for at least two consecutive rounds in order to diffuse any message, and thus, we define as *correct source* a process p_s that is correct for two consecutive rounds r_i and r_{i+1}, and computes a message m at r_i.

Another aspect to take into account is that a message may require several rounds to reach a target process, due to the network topology and to the protocol employed to diffuse it. As a matter of fact, the state of a process may change over time and a target process must not be permanently faulty in order to deliver a

message sent by a source. Therefore, we say that a process p_j is *not permanently faulty* if for every round r_i there always exists a round $r' \geq r_i$ where p_j is correct.

Given all considerations stated above, we define a specification for the reliable communication problem with Mobile Byzantine Faults.

Reliable communication with MBF specification. Given a *correct source* process p_s and *target* process p_t, a reliable communication primitive guarantees that:

- *safety* - if p_t is correct at r_i and it delivers a message m from p_s, then m has been sent by p_s;
- *liveness*: if a correct source p_s sends a message m to a not permanently faulty process p_t, then p_t eventually delivers m.

5 Reliable communication in asynchronous systems

In this section, we show that it is impossible to design a protocol $\mathcal{P}$ that is able to solve the reliable communication problem between a correct source p_s and a target p_t when the distributed system is asynchronous and there is only one mobile Byzantine agent. This motivates the subsequent assumptions for analyzing synchronous systems (see section 6).

When assuming a fully asynchronous system, we consider that correct processes still execute a deterministic distributed protocol $\mathcal{P}$, but there is no known upper bound on the time demanded for local computation, neither on the time required to deliver point-to-point messages.

Theorem 1. *There exists no distributed protocol $\mathcal{P}$ that is able to solve the reliable communication problem specification with Mobile Byzantine Faults in an asynchronous system even if* (i) *the source process p_s is permanently correct,* (ii) *there exists only one mobile Byzantine agent, and* (iii) *processes are aware of their failure state.*

Proof. The reliable communication specification requires both safety and liveness property to be satisfied. We show that no protocol $\mathcal{P}$ can ensures the liveness property, even assuming an always correct source, only one mobile Byzantine agent and the aware failure model.

The reliable delivery property enforced by reliable and authenticated links is guaranteed only between correct processes. Given that there is no constraint on the link delay, even assuming a permanently correct source that continuously sends a message m, such a message may never be delivered by the link, because a target process p_t may be compromised during each transmission of m. □

On the other hand, we highlight on the solvability of safe communication (i.e. enforcing only the safety property) in case of an asynchronous system.

The immediate consequence is that in the aware failure model, it is possible to design a "best-effort" protocol that ensures safety while trying to maximize the number of delivered messages.

Theorem 2. *Safe communication can be achieved with a non-degenerated protocol in an asynchronous distributed system in the aware mobile Byzantine failure model.*

Proof. We show a "best-effort" solution for the safe communication problem. Let us assume that every process p_j has access to a local clock T_j. It is reasonable to assume that a Byzantine agent which is forcing a process p_k to send a message m must remain on p_k till the end of its transmission to guarantee the link message delivery. Let us consider the following protocol:

- the source process p_s continuously sends $\langle s, t, m\rangle$;
- every process p_j stores every message $\langle s, t, m\rangle$ received from a process p_k jointly with timestamp $t^k_{\langle s,t,m\rangle}$ containing the value of T_j at the reception of $\langle s, t, m\rangle$;
- every process p_j stores and continuously relays any message $\langle s, t, m\rangle$ received from p_s;
- every process p_j that stores a set of $2f + 1$ tuples $M := [\langle\langle s, t, m\rangle, t^1_{\langle s,t,m\rangle}\rangle, \langle\langle s, t, m\rangle, t^2_{\langle s,t,m\rangle}\rangle, \ldots, \langle\langle s, t, m\rangle, t^{2f+1}_{\langle s,t,m\rangle}\rangle]$ received from distinct neighbors such that $\forall_{i<j}, t^i_{\langle s,t,m\rangle} < t^j_{\langle s,t,m\rangle}$ and $t^{2f+1}_{\langle s,t,m\rangle} - t^1_{\langle s,t,m\rangle} < \rho$ continuously relays $\langle s, t, m\rangle$;
- if process p_t relays $\langle s, t, m\rangle$ then it delivers m.

We show that the protocol defined above guarantees safety of reliable communication in an asynchronous system. Let us consider a single agent initially placed on a process $p_1 \neq p_s$, that starts the transmission of a spurious message $\tilde{m}$ to a process p_q at time t_1^{start} and concludes at time t_1^{end} when $\tilde{m}$ is received by p_q. Process p_q then stores $\tilde{m}$ and a timestamp $t^1_{\tilde{m}}$ obtained by its local clock at the reception of $\tilde{m}$. Subsequently, the Byzantine agent may move on a different process p_2 and start sending another copy of $\tilde{m}$ to p_q at time t_2^{start}, that it concludes at time t_2^{end} when the message is received by p_q. Again, process p_q stores $\tilde{m}$ and a timestamp $t^2_{\tilde{m}}$. And once more, the agent can move another time on a process p_3 and iterate again the transmission of $\tilde{m}$. According with the absence of link latency guarantees, it could happen that $t_i^{end} - t_i^{start} \to 0$. On the other hand, $t^{j+2}_{\tilde{m}}$ - $t^j_{\tilde{m}} > \rho$, because a mobile agent must move twice in order to send a spurious message for three distinct processes. It follows that, assuming f mobile Byzantine agents, if a process q receives more than $2f$ copies of a message m in a time windows shorter than ρ, then it can safely accept m. For ease of explanation, the execution stated above is depicted in Figure 2. □

6 Reliable communication in synchronous systems

In this section, we briefly present the seminal reliable communication protocol defined by Sasaki et al. [24], and we define a new parameterized algorithm, *RCMB*.

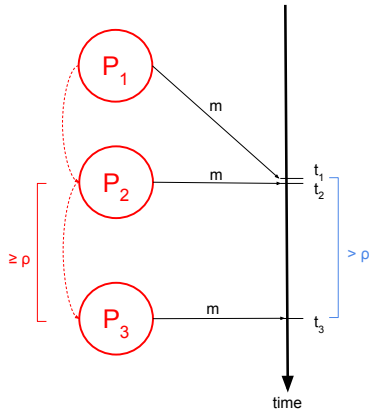

Fig. 2: Graphical execution example of Theorem 2.

Sasaki et al. [24] proposed a reliable communication protocol aimed to enable mobile Byzantine agreement on multi-hop networks. Their solution is based on the fact that mobile Byzantine agents may compromise at most f processes at every round: leveraging the disjoint paths available between all pairs of processes, they defined a reliable communication protocol that enables mobile Byzantine agreement in the unaware failure model in graphs $G(\alpha, \beta)$ where the inequality $\alpha > 2\beta f$ is satisfied. Specifically, messages between every pair of processes are routed over α disjoint paths and Byzantine agents may at most compromise βf of them.

RCMB Algorithm. We define a new protocol addressing the reliable communication problem, *RCMB*. With respect to the one proposed by Sasaki et al., it aims to keep the number of processes that concurrently send spurious messages bounded over time.

Algorithm *Reliable Communication Mobile Byzantines - RCMB*:

- the source process p_s computes message m addressed to a target process p_t at round r_i, and saves $\langle s, t, m\rangle$ in a set variable *delivered*.
- any message $\langle s, t, m\rangle$ stored in *delivered* is removed after τ rounds.
- every process p_j queues every message stored in *delivered* at round r_i to be sent in round r_{i+1} to itself and to all of its neighbors;
- if a correct process p_j receives a message $\langle s, t, m\rangle$ from p_s at round r_i, then p_j saves $\langle s, t, m\rangle$ in a set variable *delivered*, and delivers m from p_s if $j = t$;
- if a correct process p_j receives more than σ copies of a message $\langle s, t, m\rangle$ from distinct neighbors at round r_i, then p_j saves $\langle s, t, m\rangle$ in a set variable *delivered*, and delivers m from p_s if $j = t$;

The parameter σ is a safety threshold, corresponding to the number of copies of the same message that must concurrently be received to deliver it. The parameter τ allows processes that were faulty in the unaware failure model to remove

spurious messages that may have been injected by malicious agents. It can be ignored in the aware failure model, because cured processes directly wipe their local variables. Notice that, in case of $\tau = 1$, every message stored in *delivered* at round r_i is queued to be sent at round r_{i+1} and then dropped.

6.1 Reliable communication correctness conditions

We provide in this section several correctness conditions that enable to solve the reliable communication problem with one of the protocols presented in the previous subsection. We investigate the solvability of reliable communication in two scenarios: a correct source and a permanently correct source (that is, a source that is correct in every round r_i). The latter case is motivated by the fact that such additional assumption enables to solve the reliable communication problem in further topologies.

Unaware failure model

Theorem 3. *Reliable communication cannot be achieved in the unaware mobile Byzantine failure model with $n \leq 4f$.*

Proof. The result can be deduced from the lower bound implementing the safe register abstraction in the unaware mobile Byzantine failure model [4]. Let us consider a set of $4f$ processes connected through a complete communication network. Let us assume a correct source p_s that computes a message m at round r_0, that p_s sends it to all other processes at round r_1 and that p_t and other $f-1$ processes are faulty at r_1. Thus, p_t is faulty while the reliable communication protocol is diffusing m according to a distributed protocol $\mathcal{P}$. Subsequently, the mobile Byzantine agents move on process p_s and on $f-1$ other processes between rounds r_1 and r_2. It follows that at round r_2 there are $2f$ processes that share a state that contains m and $2f$ processes (f Byzantine faulty at r_2 and f that were faulty in r_1) that may share a state injected by the adversary, thus it is not possible to distinguish which set of processes is storing the message sent by the correct source. □

Theorem 4. *The `RCMB` protocol with $\sigma = (\tau + 1)f$ guarantees safety of reliable communication in the unaware mobile Byzantine failure model.*

Proof. Let us consider a set of n process connected through a complete network. Let us assume, for the ease of contradiction, that a target process p_t delivers a message m at round r_i from p_s but m has not been sent by its source (i.e. m is a spurious message).

The delivery of a message m in the *RCMB* protocol is independent from the process local variables and it is only determined by the messages that are currently received in a round. On the other hand, the messages that are diffused in a round depends on the content of the *delivered* variable.

The message m has not been received by a process through a link with the source process p_s according to our hypothesis. It follows that there have been

more than $\sigma = (\tau + 1)f$ processes that sent $\langle s, t, m \rangle$ to p_t at round r_i. Mobile Byzantine agents can force f processes to send $\langle s, t, m \rangle$ at round r_i and they can inject $\langle s, t, m \rangle$ in the *delivered* sent of the processes that were faulty at $r_{i-k}, k \in [1, \tau]$ if $\tau \geq 1$. Thus, at most τf correct processes may potentially relay $\langle s, t, m \rangle$ in a round because they were previously faulty, since after τ rounds $\langle s, t, m \rangle$ is dropped from the *delivered* set. Every other correct processes process $p_j \neq p_t$ sends $\langle s, t, m \rangle$ at r_i only if either p_j received such a message through a link with process p_s, or from more than $(\tau + 1)f$ neighibors. It follows that at most $(\tau + 1)f$ processes in the system may concurrently send $\langle s, t, m \rangle$. Thus message m has been sent by its source. This leads to a contradiction and the claim follows. □

Theorem 5. *The RCMB protocol with $\tau = 1$ and $\sigma = (\tau+1)f$ provides reliable communication in complete networks of size $n > 4f$ in the unaware mobile Byzantine failure model.*

Proof. We verified the safety property of the *RCMB* protocol with $\sigma = (\tau+1)f$ in the unaware mobile Byzantine failure model in Theorem 4. We need to prove the liveness property of reliable communication in a complete networks of size $n > 4f$ considering $\tau = 1$.

Let us assume a correct source p_s that computes a message m at r_0 and sends it at r_1 to itself and to all of its neighbors according to the *RCMB* algorithm. It follows that more than $3f$ processes queue m to be sent at r_2, because m has been received through a link from its source. At any round r_i there are at most f processes that get faulty and at most f ones that were faulty in r_{i-1}. Thus, all correct processes receive at least $2f + 1$ copies of m from distinct nodes at any round $r_j \geq r_2$ and they relay it at the subsequent round. It follows that message m is relayed by at least $2f + 1 > \sigma$ processes on any round $r_j \geq r_2$, and that process p_t delivers it in a round $r_j \geq r_2$ it is correct. □

Theorem 6. *The RCMB protocol with $\tau = 1$ and $\sigma = (\tau + 1)f$ provides reliable communication in the unaware mobile Byzantine failure model in a k-clique community network topology with $k > 4f + 1$.*

Proof. We verified the safety property of the *RCMB* protocol with $\sigma = (\tau+1)f$ in the unaware mobile Byzantine failure model in Theorem 4. We need to prove the liveness property of reliable communication in a k-clique community network topology with $k > 4f + 1$ considering $\tau = 1$.

Let us assume a correct source p_s that computes a message m at round r_0 and sends it at round r_1. Given a k-clique community network, two processes p_s and p_t are either both part of a k-clique or they are are included in two distinct k-cliques that are connected through a sequence of adjacent ones.

Let us assume that p_s and p_t are both part of a k-clique $\mathcal{K}_0$. We showed in Theorem 5 that all correct processes in a complete network of at least $4f + 1$ nodes continuously relay a message m sent by a correct sender. It follows that t delivers m in a round $r_i \geq r_1$ it is correct.

Let us assume that p_t is part of a k-clique $\mathcal{K}_1$ adjacent to $\mathcal{K}_0$. All correct processes but $2f$ in $\mathcal{K}_0$ sends m at every round $r_j \geq r_{i+2}$ to all of their neighbors. It follows that p_t receives at least $2f+1 > \sigma$ copies of m on every round $r_j \geq r_{i+2}$ because it is connected to at least $4f+1$ nodes in $\mathcal{K}_0$. Thus, it delivers m in a round $r_j \geq r_{i+2}$ it is correct.

Such an argumentation extends to any process in a k-clique reachable through a sequence of adjacent k-cliques. □

Theorem 7. *The RCMB protocol with $\tau = 2$ and $\sigma = (\tau+1)f$ provides reliable communication in the unaware mobile Byzantine failure model in a network topology G where $n > 6f$ and $X(G) > 6f$.*

Proof. The condition $X(G) > 6f$ allows to arrange nodes of a graph G in a level ordering of two or more levels $[L_0, \cdots L_k]$. Let us consider a correct source p_s that computes a message m at r_0 and sends it at round r_1.

Let us assume that the level ordering with respect to p_s is composed by 2 levels. It follows that all processes have a link with the source and that all the correct ones receive m at round r_1 directly from the source, thus they save it into their *delivered* set and relay it at r_2. Subsequently, the mobile Byzantine agents can move between r_1 and r_2. At round r_2 all correct processes are connected to at least $4f+1 > \sigma$ processes that relays m. It follows they relay m at round r_3 and at all the subsequent rounds.

Let us assume that the level ordering with respect to p_s is composed by 3 or more levels. At round r_1 all correct processes in L_1 receive m directly from the source, thus they save it into their *delivered* set and they relay it at r_2. Subsequently, the mobile Byzantine agents can move between r_1 and r_2, and at round r_2 all correct processes in L_1 relay m to all nodes in L_2. Every process in L_2 has at least $6f+1$ neighbors in L_1 and at least $4f+1 > \sigma$ of them relay m. It follows they all save and relay m at round r_3. Between rounds r_2 and r_3 the mobile Byzantine agents move and compromise further f processes. It follows that at round r_3 every process in levels L_1, L_2 and L_3 receives m from at least $3f+1 > \sigma$ processes, because each of them has at least $6f+1$ neighbors inside the first three levels and at most $3f$ processes may have been compromised from the beginning of the transmission. It follows that all correct processes in the first three levels relay m at every round $r_i \geq r_4$. This reasoning extends considering more levels. □

Theorem 8. *The RC Sasaki et al. protocol with $\sigma = \beta f$ provides reliable communication in the unaware mobile Byzantine failure model in networks where the inequality $\alpha > 2\beta f$ is satisfied [24].*

Proof. Every reliable communication instance between a source process p_s and a destination process p_t lasts exactly β rounds in the *RC Sasaki et al.* protocol. The inequality $\alpha > 2\beta f$ guarantees that between every pair of processes there exist at least $2\alpha+1$ disjoint paths of length at most β. Any process can relay messages between peers p_s and p_t at only one defined round every β ones. It follows that the mobile Byzantine agents can compromise at most βf processes (and thus

disjoint paths) in β rounds, and thus no correct process receive more than σ copies of a spurious message in a round. The assumption $\alpha > 2\beta f$ guarantees instead that there always exist $\beta f + 1$ disjoint paths that are not compromised by Byzantine agents in every communication instance. □

Theorem 9. *The RCMB protocol with $\tau = 1$ and $\sigma = (\tau + 1)f$ provides reliable communication from a permanently correct source in the unaware mobile Byzantine failure model in networks where $\Psi(G) > 4f$.*

Proof. We verified the safety property of the *RCMB* algorithm with $\sigma = (\tau+1)f$ in the unaware mobile Byzantine failure model in Theorem 4. We need to prove the liveness property of reliable communication in networks where $\Psi(G) > 4f$ in case of a permanently correct source and $\tau = 1$.

The condition $\Psi(G) > 4f$ allows to arrange the nodes of a network G in a $(4f + 1)$-minimum level ordering with respect to every vertex of G.

Let us assume that process p_s sends a message m employing *RCMB* to process p_t at round r_i. Process p_t can either be in L_1 or in $L_{i>1}$. In the former case it receives m through a link from s starting from round $r_j \geq r_{i+1}$ it is correct, and thus it eventually delivers the message m. In the latter case, all correct processes in L_1 receive m at every round $r_j \geq r_{i+1}$. Thus, they queue m to be sent at every round $r_j \geq r_{i+2}$. At every round, there are at most f processes that can be faulty among all levels. It follows that at least $2f + 1$ processes in L_1 relay m to processes in L_2 at every round r_j, because f nodes may have been faulty at r_{j-1} and f ones are faulty at r_j. Therefore, all correct processes in L_2 relays m to all of their neighbors at every round $r_j \geq r_{i+3}$, and if process t is in L_2 then it delivers m at $r_j \geq r_{i+3}$ when it is correct. The reasoning extends to any other level given the assumption of $\Psi(G) > 4f$, and the claim follows. □

Aware failure model

Theorem 10. *Reliable communication cannot be achieved in the aware mobile Byzantine failure model with $n \leq 3f$.*

Proof. The result can be deduced from the lower bound implementing the safe register abstraction in the aware mobile Byzantine failure model [4]. Let us consider a set of $3f$ processes connected through a complete communication network. Let us assume a correct source p_s that computes a message m at round r_0, that p_s sends it to all other processes at round r_1 and that p_t and other $f-1$ processes are faulty at r_1. Thus, p_t is faulty while the reliable communication protocol is diffusing m according to a distributed protocol $\mathcal{P}$. Subsequently, the mobile Byzantine agents moves on process p_s and on $f - 1$ other processes between rounds r_1 and r_2. It follows that at round r_2 there are f processes that share a state that contains m, f cured processes (i.e. with wiped local variables) and f faulty processes. Thus, it is not possible to distinguish which set of processes (the f faulty or the f not cured ones) is storing the message sent by the correct source. □

Theorem 11. *The RCMB protocol with $\sigma = f$ guarantees safety of reliable communication in the aware mobile Byzantine failure model.*

Proof. Let us consider a set of n process connected through a complete network. Let us assume, for the ease of contradiction, that a target process p_t has delivered a message m at round r_i from p_s but m has not been sent by its source (i.e. m is a spurious message).

The delivery of a message m in *RCMB* is independent from the process local variables and it is only determined by the messages that are received in a round. The message m has been received by no process through a link with the source process p_s according to our hypothesis. It follows there have been more than $\sigma = f$ processes that sent $\langle s, t, m\rangle$ to p_t at round r_i. The mobile Byzantine agents can force f processes to send $\langle s, t, m\rangle$ at round r_i. The correct processes at r_i that were faulty at r_{i-1} turn to the cured state, thus they wipe their local variables (and thus their *delivered* set) and remove any message previously queued for the submission. Any correct process $p_j \neq p_t$ sends $\langle s, t, m\rangle$ at r_i only if either p_j has received such a message through a link with process p_s, or from more than f neighbors in a round. It follows that at most f processes in the system may concurrently send $\langle s, t, m\rangle$. Thus message m has been sent by its source. This leads to a contradiction and the claim follows. □

Theorem 12. *The RC Sasaki et al. protocol with $\sigma = (\beta-1)f$ provides reliable communication in the aware mobile Byzantine failure model in networks where the inequality $\alpha > (2\beta - 1)f$ is satisfied.*

Proof. The cured processes remain silent, namely they drop every message previously queued for the submission. In the first round of a reliable communication instance, only the source is allowed to transmit. It follows that no process can diffuse spurious messages in such a round. Therefore, spurious messages can only traverse $(\beta-1)f$ disjoint paths in a communication instance. On the other hand, Byzantine agents can still compromise f processes per round, preventing peers from receiving and relaying messages, and thus up to βf ones may be compromised in every communication instance. The inequality follow considering that $(\beta-1)f+1$ copies of a message received in a single round are sufficient to ensure safety and that at most βf process can be compromised during a communication instance. □

Theorem 13. *The RCMB protocol with $\sigma = f$ provides reliable communication in complete networks of size $n > 3f$ in the aware mobile Byzantine failure model.*

Proof. We verified the safety property of the *RCMB* algorithm with $\sigma = f$ in the aware mobile Byzantine failure model in Theorem 11. We need to prove the liveness property of reliable communication in a complete networks of size $n > 3f$.

Let us assume a correct source p_s that computes a message m at r_0 and sends it at r_1 to itself and to all of its neighbors according to the *RCMB* algorithm. It follows that p_s and at least $2f$ processes queue m to be sent at r_2, because

m has been received through a link from its source. At any round r_i there are at most f processes that are faulty and at most f ones that were faulty in r_{i-1}. Thus, all correct processes receive at least $f + 1 > \sigma$ copies of m from distinct nodes at any round $r_j \geq r_2$ and they relay it in the subsequent round. It follows that message m is relayed by at least $f + 1$ processes at any round $r_j \geq r_2$, and that process p_t delivers it in a round $r_j \geq r_2$ it is correct. □

Theorem 14. *The `RCMB` protocol with $\sigma = f$ provides reliable communication in the aware mobile Byzantine failure model in i) a k-clique community network topology with $k > 3f + 1$ and ii) in topologies where $X(G) > 5f$.*

Proof. We verified the safety property of the *RCMB* algorithm with $\sigma = f$ in the unaware mobile Byzantine failure model in Theorem 11.

The liveness property in case of k-clique community networks with $k > 3f+1$ or networks where $X(G) > 5f$ follows from the same argumentation provided respectivelly in Theorems 6 and 7 considering that σ is reduced to f. □

Theorem 15. *The `RCMB` protocol with $\sigma = f$ provides reliable communication from a permanent correct source in the aware mobile Byzantine failure model in networks where $\Psi(G) > 3f$.*

Proof. We verified the safety property of the *RCMB* algorithm with $\sigma = f$ in Theorem 11.

The liveness property in networks where $\Psi(G) > 3f$ follows from the same argumentation provided in Theorem 9 considering that σ is reduced to f. □

6.2 Graph parameters comparison

In this section we provide some examples of topology where the condition $\alpha > 2\beta f$ by Sasaki et al. [24] is not satisfied, but the reliable communication problem remains solvable.

Theorems 6 and 14 identify k-clique communities as a topology where the reliable communication problem is solvable. There exist topologies where $\alpha \leq 2\beta f$ but $k > 4f + 1$, and an example is depicted in Figure 3a: a 6-clique community graph. According with Theorem 6, it is possible to provide reliable communication tolerating one mobile Byzantine agents ($f = 1$) in such a topology (indeed, $k > 4f + 1 = 5$) considering the unaware failure model with algorithm RCMB. On the other hand, in such a graph $\beta = 3$ and $\alpha = 5$, thus the inequality $\alpha > 2\beta f$ is not satisfied for $f \geq 1$, so the algorithm by Sasaki et al. [24] does not guarantee reliable communication in such a network.

Theorems 9 and 15 identify graphs where the parameter $\Psi(G)$ is greater than certain values as topologies where the reliable communication problem is solvable from a permanent correct source. There exist topologies where $\alpha \leq 2\beta f$ but $\Psi(G) > 4f$, and an example is depicted in Figure 3b. According with Theorem 8, one mobile Byzantine agent ($f = 1$) cannot be tolerated by the algorithm by Sasaki et al [24], indeed $\alpha = 5$ and $\beta = 3$. Instead, $\Psi(G) > 4f$ in such an example,

allowing to achieve reliable communication against one mobile Byzantine agent with algorithm RCMB.

The conditions defined in Theorems 7 and 14 identify new topologies where it is possible to solve the reliable communication problem. Specifically, there exist topologies where $\alpha \leq 2\beta f$ but $X(G) > 6f$. An example is depicted in Figure 3c: a $\langle 7, 14 \rangle$-multipartite cycle. In such a network, $X(G) = 7$, $\alpha = 14$ and $\beta = 7$. According with Theorem 7 it is possible to achieve reliable communication against one mobile Byzantine agents (indeed, $X(G) > 6f$) with Algorithm RCMB. On the other hand the inequality $\alpha > 2\beta f$ is not satisfied in such a topology, so the algorithm by Sasaki et al. [24] cannot guarantee reliable communication in such a setting.

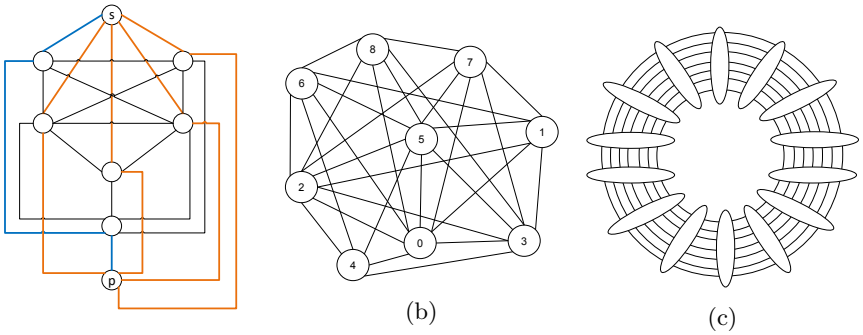

(a) (b) (c)

Fig. 3: (a) 6-clique community example. (b) $\Psi(G) = 5$, $\alpha > 2\beta f$ not satisfied with $f > 1$. (c) $\langle 7, 14 \rangle$-multipartite cycle.

7 Conclusion

We analyzed the reliable communication problem in distributed systems affected by Mobile Byzantine Faults. We highlighted the specific difficulties that arise when considering mobile malicious agents able to move in the system and to continuously compromise nodes. We shown that the reliable communication problem arises even in complete communication networks, and that it is not possible to address it in an asynchronous system. Then, starting from the only solution available in the literature (the one proposed by Sasaki et al. [24]), we provided additional insights about the specific properties that such protocols are able to guarantee. In more details, we defined a new reliable communication protocol, *RCMB*, and we identified new multi-hop topologies where reliable communication primitives remain feasible.

Our work paves the way toward deeper analyzes about reliable communication and others related distributed system problems with mobile Byzantine

faults in multi-hop networks. A particularly interesting question is the feasability of tolerating both mobile Byzantine failures and self-stabilization (as in the register construction of Bonomi et al. [6]) for the purpose of reliable communication. To our knowledge, this problem was only shown solvable by Maurer et al. [14] for the static Byzantine case.

References

1. Bonnet, F., Défago, X., Nguyen, T.D., Potop-Butucaru, M.: Tight bound on mobile byzantine agreement. Theor. Comput. Sci. **609**, 361–373 (2016). https://doi.org/10.1016/j.tcs.2015.10.019
2. Bonomi, S., Farina, G., Tixeuil, S.: Reliable broadcast in dynamic networks with locally bounded byzantine failures. In: Izumi, T., Kuznetsov, P. (eds.) Stabilization, Safety, and Security of Distributed Systems - 20th International Symposium, SSS 2018, Tokyo, Japan, November 4-7, 2018, Proceedings. Lecture Notes in Computer Science, vol. 11201, pp. 170–185. Springer (2018). https://doi.org/10.1007/978-3-030-03232-6_12
3. Bonomi, S., Farina, G., Tixeuil, S.: Multi-hop byzantine reliable broadcast with honest dealer made practical. J. Braz. Comp. Soc. **25**(1), 9:1–9:23 (2019). https://doi.org/10.1186/s13173-019-0090-x
4. Bonomi, S., Pozzo, A.D., Potop-Butucaru, M.: Optimal self-stabilizing synchronous mobile byzantine-tolerant atomic register. Theor. Comput. Sci. **709**, 64–79 (2018). https://doi.org/10.1016/j.tcs.2017.08.020
5. Bonomi, S., Pozzo, A.D., Potop-Butucaru, M., Tixeuil, S.: Optimal storage under unsynchronized mobile byzantine faults. In: 36th IEEE Symposium on Reliable Distributed Systems, SRDS 2017, Hong Kong, Hong Kong, September 26-29, 2017. pp. 154–163. IEEE Computer Society (2017). https://doi.org/10.1109/SRDS.2017.20
6. Bonomi, S., Pozzo, A.D., Potop-Butucaru, M., Tixeuil, S.: Brief announcement: Optimal self-stabilizing mobile byzantine-tolerant regular register with bounded timestamps. In: Izumi, T., Kuznetsov, P. (eds.) Stabilization, Safety, and Security of Distributed Systems - 20th International Symposium, SSS 2018, Tokyo, Japan, November 4-7, 2018, Proceedings. Lecture Notes in Computer Science, vol. 11201, pp. 398–403. Springer (2018). https://doi.org/10.1007/978-3-030-03232-6_28
7. Bonomi, S., Pozzo, A.D., Potop-Butucaru, M., Tixeuil, S.: Approximate agreement under mobile byzantine faults. Theor. Comput. Sci. **758**, 17–29 (2019). https://doi.org/10.1016/j.tcs.2018.08.001
8. Buhrman, H., Garay, J.A., Hoepman, J.: Optimal resiliency against mobile faults. In: Digest of Papers: FTCS-25, The Twenty-Fifth International Symposium on Fault-Tolerant Computing, Pasadena, California, USA, June 27-30, 1995. pp. 83–88. IEEE Computer Society (1995). https://doi.org/10.1109/FTCS.1995.466995
9. Cachin, C., Guerraoui, R., Rodrigues, L.E.T.: Introduction to Reliable and Secure Distributed Programming (2. ed.). Springer (2011). https://doi.org/10.1007/978-3-642-15260-3
10. Dolev, D.: Unanimity in an unknown and unreliable environment. In: 22nd Annual Symposium on Foundations of Computer Science, Nashville, Tennessee, USA, 28-30 October 1981. pp. 159–168 (1981). https://doi.org/10.1109/SFCS.1981.53
11. Garay, J.A.: Reaching (and maintaining) agreement in the presence of mobile faults (extended abstract). In: Tel, G., Vitányi, P.M.B. (eds.) Distributed Algorithms, 8th International Workshop, WDAG '94, Terschelling, The Netherlands, September 29

- October 1, 1994, Proceedings. Lecture Notes in Computer Science, vol. 857, pp. 253–264. Springer (1994). https://doi.org/10.1007/BFb0020438
12. Litsas, C., Pagourtzis, A., Sakavalas, D.: A graph parameter that matches the resilience of the certified propagation algorithm. In: Cichon, J., Gebala, M., Klonowski, M. (eds.) Ad-hoc, Mobile, and Wireless Network - 12th International Conference, ADHOC-NOW 2013, Wrocław, Poland, July 8-10, 2013. Proceedings. Lecture Notes in Computer Science, vol. 7960, pp. 269–280. Springer (2013). https://doi.org/10.1007/978-3-642-39247-4_23
13. Maurer, A., Tixeuil, S.: Byzantine broadcast with fixed disjoint paths. J. Parallel Distrib. Comput. **74**(11), 3153–3160 (2014). https://doi.org/10.1016/j.jpdc.2014.07.010
14. Maurer, A., Tixeuil, S.: Self-stabilizing byzantine broadcast. In: 33rd IEEE International Symposium on Reliable Distributed Systems, SRDS 2014, Nara, Japan, October 6-9, 2014. pp. 152–160. IEEE Computer Society (2014). https://doi.org/10.1109/SRDS.2014.10
15. Maurer, A., Tixeuil, S.: Containing byzantine failures with control zones. IEEE Trans. Parallel Distrib. Syst. **26**(2), 362–370 (2015). https://doi.org/10.1109/TPDS.2014.2308190
16. Maurer, A., Tixeuil, S.: Tolerating random byzantine failures in an unbounded network. Parallel Processing Letters **26**(1), 1650003:1–1650003:12 (2016). https://doi.org/10.1142/S0129626416500031
17. Maurer, A., Tixeuil, S., Défago, X.: Communicating reliably in multihop dynamic networks despite byzantine failures. In: 34th IEEE Symposium on Reliable Distributed Systems, SRDS 2015, Montreal, QC, Canada, September 28 - October 1, 2015. pp. 238–245. IEEE Computer Society (2015). https://doi.org/10.1109/SRDS.2015.10
18. Ostrovsky, R., Yung, M.: How to withstand mobile virus attacks (extended abstract). In: Logrippo, L. (ed.) Proceedings of the Tenth Annual ACM Symposium on Principles of Distributed Computing, Montreal, Quebec, Canada, August 19-21, 1991. pp. 51–59. ACM (1991). https://doi.org/10.1145/112600.112605
19. Pagourtzis, A., Panagiotakos, G., Sakavalas, D.: Reliable broadcast with respect to topology knowledge. Distributed Computing **30**(2), 87–102 (2017). https://doi.org/10.1007/s00446-016-0279-6
20. Pelc, A.: Reliable communication in networks with byzantine link failures. Networks **22**(5), 441–459 (1992). https://doi.org/10.1002/net.3230220503
21. Pelc, A., Peleg, D.: Broadcasting with locally bounded byzantine faults. Inf. Process. Lett. **93**(3), 109–115 (2005). https://doi.org/10.1016/j.ipl.2004.10.007
22. Reischuk, R.: A new solution for the byzantine generals problem. Information and Control **64**(1-3), 23–42 (1985). https://doi.org/10.1016/S0019-9958(85)80042-5
23. Sakavalas, D., Tseng, L.: Delivery delay and mobile faults. In: 17th IEEE International Symposium on Network Computing and Applications, NCA 2018, Cambridge, MA, USA, November 1-3, 2018. pp. 1–8. IEEE (2018). https://doi.org/10.1109/NCA.2018.8548345
24. Sasaki, T., Yamauchi, Y., Kijima, S., Yamashita, M.: Mobile byzantine agreement on arbitrary network. In: Principles of Distributed Systems - 17th International Conference, OPODIS 2013, Nice, France, December 16-18, 2013. Proceedings. pp. 236–250 (2013). https://doi.org/10.1007/978-3-319-03850-6_17
25. 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). https://doi.org/10.1016/j.ipl.2014.11.010