

Self-Stabilizing Byzantine Resilient Topology Discovery and Message Delivery

Shlomi Dolev ^{*}, Omri Liba ^{*}, and Elad M. Schiller ^{**}

¹ Department of Computer Science, Ben-Gurion University of the Negev, Beer-Sheva, Israel.
{dolev, liba}@cs.bgu.ac.il

² Department of Computer Science and Engineering, Chalmers University of Technology,
Goeteborg, Sweden. elad@chalmers.se

Abstract. Traditional Byzantine resilient algorithms use $2f + 1$ vertex-disjoint paths to ensure message delivery in the presence of up to f Byzantine nodes. The question of how these paths are identified is related to the fundamental problem of topology discovery. Distributed algorithms for topology discovery cope with a never ending task: dealing with frequent changes in the network topology and unpredictable transient faults. Therefore, algorithms for topology discovery should be self-stabilizing to ensure convergence of the topology information following any such unpredictable sequence of events. We present the first such algorithm that can cope with Byzantine nodes. Starting in an arbitrary global state, and in the presence of f Byzantine nodes, each node is eventually aware of all the other non-Byzantine nodes and their connecting communication links. Using the topology information, nodes can, for example, route messages across the network and deliver messages from one end user to another. We present the first deterministic, cryptographic-assumptions-free, self-stabilizing, Byzantine-resilient algorithms for network topology discovery and end-to-end message delivery. We also consider the task of r -neighborhood discovery for the case in which r and the degree of nodes are bounded by constants. The use of r -neighborhood discovery facilitates polynomial time, communication and space solutions for the above tasks. The obtained algorithms can be used to authenticate parties, in particular during the establishment of private secrets, thus forming public key schemes that are resistant to man-in-the-middle attacks of the compromised Byzantine nodes. A polynomial and efficient end-to-end algorithm that is based on the established private secrets can be employed in between periodical secret re-establishments.

1 Introduction

Self-stabilizing Byzantine resilient topology discovery is a fundamental distributed task that enables communication among parties in the network even if some of the components are compromised by an adversary. Currently, such topology discovery is becoming extremely important where countries' main infrastructures, such as the electrical

^{*} Partially supported by Deutsche Telekom, Rita Altura Trust Chair in Computer Sciences, Lynne and William Frankel Center for Computer Sciences, Israel Science Foundation (grant number 428/11) and Cabarnit Cyber Security MAGNET Consortium.

^{**} Partially supported by the EC, through project FP7-STREP-288195, KARYON (Kernel-based ARchitecture for safetY-critical cONtrol) and the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 257007.

smart-grid, water supply networks and intelligent transportation systems are subject to cyber-attacks. Self-stabilizing Byzantine resilient algorithms naturally cope with mobile attacks [e.g., ?]. Whenever the set of compromised components is fixed (or dynamic, but small) during a period that suffices for convergence of the algorithm, the system starts demonstrating useful behavior following the convergence. For example, consider the case in which nodes of the smart-grid are constantly compromised by an adversary while local recovery techniques, such as local node reset and/or refresh, ensure the recovery of a compromised node after a bounded time. Once the current compromised set does not imply a partition of the communication graph, the distributed control of the smart grid automatically recovers. Self-stabilizing Byzantine resilient algorithms for topology discovery and message delivery are important for systems that have to cope with unanticipated transient violations of the assumptions that the algorithms are based upon, such as unanticipated violation of the upper number of compromised nodes and unanticipated transmission interferences that is beyond the error correction code capabilities.

The dynamic and difficult-to-predict nature of electrical smart-grid and intelligent transportation systems give rise to many fault-tolerance issues and require efficient solutions. Such networks are subject to transient faults due to hardware/software temporal malfunctions or short-lived violations of the assumed settings for the location and state of their nodes. Fault-tolerant systems that are *self-stabilizing* [?] can recover after the occurrence of transient faults, which can drive the system to an arbitrary system state. The system designers consider *all* configurations as possible configurations from which the system is started. The self-stabilization design criteria liberate the system designer from dealing with specific fault scenarios, risking neglecting some scenarios, and having to address each fault scenario separately.

We also consider Byzantine faults that address the possibility of a node to be compromised by an adversary and/or to run a corrupted program, rather than merely assuming that they start in an arbitrary local state. Byzantine components may behave arbitrarily (selfishly, or even maliciously) as message senders and as relaying nodes. E.g., Byzantine nodes may block messages, selectively omit messages, redirect message routes, playback messages, or modify messages. Any system behavior is possible, when all (or one third or more of) the nodes are Byzantine nodes. Thus, the number of Byzantine nodes, f , is usually restricted to be less than one third of the nodes [??].

The task of *r-neighborhood network discovery* allows each node to know the set of nodes that are at most r hops away from it in the communication network. Moreover, the task provides information about the communication links attached to these nodes. The task *topology discovery* considers knowledge regarding the node's entire connected component. The *r-neighborhood network discovery* and *network topology discovery* tasks are identical when r is the communication graph radius.

This work presents the first deterministic self-stabilizing algorithms for *r-neighborhood discovery* in the presence of Byzantine nodes. We assume that every *r-neighborhood* cannot be partitioned by the Byzantine nodes. In particular, we assume the existence of at least $2f + 1$ vertex-disjoint paths in the *r-neighborhood*, between any two non-Byzantine nodes, where at most f Byzantine nodes are present in the

r -neighborhood, rather than in the entire network.³ Note that by the self-stabilizing nature of our algorithms, recovery is guaranteed after a temporal violation of the above assumption. When r is defined to be the communication graph radius, our assumptions are equivalent to the standard assumption for Byzantine agreement in general (rather than only complete) communication graphs. In particular the standard assumption is that $2f + 1$ vertex disjoint paths exist and *are known* (see e.g., [?]) while we present distributed algorithms to find these paths starting in an arbitrary state.

Related work. Self-stabilizing algorithms for finding vertex-disjoint paths for at most two paths between any pair of nodes, and for all vertex-disjoint paths in anonymous mesh networks appear in [?] and in [?], respectively. We propose self-stabilizing Byzantine resilient procedures for finding $f+1$ vertex-disjoint paths in $2f+1$ -connected graphs. In [?], the authors study the problem of spanning tree construction in the presence of Byzantine nodes. Nesterenko and Tixeuil [?] presented a *non-stabilizing* and inconsistent algorithm for topology discovery in the presence of Byzantine nodes – see the paper’s errata for further details about the algorithm’s flaws.⁴ Awerbuch and Sipser [?] consider algorithms that were designed for synchronous static network and give topology update as an example. They show how to use such algorithms in asynchronous dynamic networks. Unfortunately, their scheme starts from a consistent state and cannot cope with transient faults or Byzantine nodes.

The problems of *Byzantine gossip* [? ? ? ? ?] and *Byzantine Broadcast* [? ?] consider the dissemination of information in the presence of Byzantine nodes rather than self-stabilizing topology discovery. Non-self-stabilizing Byzantine resilient gossip in the presence of one selfish node is considered in [? ?]. In [?] the authors study oblivious deterministic gossip algorithms for multi-channel radio networks with a malicious adversary. They assume that the adversary can disrupt one channel per round, preventing communication on that channel. In [?] the authors consider probabilistic gossip mechanisms for reducing the redundant transmissions of flooding algorithms. They present several protocols that exploit local connectivity to adaptively correct propagation failures and protect against Byzantine attacks. Probabilistic gossip mechanisms in the context of recommendations and social networks are considered in [?]. In [?] the authors consider rules for avoiding a combinatorial explosion in (non-self-stabilizing) gossip protocol. Note that deterministic and self-stabilizing solutions are not presented in [? ? ? ? ?]. Drabkin et al. [?] consider non-self-stabilizing broadcast protocols that overcome Byzantine failures by using digital signatures, message signature gossiping, and failure detectors. Our deterministic self-stabilizing algorithm merely use the topological properties of the communication graph to ensure correct message delivery to the application layer in the presence of message omission, modifications and Byzantine nodes. A non-self-stabilizing broadcasting algorithm is considered in [?]. The authors assume the restricted case in which links and nodes of a communication

³ Section 4 considers cases in which r and an upper bound on the node degree, Δ , are constants. For these cases, we have $\mathcal{O}(n)$ disjoint r -neighborhoods. Each of these (disjoint) r -neighborhoods may have up to f Byzantine nodes, and yet the above assumptions about at least $2f + 1$ vertex-disjoint paths in the r -neighborhood, hold.

⁴ <http://vega.cs.kent.edu/~mikhail/Research/topology.errata.html>

network are subject to Byzantine failures, and that faults are distributed randomly and independently.

Our contribution. We present two cryptographic-assumptions-free yet secure algorithms that are deterministic, self-stabilizing and Byzantine resilient.

We start by showing the existence of deterministic, self-stabilizing, Byzantine resilient algorithms for network topology discovery and end-to-end message delivery. The algorithms convergence time is in $\mathcal{O}(n)$. They take in to account every possible path and requiring bounded (yet exponential) memory and bounded (yet exponential) communication costs. Therefore, we also consider the task of r -neighborhood discovery, where r is a constant. We assume that if the r -neighborhood of a node has f Byzantine nodes, there are $2f+1$ vertex independent paths between the node and any non-Byzantine node in its r -neighborhood. The obtained r -neighborhood discovery algorithm requires polynomial memory and communication costs and supports deterministic, self-stabilizing, Byzantine-resilient algorithm for end-to-end message delivery across the network. Unlike topology update, the proposed end-to-end message delivery algorithm establishes message exchange synchronization between end-users that is based on message reception acknowledgments.

Detailed proofs appear in the Appendix.

2 Preliminaries

We consider settings of a standard asynchronous system [cf. ?]. The system consists of a set, $N = \{p_i\}$, of communicating entities, chosen from a set, P , which we call *nodes*. The upper bound on the number of nodes in the system is $n = |P|$. Each node has a unique identifier. Sometime we refer to a set, $P \setminus N$, of nonexisting nodes that a false indication on their existence can be recorded in the system. A node p_i can directly communicate with its *neighbors*, $N_i \subseteq N$. The system can be represented by an undirected network of directly communicating nodes, $G = (N, E)$, named the *communication graph*, where $E = \{(p_i, p_j) \in N \times N : p_j \in N_i\}$. We denote N_k 's set of indices by $indices(N_k) = \{m : p_m \in N_k\}$ and N_k 's set of edges by $edges(N_j) = \{p_j\} \times N_j$.

The r -neighborhood of a node $p_i \in N$ is the connected component that includes p_i and all nodes that can be reached from p_i by a path of length r or less. The problem of r -neighborhood topology discovery considers communication graphs in which p_i 's degree, δ_i , is bounded by a constant Δ . Hence, when both the neighborhood radius, r , and the node degree, Δ , are constants the number of nodes in the r -neighborhood is also bounded by a constant, namely by $\mathcal{O}(\Delta^{r+1})$.

We model the communication channel, $queue_{i,j}$, from node p_i to node $p_j \in N_i$ as a FIFO queuing list of the messages that p_i has sent to p_j and p_j is about to receive. When p_i sends message m , the operation **send** inserts a copy of m to the queue $queue_{i,j}$ of the one destination p_j , such that $p_j \in N_i$. We assume that the number of messages in transit, i.e., stored in $queue_{i,j}$, is at most *capacity*. Once m arrives, p_j executes **receive** and m is dequeued.

We assume that p_i is completely aware of N_i , as in [?]. In particular, we assume that the identity of the sending node is known to the receiving one. In the context of the

studied problem, we say that node $p_i \in N$ is *correct* if it reports on its genuine neighborhood, N_i . A *Byzantine* node, $p_b \in N$, is a node that can send arbitrarily corrupted messages. Byzantine nodes can introduce new messages and modify or omit messages that pass through them. This way they can, e.g., disinform correct nodes about their neighborhoods, about the neighborhood of other correct nodes, or the path through which messages travel, to name a very few specific misleading actions that Byzantine nodes may exhibit. Note that our assumptions do not restrict system settings in which a *duplicitious Byzantine* node, p_b , reports about N_b differently to its correct neighbors. In particular, p_b can have $\{N_{b_1}, \dots, N_{b_{\delta_b}}\}$ reports, such that p_b 's identity in N_{b_i} is different than the one in N_{b_j} , where δ_x is the degree of node p_x . One may use a set of non-duplicitious Byzantine nodes, $\{p_{b_1}, \dots, p_{b_{\delta}}\}$, to model each of p_b 's reports. Thus, for a $2k + 1$ connected graph, the system tolerates no more than $\lfloor k/\Delta \rfloor$ duplicitious Byzantine nodes, where Δ is an upper bound on the node degree.

We denote C and B to be, respectively, the set of correct and Byzantine nodes. We assume that $|B| = f$, the identity of B 's nodes is unknown to the ones in C , and B is fixed throughout the considered execution segment. These execution segments are long enough for convergence and then for obtaining sufficient useful work. We assume that between any pair of correct nodes there are at least $2f + 1$ vertex-disjoint paths. We denote by $G_c = (C, E \cap C \times C)$ the *correct graph* induced by the set of correct nodes.

Self-stabilizing algorithms never terminate [?]. The non-termination property can be easily identified in the code of a self-stabilizing algorithm: the code is usually a do forever loop that contains communication operations with the neighbors. An iteration is said to be complete if it starts in the loop's first line and ends at the last (regardless of whether it enters branches).

Every node, p_i , executes a program that is a sequence of (*atomic*) *steps*. For ease of description, we assume the interleaving model with atomic step execution; a single step at any given time. An input event can either be the receipt of a message or a periodic timer going off triggering p_i to *send*. Note that the system is totally asynchronous and the (non-fixed) node processing rates are irrelevant to the correctness proof.

The *state* s_i of a node p_i consists of the value of all the variables of the node (including the set of all incoming communication channels, $\{queue_{j,i} | p_j \in N_i\}$). The execution of a step in the algorithm can change the state of a node. The term (*system*) *configuration* is used for a tuple of the form (s_1, s_2, \dots, s_n) , where each s_i is the state of node p_i (including messages in transit for p_i). We define an *execution* $E = c[0], a[0], c[1], a[1], \dots$ as an alternating sequence of system configurations $c[x]$ and steps $a[x]$, such that each configuration $c[x + 1]$ (except the initial configuration $c[0]$) is obtained from the preceding configuration $c[x]$ by the execution of the step $a[x]$. We often associate the notation of a step with its executing node p_i using a subscript, e.g., a_i . An execution R (run) is *fair* if every correct node, $p_i \in C$, executes a step infinitely often in R . Time (e.g. needed for convergence) is measured by the number of *asynchronous rounds*, where the first asynchronous round is the minimal prefix of the execution in which every node takes at least one step. The second asynchronous round is the first asynchronous round in the suffix of the run that follows the first asynchronous round, and so on. The message complexity (e.g. needed for convergence) is the number of messages measured in the specific case of synchronous execution.

We define the system's task by a set of executions called *legal executions* (LE) in which the task's requirements hold. A configuration c is a *safe configuration* for an algorithm and the task of LE provided that any execution that starts in c is a legal execution (belongs to LE). An algorithm is *self-stabilizing* with relation to the task LE when every infinite execution of the algorithm reaches a safe configuration with relation to the algorithm and the task.

3 Topology Discovery

The algorithm learns about the neighborhoods that the nodes report. Each report message contains an ordered list of nodes it passed so far, starting in a source node. These lists are used for verifying that the reports are sent over $f + 1$ vertex-disjoint paths.

When a report message, m , arrives to p_i , it inserts m to the queue $informedTopology_i$, and tests the queue consistency until there is enough independent evidence to support the report. The consistency test of p_i iterates over each node p_k such that, p_k appears in at least one of the messages stored in $informedTopology_i$. For each such node p_k , node p_i checks whether there are at least $f + 1$ messages from the same source node that have mutually vertex-disjoint paths and report on the same neighborhood. The neighborhood of each such p_k , that has at least $f + 1$ vertex-disjoint paths with identical neighborhood, is stored in the array $Result_i[k]$ and the total number of paths that relayed this neighborhood is kept in $Count[k]$.

We note that there may still be nodes $p_{fake} \in P \setminus (N)$, for which there is an entry $Result[fake]$. For example, $informedTopology$ may contain f messages, all originated from different Byzantine nodes, and a message m' that appears in the initial configuration and supports the (false) neighborhood the Byzantine messages refer to. These $f + 1$ messages can contain mutually vertex-disjoint paths, and thus during the consistency test, a result will be found for $Result[fake]$. We show that during the next computations, the message m' will be identified and ignored. The $Result$ array should include two reports for each (undirected) edge; the two nodes that are attached to the edge, each send a report. Hence, $Result$ includes a set of directed (report) edges. The term *contradicting edge* is needed when examining the $Result$ set consistency.

Definition 1 (Contradicting edges). Given two nodes, $p_i, p_j \in P$, we say that the edge (p_i, p_j) is contradicting with the set $evidence \subseteq edges(N_j)$, if $(p_i, p_j) \notin evidence$.

Following the consistency test, p_i examines the $Result$ array for contradictions. Node p_i checks the path of each message $m \in informedTopology_i$ with source p_r , neighborhood $neighborhood_r$ and $Path_r$. If every edge (p_s, p_j) on the path appears in $Result[s]$ and $Result[j]$, then we move to the next message. Otherwise, we found a fake supporter, and therefore we reduce $Count[r]$ by one. If the resulting $Count[r]$ is smaller than $f + 1$, we nullify the r 'th entry of the $Result$ array. Once all messages are processed, the $Result$ array consisting of the (confirmed) local topologies is the output. At the end, p_i forwards the arriving message, m , to each neighbor that does not appear in the path of m . The message sent by p_i includes the node from which m arrived as part of the path m .

The pseudocode of Algorithm 1 In every iteration of the infinite loop, p_i starts to compute its preliminary topology view by calling *ComputeResults* in line 2. Then,

every node p_k in the queue *InformedTopology*, node p_i goes over the messages in the queue from head to bottom. While iterating the queue, for every message m with source p_k , neighborhood N_k and visited path $Path_k$, p_i inserts $Path_k$ to $opinion[N_k]$, see line 18. After inserting, p_i checks if there is a neighborhood $Neig_k$ for which $opinion[Neig_k]$ contains at least $f + 1$ vertex-disjoint paths, see line 19. When such a neighborhood is found, it is stored in the *Result* array (line 19). In line 20, p_i stores the number of vertex disjoint paths relayed messages that contained the selected neighborhood for p_k . After computing an initial view of the topology, in line 3, p_i removes non-existing nodes from the computed topology. For every message m in *InformedTopology*, node p_i aims at validating its visited path. In line 24, p_i checks if there exists a node p_k whose neighborhood contradicts the visited path of m . If such a node exists, p_i decreases the associated entry in the *Count* array (line 25). This decrease may cause $Count[r]$ to be smaller than $f + 1$, in this case p_i considers p_k to be fake and deletes the local topology of p_k from *Result[r]* (line 26).

Upon receiving a message m , node p_i inserts the message to the queue, in case it does not already exist, and just moves it to the top of the queue in case it does. The node p_i now needs to relay the message p_i got to all neighbors that are not on the message visited path (line 9). When sending, p_i also attaches the identifier of the node, from which the message was received, to the visited path of the message.

- *Insert(m)*: Insert item m to the queue head.
- *Remove(Message)*: Remove item m from the queue.
- *Iterator()*: Returns an pointer for iterating over the queue's items by their residence order in the queue.
- *HasNext()*: Tests whether the Iterator is at the queue end.
- *Next()*: Returns the next element to iterate over.
- *SizeOf()*: Returns the number of elements in the calling set.
- *MoveToHead(m)*: Move item m to the queue head.
- *IsAfter(m, S)*: Test that item m is after the items $m' \in S$, where S is the queue item set.

Fig. 1. *Queue*: general purpose data structure for queuing items, and its operation list.

Algorithm's correctness proof. We now prove that within a linear amount of asynchronous rounds, the system stabilizes and every output is legal. The proof considers an arbitrary starting configuration with arbitrary messages in transit that could be actually in the communication channel or already stored in p_j 's message queue and will be forwarded in the next steps of p_j . Each message in transit that traverse correct nodes can be forwarded within less than $\mathcal{O}(|C|)$ asynchronous rounds. Note that any message that traverses Byzantine nodes and arrives to a correct node that has at least one Byzantine node in its path. The reason is that the correct neighbor to the last Byzantine in the path lists the Byzantine node when forwarding the message. Thus, f is at most the number of messages that encode vertex-disjoint paths from a certain source that are initiated or corrupted by a Byzantine node. Since there are at least $f + 1$ vertex-disjoint paths with no Byzantine nodes from any source p_k to any node p_i and since p_k repeatedly sends messages to all nodes on all possible paths, p_i receives at least $f + 1$ messages from p_k with vertex-disjoint paths.

The FIFO queue usage and the repeated send operations of p_k ensure that the most recent $f + 1$ messages with vertex-disjoint paths in *InformedTopology* queue are uncorrupted messages. Namely, misleading messages that were present in the initial configuration will be pushed to appear below the new $f + 1$ uncorrupted messages. Thus, each node p_i eventually has the local topology of each correct node (stored in

Algorithm 1: Topology discovery (code for node p_i)

Input: $Neighborhood_i$: The ids of the nodes with which node p_i can communicate directly;
Output: $ConfirmedTopology \subset P \times P$: Discovered topology, which is represent by a directed edge set;
Variable $InformedTopology$: *Queue*, see Figure 1: topological messages,
 $\langle node, neighborhood, path \rangle$;
Function: $NodeDisjointPaths(S)$: Test $S = \{\langle node, neighborhood, path \rangle\}$ to encode at least $f + 1$ vertex-disjoint paths;
Function: $PathContradictsNeighborhood(k, Neighborhood_k, path)$: Test that there is no node $p_j \in N$ for which there is an edge (p_k, p_j) in the message's visited path, $path \subseteq P \times N$, such that (p_k, p_j) is contradicting with $Neighborhood_k$;

```

1 while true do
2   Result  $\leftarrow$  ComputeResults()
3   let Result  $\leftarrow$  RemoveContradictions(Result)
4   RemoveGarbage(Result)
5   ConfirmedTopology  $\leftarrow$  ConfirmedTopology  $\cup (\bigcup_{p_k \in P} Result[k])$ 
6   foreach  $p_k \in N_i$  do send( $i, Neighborhood_i, \emptyset$ ) to  $p_k$ 
7 Upon Receive ( $\langle \ell, Neighborhood_\ell, VisitedPath_\ell \rangle$ ) from  $p_j$ ;
  begin
8   Insert( $p_\ell, Neighborhood_\ell, VisitedPath_\ell \cup \{j\}$ )
9   foreach  $p_k \in N_i$  do if  $k \notin VisitedPath_\ell$  then send( $p_\ell, Neighborhood_\ell, VisitedPath_\ell \cup \{j\}$ )
    to  $p_k$ 
10 Procedure: Insert( $k, Neighborhood_k, VisitedPath_k$ );
    begin
11     if  $\langle k, Neighborhood_k, VisitedPath_k \rangle \in InformedTopology$  then
12       InformedTopology.MoveToHead( $m$ )
13     else if  $p_k \in N \wedge Neighborhood_k \subseteq indices(N) \wedge VisitedPath_k \subseteq indices(N)$  then
14       InformedTopology.Insert( $\langle k, Neighborhood_k, VisitedPath_k \rangle$ )
15 Function: ComputeResults();
    begin
16     foreach  $p_k \in P : \langle k, Neighborhood_k, VisitedPath_k \rangle \in InformedTopology$  do
17       let ( $FirstDisjointPathsFound, Message, opinion[]$ )  $\leftarrow$ 
18         ( $false, InformedTopology.Iterator(), [\emptyset]$ )
19       while Message.hasNext() do
20          $\langle \ell, Neighborhood_\ell, VisitedPath_\ell \rangle \leftarrow Message.Next()$ 
21         if  $\ell = k$  then  $opinion[Neighborhood_\ell].Insert(\langle \ell, Neighborhood_\ell, VisitedPath_\ell \rangle)$ 
22         if  $FirstDisjointPathsFound = false \wedge$ 
23            $NodeDisjointPaths(opinion[Neighborhood_\ell])$  then
24           ( $Result[k], FirstDisjointPathsFound$ )  $\leftarrow$  ( $Neighborhood_\ell, true$ )
25       Count[k]  $\leftarrow$   $opinion[k][Result[k].SizeOf()]$ 
26     return Result
27 Function: RemoveContradictions(Result);
    begin
28     foreach  $\langle r, Neighborhood_r, VisitedPath_r \rangle \in InformedTopology$  do
29       if  $\exists p_k \in P : PathContradictsNeighborhood(p_k, Result[k], VisitedPath_r) = true$ 
30       then
31         if  $Neighborhood_r = Result[r]$  then Count[r]  $\leftarrow$  Count[r] - 1
32         if Count[r]  $\leq f$  then Result[r]  $\leftarrow \emptyset$ 
33     return Result
34 Procedure: RemoveGarbage(Result);
    begin
35     foreach  $p_k \in N$  do
36       foreach  $m = \langle k, Neighborhood_k, VisitedPath_k \rangle \in InformedTopology :$ 
37          $\{k\} \cup Neighborhood_k \cup VisitedPath_k \not\subseteq P \vee InformedTopology.IsAfter(m,$ 
38          $opinion[k][Result[k]])$  do InformedTopology.Remove( $m$ )

```

the $Result_i$ array). The opposite is however not correct as local topologies of non-existing nodes may still appear in the result array. For example, $InformedTopology_i$

may include in the first configuration $f + 1$ messages with vertex-disjoint paths for a non-existing node. Since after *ComputeResults* we know the correct neighborhood of each correct node p_k , we may try to ensure the validity of all messages. For every message that encodes a non-existing source node, there must be a node p_ℓ on the message path, such that p_ℓ is correct and p_ℓ 's neighbor is non-existing, this is true since p_i itself is correct. Thus, we may identify these messages and ignore them. Furthermore, no valid messages are ignored because of this validity check.

We also note that, since we assume that the nodes of the system are a subset of P , the size of the queue *InformedTopology* is bounded. Lemma 1 bounds the needed amount of node memory (the proof details appear in Section A of the Appendix).

Lemma 1 (Bounded memory). *At any time, there are at most $n \cdot 2^{2n}$ messages in *InformedTopology*, where $p_i \in C$, $n = |P|$ and $\mathcal{O}(n \log(n))$ is the message size.*

r -neighborhood discovery. Algorithm 1 demonstrates the existence of a deterministic self-stabilizing Byzantine resilient algorithm for topology discovery. Lemma 1 shows that the memory costs are high when the entire system topology is to be discovered. We note that one may consider the task of r -neighborhood discovery. Recall that in the r -neighborhood discovery task, it is assumed that every r -neighborhood cannot be partitioned by Byzantine nodes. Therefore, it is sufficient to constrain the maximal path length in line 9. The correctness proof of the algorithm for the r -neighborhood discovery follows similar arguments to the correctness proof of Algorithm 1.

4 End-to-End Delivery

We present a design for a self-stabilizing Byzantine resilient algorithm for the transport layer protocol that uses the output of Algorithm 1. The design is based on a function (named *getDisjointPaths()*) for selecting vertex-disjoint paths that contain a set of $f + 1$ correct vertex-disjoint paths. We use *getDisjointPaths()* and ARQ (Automatic Repeat reQuest) techniques for designing Algorithm 2, which ensures safe delivery between sender and receiver.

Exchanging messages over $f + 1$ correct vertex-disjoint paths We guarantee correct message exchange by sending messages over a polynomial number of vertex-disjoint paths between the sender and the receiver. We consider a set, *CorrectPaths*, that includes $f + 1$ correct vertex-disjoint paths. Suppose that *ConfirmedTopology* (see the output of Algorithm 1) encodes a set, *Paths*, of $2f + 1$ vertex-disjoint paths between the sender and the receiver. It can be shown that *Paths* includes at most f incorrect paths that each contain at least one Byzantine node, i.e., $Paths \supseteq CorrectPaths$. As we see next, *ConfirmedTopology* does not always encode *Paths*, thus, one needs to circumvent this difficulty.

Note that even though $2f + 1$ vertex-disjoint paths between the sender and the receiver are present in the communication graph, the discovered topology in *ConfirmedTopology* may not encode the set *Paths*, because f of the paths in the set *Paths* can be controlled by Byzantine nodes. Namely, the information about at least one edge in f of the paths in the set *Paths*, can be missing in *ConfirmedTopology*.

We consider the problem of relaying messages over the set *CorrectPaths* when only *ConfirmedTopology* is known, and propose three implementations to the func-

The case of constant r and Δ . The sender and the receiver exchange messages by using all possible paths between them; feasible considering r -neighborhoods, where the neighborhood radius, r , and the node degree Δ are constants.

The case of constant f . For each possible choice of f system nodes, p_1, p_2, \dots, p_f , the sender and the receiver compute a new graph $G(p_1, p_2, \dots, p_f)$ that is the result of removing p_1, p_2, \dots, p_f , from G_{out} , which is the graph defined by the discovered topology, $ConfirmedTopology$. Let $\mathcal{P}(p_1, p_2, \dots, p_f)$ be a set of $f + 1$ vertex-disjoint paths in $G(p_1, p_2, \dots, p_f)$ (or the empty set when $\mathcal{P}(p_1, p_2, \dots, p_f)$ does not exist) and $Paths = \bigcup_{p_1, p_2, \dots, p_f} \mathcal{P}(p_1, p_2, \dots, p_f)$. The sender and the receiver can exchange messages over $Paths$, because $|Paths|$ is polynomial at least one choice of p_1, p_2, \dots, p_f , has a corresponding set $\mathcal{P}(p_1, p_2, \dots, p_f)$ that contains *CorrectPaths* (Section B of the Appendix).

The case of no Byzantine neighbors The procedure assumes that any Byzantine node has no directly connected Byzantine neighbor in the communication graph. Specifically, this polynomial cost solution considers the (extended) graph, G_{ext} , that includes all the edges in *confirmedTopology* and *suspicious edges*. Given three nodes, $p_i, p_j, p_k \in P$, we say that node p_i considers the undirected edge (p_k, p_j) suspicious, if the edge appears as a directed edge in $ConfirmedTopology_i$ for only one direction, e.g., (p_j, p_k) .

The extended graph, G_{ext} , may contain fake edges that do not exist in the communication graph, but Byzantine nodes reports on their existence. Nevertheless, G_{ext} includes all the correct paths of the communication graph, G . Therefore, the $2f + 1$ vertex-disjoint paths that exist in G also exist in G_{ext} and they can facilitate a polynomial cost solution for the message exchange task (Section B of the Appendix).

Fig. 2. Implementation proposals for the function *getDisjointPaths()*.

tion *getDisjointPaths()* in Figure 2. The value of *ConfirmedTopology* is a set of directed edges (p_i, p_j) . An undirected edge is approved if both (p_i, p_j) and (p_j, p_i) appear in *ConfirmedTopology*. Other edges in *ConfirmedTopology* are said to be suspicious. For each of the proposed implementations, we show in Section B of the Appendix that a polynomial number of paths are used and that they contain *CorrectPaths*. Thus, the sender and the receiver can exchange messages using a polynomial number of paths and message send operations, because each path is of linear length.

Ensuring safe message delivery We propose a way for the sender and the receiver, that exchange a message over the paths in *getDisjointPaths()*, to stop considering messages and acknowledgments sent by Byzantine nodes. They repeatedly send messages and acknowledgments over the selected vertex-disjoint paths. Before message or acknowledgment delivery, the sender and the receiver expect to receive each message and acknowledgment at least $(capacity \cdot n + 1)$ consecutive times over at least $f + 1$ vertex independent paths, and by that provide evidence that their messages and acknowledgments were indeed sent by them.

We employ techniques for labeling the messages (in an ARQ style), recording visited path of each message, and counting the number of received message over each path. The sender sends messages to the receiver, and the receiver responds with acknowledgments after these messages are delivered to the application layer. Once the sender receives the acknowledgment, it can fetch the next message that should be sent

to the receiver. The difficulty here is to guarantee that the sender and receiver can indeed exchange messages and acknowledgments between them, and stop considering messages and acknowledgments sent by Byzantine nodes.

The sender repeatedly sends message m , which is identified by $m.ARQLabel$, to the receiver over all selected paths. The sender does not stop sending m before it is guaranteed that m was delivered to the application layer of the receiving-side. When the receiver receives the message, the set $m.VisitedPath$ encodes the path along which m was relayed over. Before delivery, the receiver expects to receive m at least $(capacity \cdot n + 1)$ consecutive times from at least $f + 1$ vertex independent paths. Waiting for $(capacity \cdot n + 1)$ consecutive messages on each path, ensures that the receiver gets at least one message which was actually sent recently by the sender. Once the receiver delivers m to the application layer, the receiver starts to repeatedly acknowledge with the label $m.ARQLabel$ over the selected paths (while recording the visited path). The sender expects to receive m 's acknowledgment at least $capacity \cdot n + 1$ consecutive times from at least $f + 1$ vertex independent paths before concluding that m was delivered to the application layer of the receiving-side.

Once the receiver delivers a message to the application layer, the receiver starts to repeatedly acknowledge the recently delivered message over the selected paths. In addition, the receiver also restarts its counters and the log of received messages upon a message delivery to the application layer. Similarly the sender count acknowledgments to the current label used, when the sender receives at least $capacity \cdot n + 1$ acknowledgments over $f + 1$ vertex-disjoint paths, the sender fetches the next message from the application layer, changes the label and starts to send the new message.

The pseudocode of Algorithm 2 In every iteration of the infinite loop, p_i fetches *Message*, prepares *Message*'s label (line 3) and starts sending *Message* over the selected paths, see the procedure *ByzantineFaultTolerantSend(Message)*. When p_i gets enough acknowledgments for *Message* (line 4), p_i stops sending the current message and fetches the next. Upon receiving a message *msg*, node p_i tests *msg*'s destination (line 6). When p_i is not *msg*'s destination, it forwards *msg* to the next node on *msg*'s intended path, after updating *msg*'s visited path. When p_i is *msg*'s destination, p_i checks *msg*'s type (line 9). When *msg*'s type is *Data*, p_i inserts the message payload and label to the part of the data structure associated with the message source, i.e., the sender, and the message visited path (line 10). In line 12, node p_i checks whether $f + 1$ vertex-disjoint paths relayed the message at least $capacity \cdot n + 1$ times, where *capacity* is an upper bound on the number of messages in transit over a communication link. If so, p_i delivers the *msg* to the application layer (line 20), clears the entire data structure and finally sends acknowledgments on the selected paths until a new message is confirmed. Moreover, in line 21 we signal that we are ready to receive a new message. When *msg*'s type is *ACK*, we act almost as when the message is of type *Data*. When the condition in line 18 holds, we signal that the message was confirmed at the receiver by setting *Approved* to be *true*, in line 18. We note that the code of Algorithm 2 considers only one possible pair of source and destination. A many-source to many-destination version of this algorithm can simply use a separate instantiation of this algorithm for each pair of source and destination.

Algorithm 2: Self-stabilizing Byzantine resilient end-to-end delivery (p_i 's code)

Interface: *FetchMessage()*: Gets messages from the upper layer. We denote by *InputMessageQueue* the unbounded queue of all messages that are to be delivered to the destination;

Interface: *DeliverMessage(Source, Message)*: Deliver an arriving message to the higher layer. We denote by *OutputMessageQueue* the unbounded queue of all messages that are to be delivered to the higher layer. We assume that it always contains at least the last message inserted to it;

Input: *ConfirmedTopology*: The discovered topology (represented by a directed edge set, see Algorithm 1);

Data Structure: Transport layer messages: $\langle Source, Destination, VisitedPath, IntendedPath, ARQLabel, Type, Payload \rangle$, where *Source* is the sending node, *Destination* is the target node, *VisitedPath* is the actual relay path, *IntendedPath* is the planned relay path, *ARQLabel* is the sequence number of the stop-and-wait ARQ protocol, and *Type* $\in \{Data, ACK\}$ message type, where DATA and ACK are constant;

Variable *Message*: the current message being sent;

Variable *ReceivedMessages* $[j][Path]$: queue of p_j 's messages that were relayed over path *Path*;

Variable *Confirmations* $[j][Path]$: p_j 's acknowledgment queue for messages that were relayed over *Path*;

Variable *label*: the current sequence number of the stop-and-wait ARQ protocol;

Variable *Approved*: A Boolean variable indicating whether *Message* was accepted at the destination;

Function: *NodeDisjointPaths(S)*: Test *S*, a set of paths, to encode at least $f + 1$ vertex-disjoint paths;

Function: *FloodedPath(MessageQueue, m)*: Test whether *m* is encoded by the first $capacity \cdot n + 1$ messages in *MessageQueue*;

Function: *getDisjointPaths(ReportedTopology, Source, Destination)*: Get a set of vertex-disjoint paths between *Source* and *Destination* in the discovered graph, *ReportedTopology* (Figure 2);

Function: *ClearQueue(Source)*: Delete all data in *ReceivedMessages* $[Source][*]$;

Function: *ClearAckQueue(Destination)*: Delete all data in *Confirmations* $[Destination][*]$;

```
1 while true do
2   ClearAckQueue(Message.Destination)
3   (Message, label)  $\leftarrow$  (FetchMessage(), label + 1 modulo 3)
4   while Approved = false do ByzantineFaultTolerantSend(Message)
5 Upon Receive (msg) From  $p_j$ ;
  begin
6   if msg.Destination  $\neq i$  then
7     msg.VisitedPath  $\leftarrow$  msg.VisitedPath  $\cup \{j\}$ 
8     send(msg) to next (msg.IntendedPath)
9   else if msg.Type = Data then
10    ReceivedMessages[msg.Source][msg.VisitedPath].insert( $\langle$  msg.Payload,
11    msg.ARQLabel  $\rangle$ )
12    let Paths  $\leftarrow \{Path : FloodedPath(Confirmations[msg.Source][Path], msg)\}$ 
13    if NodeDisjointPaths(Paths) then
14      NewMessage  $\leftarrow$  true
15      Confirm(msg.Source, m.ARQLabel, m.Payload)
16   else if msg.Type = ACK then
17     if label = msg.ARQLabel then
18       Confirmations[msg.Source][msg.VisitedPath].insert( $\langle$  msg.Payload, msg.ARQLabel  $\rangle$ )
19       let Paths  $\leftarrow \{Path : FloodedPath(Confirmations[msg.Source][Path],$ 
20        $\langle$  msg.Payload, msg.ARQLabel  $\rangle)\}$ 
21       if NodeDisjointPaths(Paths) then Approved  $\leftarrow$  true
22   Function: Confirm(Source, ARQLabel, Payload);
  begin
23   if CurrentLabel  $\neq$  ARQLabel then DeliverMessage(Source, Payload)
24   (CurrentLabel, NewMessage)  $\leftarrow$  (ARQLabel, false)
25   ClearQueue(Source)
26   while NewMessage = false do ByzantineFaultTolerantSend( $\langle$  Source, ARQLabel,
  ACK, Payload  $\rangle$ )
27   Function: ByzantineFaultTolerantSend( $\langle$  Destination, ARQLabel, Type, Payload  $\rangle$ );
  begin
28   let Paths  $\leftarrow$  getDisjointPaths(ConfirmedTopology, i, Destination)
29   foreach Path  $\in$  Paths do send( $\langle$  i, Destination,  $\emptyset$ , Path, ARQLabel, Type, Payload  $\rangle$ ) to
  first(Path)
```

Correctness proof. We show that message delivery guarantees hold after a bounded convergence period. The proof is based on the system’s ability to relay messages over $f + 1$ correct vertex-disjoint messages (Figure 2), and focuses on showing safe message delivery between the sender and the receiver. After proving that the sender fetches messages infinitely often, we show that within four such fetches, the message delivery guarantees hold; receiver-side delivers all of the sender’s messages and just them. The proof in detail appears in Section C of the Appendix.

Let us consider messages, m , and their acknowledgements, that arrive at least $(\text{capacity} \cdot n + 1)$ times over $f + 1$ vertex-independent paths, to the receiver-side, and respectively the sender-side, with identical payloads and labels. The receiver, and respectively the sender, has the *evidence* that m was *indeed sent by the sender*, and respectively, *acknowledged by the receiver*. The sender and the receiver *clear* their *logs* whenever they have such evidences about m . The proof shows that, after a finite convergence period, the system reaches an execution in which the following events reoccur: (Fetch) the sender clears its log, fetches message m , and sends it to the receiver, (R-Get) the receiver gets the evidence that m was indeed sent by the sender, (Deliver) the receiver clears its log, delivers m , and acknowledge it to the sender, and (S-Get) the sender gets the evidence that m was acknowledged by the receiver. Namely, the system reaches a legal execution.

First we prove that event **Fetch** occurs infinitely often, in the way of proof by contradiction. Let us assume (towards a contradiction) that the sender fetches message m and then never fetches another message m' . The sender sends m and counts acknowledgments that has m ’s label. According to the algorithm, the sender can fetch the next message, $m' \neq m$, when it has the evidence that m was indeed acknowledged by the receiver. The receiver acknowledges m ’s reception when it has the evidence that m was indeed sent by the sender. After nullifying its logs, the receiver repeatedly sends m ’s acknowledgments until it has evidences for other messages, m' , that were indeed sent by the sender after m . By the assumption that the sender never fetches $m' \neq m$, we have that the receiver keeps on acknowledging m until $m' \neq m$ arrives from the sender. Therefore, m arrives from the sender to the receiver, and the receiver acknowledges m to the sender. Thus, a contradiction that the sender never fetches $m' \neq m$.

The rest of the proof shows that (eventually) between every two event of type **Fetch**, also the events **R-Get**, **Deliver** and **S-Get** occur (and in that order). We show that this is guaranteed within four occurrences of event **Fetch**. Following the fetch of each of the first three messages and before the next one, the sender must have evidence that the receiver executed event **Deliver**, i.e., clearing the receiver’s log. Note that during convergence, this may surely be false evidence. Just before fetching a new message in event **Fetch**, the sender must clear its logs and reassign a label value, say, the value is 0. There must be a subsequent fetch with label 1, because, as explained above, event **Fetch** occurs (infinitely often). Since the sender clears its logs in event **Fetch**, from now on and until the next event **Fetch**, any corrupted message found in the sender’s log must be of Byzantine origin. Therefore, the next time sender gets the evidence that m was acknowledged by the receiver, the receiver has truly done so. Note that between any such two (truthful) acknowledgments (with different labels), say with label, 1, 2, \dots , the receiver must execute event **Deliver** and clean its log, see Algorithm 2, line 22. Since

the sender sends over $f + 1$ correct paths, and the receiver's logs are clear, eventually the receiver will have evidence for the message with label 0. As corrupted messages originate only from Byzantine nodes and there are at most f such nodes, the receiver's log may not contain evidence for non-sender messages. To conclude, starting from the 4-th message, the receiver will confirm all of the sender's messages, and will not confirm non-sender messages.

5 Extensions and Conclusions

As an extension to this work, we suggest to combine the algorithms for r -neighborhood network discovery and the end-to-end capabilities in order to allow the use of end-to-end message delivery within the r -neighborhoods. These two algorithms can be used by the nodes, under reasonable node density assumptions, for discovering their r -neighborhoods, and, subsequently, extending the scope of their end-to-end capabilities beyond their r -neighborhood, as we describe in the following. We instruct further remote nodes to relay topology information, and in this way collect information on remote neighborhoods. One can consider an algorithm for studying specific remote neighborhoods that are defined, for example, by their geographic region, assuming the usage of GPS inputs; a specific direction and distance from the topology exploring node defines the exploration goal. The algorithm nominates $2f + 1$ nodes in the specific direction to return further information towards the desired direction. The sender uses end-to-end communication to the current $2f + 1$ nodes in the *front* of the current exploration, asks them for their r -neighborhood, and chooses a new set of $2f + 1$ nodes for forming a new front. It then instructs each of the current nodes in the current front to communicate with each node in the chosen new front, to nominate the new front nodes to form the exploration front.

To ensure stabilization, this interactive process of remote information collection should never stop. Whenever the current collection process investigates beyond the closest r -neighborhood, we concurrently start a new collection process in a pipeline fashion. The output is the result of the last finalized collection process. Thus, having a correct output after the first time a complete topology investigation is finalized.

In this work we presented two deterministic, self-stabilizing Byzantine-resilience algorithms for topology discovery and end-to-end message delivery. We have also considered an algorithm for discovering r -neighborhood in polynomial time, communication and space. Lastly, we mentioned a possible extension for exploring and communicating with remote r -neighborhoods using polynomial resources as well.

The obtained end-to-end capabilities can be used for communicating the public keys of parties and establish private keys, in spite of f corrupted nodes that may try to conduct man-in-the-middle attacks, an attack that the classical Public key infrastructure (PKI) does not cope with. Once private keys are established encrypted messages can be forwarded over any specific $f + 1$ node independent paths, one of which must be Byzantine free. The Byzantine free path will forward the encrypted message to the receiver while all corrupted messages will be discarded. Since our system should be self-stabilizing, the common private secret should be re-established periodically.

A Correctness of Algorithm 1

Lemma 1 (Bounded memory) *Let $p_i \in C$ be a correct node. At any time, there are at most $n \cdot 2^{2n}$ messages in $InformedTopologyany_i$, where $n = |P|$ and $O(|P| \log(|P|))$ is the message size.*

Proof. The queue $InformedTopologyany_i$, is made up of messages in the form $\langle node, neighborhood, visitedpath \rangle$. All nodes that appear in the message, i.e., in the first, second or third entry of the tuple are in N . The first entry, i.e. the node name is one of n possibilities. The second and third entries are subsets of N . Thus each of them has 2^n possibilities. In total there can be at most $2^n \cdot 2^n \cdot n$ messages in every $InformedTopologyany_i$.

Definition 2 specifies the requirements of the network topology discovery task. Definition 3 considers correct paths and Definition 4 considers uncorrupted graph topology messages.

Definition 2 (Legal output). *Given correct node $p_i \in C$, we say that p_i 's output is legal, if it encodes graph $G_{output} = (V_{out}, E_{out})$: (1) $C \subseteq V_{out} \subseteq C \cup B \subseteq N$, and (2) $(E \cap (C \times C)) \subseteq E_{out} \subseteq E \subseteq N \times N$.*

Definition 3 (Correct path). *We say path $\subseteq N$ is a correct one if all its nodes are correct, i.e., $path \subseteq C$.*

Definition 4 (Valid message). *In Algorithm 1, we refer to a message $m = \langle k, Neighborhood_k, VisitedPath_k \rangle$ as a valid message when: (1) $p_k \in C$ and $VisitedPath_k$ encodes a correct path in the communication graph, G , that starts in p_k , and (2) $Neighborhood_k = indices(N_k)$.*

Lemma 2 shows that eventually correct paths do not relay non valid messages. Namely, invalid messages can only exist as the result of: (1) Byzantine interventions that corrupt messages, or (2) transient faults, which occur only prior to the arbitrary starting configuration considered.⁵

We note that we consider asymptotic behavior in the following lemma and thus, *capacity* is omitted from the number of asynchronous round until stabilization.

Lemma 2 (Eventually valid messages). *Let R be a fair execution of Algorithm 1 that starts in an arbitrary configuration. Within $O(|N|)$ asynchronous rounds, the system reaches a configuration after which only valid messages are relayed on correct paths.*

⁵ This is a common way to argue about self-stabilization, we consider executions that start in an arbitrary configuration that follows the last transient fault, recalling that if additional transient faults occur a new arbitrary configuration is reached from which automatic convergence starts.

Proof. Let $c \in R$ be the starting configuration. Suppose that c includes an invalid message, $m = \langle \ell, \text{Neighborhood}_\ell, \text{VisitedPath}_\ell \rangle$, in transit between correct nodes. The lemma is obviously correct for the case that m is relayed by Byzantine nodes during the first $\mathcal{O}(|N|)$ asynchronous rounds of R . Therefore, we consider only the correct paths, $path$, over which m is relayed during the first $\mathcal{O}(|N|)$ asynchronous rounds of R . We show that, within $\mathcal{O}(|N|)$ asynchronous rounds, no correct node in $path$ relays m .

Let $p_j, p_i \in path$ be correct neighbors on the correct path. Suppose that in c , message m is in transit from p_j to p_i . Upon the arrival of message m to p_i (line 7), p_i sends $m_i = \langle \ell, \text{Neighborhood}_\ell, \text{VisitedPath}_\ell \cup \{j\} \rangle$ to any neighbor $p_k \in path$ on the path for which $p_k \in N_i \wedge k \notin \text{VisitedPath}_\ell$, see line 9.

Node p_i adds p_j 's identifier to m 's visited path VisitedPath_ℓ , see line 9. The same argument holds for any correct neighbors, $p'_j, p'_i \in path$ when p_j sends message m'_j to the next node in $path$, node p'_i . Therefore, within $|path \setminus \text{VisitedPath}_\ell|$ asynchronous rounds, it holds that $N'_i \cap (path \setminus \text{VisitedPath}_\ell) = \{p'_j, p'_i\}$.

Note that p'_i makes sure that $\text{VisitedPath}'_\ell$ does not encode loops, i.e., $p_k \notin \text{VisitedPath}'_\ell$, see line 9. Therefore, node p'_i does not relay message m' to p_k .

Definition 5 considers queues that their recent valid messages encode at least $f + 1$ vertex-disjoint paths. Moreover, the invalid ones encode at most f such paths.

Definition 5 (Valid queue). Let $p_i, p_k \in C$ be two correct nodes. We say that p_i 's queue, $\text{InformedTopology}_i$, is valid (with respect to p_k) whenever there is a prefix, $\text{ValidInformation}_{i,k}$, of messages m_k in the queue $\text{InformedTopology}_i$, such that: (1) there is a subset, $\text{Valid} = \{m_\ell = \langle k, \text{Neighborhood}_k, \text{VisitedPath}_\ell \rangle : m_\ell \text{ is valid} \} \subseteq \text{ValidInformation}_{i,k}$, for which the set $\{\text{VisitedPath}_\ell\}$ encodes at least $f + 1$ vertex-disjoint paths, and (2) the set, $\text{Invalid} = \{m_\ell = \langle k, \text{Neighborhood}_k, \text{VisitedPath}_\ell \rangle : m_\ell \text{ is invalid} \} \subseteq \text{ValidInformation}_{i,k}$, for which the set $\{\text{VisitedPath}_\ell\}$ encodes at most f vertex-disjoint paths.

Claim A shows that, within $\mathcal{O}(|C|)$ asynchronous rounds, correct paths propagate valid messages.

Claim. Let $path \subseteq C$ be a correct path from p_i to p_k . Suppose that $m_i = \langle i, N_i, \emptyset \rangle$ is a (valid) message that p_i sends, see line 6. Within $\mathcal{O}(|path|)$ asynchronous rounds, message m_i is relayed on $path$, and arrives at p_k as $m'_i = \langle i, N_i, path \rangle$. Namely, $path$ is m'_i 's visited path.

Proof. Let $c \in R$ be the first configuration that follows the start of m_i 's propagation in $path$. I.e., c is the configuration that immediately follows the step in which node p_i sends m_i by executing line 6. Let $p_r, p_j \in path$ be two correct neighbors on the path. Without the loss of generality, suppose that node p_i sends message m_i directly to node p_r , i.e., in c , node p_r is just about to receive m_i . The proof arguments hold also when assuming that p_j sends message $m_j = \langle i, N_i, \{r\} \rangle$ to the next node in $path$. Thus, generality is not lost.

We show that, within one asynchronous round, p_r sends m_r to p_j . Upon the arrival of message m_i to p_r (line 7), node p_r sends the message m_r to any neighbor, such as p_j ,

for which $p_j \in N_r \wedge r \notin \text{VisitedPath}_i = \emptyset$, see line 9. Since the same argument holds when p_j sends m_j to the next node in path , we have that within $|\text{path}|$ asynchronous rounds, m'_i is delivered to node p_k . \square

Lemma 3 shows that queues get to become valid.

Lemma 3 (Eventually valid queues). *Let R be a fair execution of Algorithm 1 that starts in an arbitrary configuration and $p_i, p_k \in C$ be any pair of correct nodes. The system reaches a configuration in which the queue, $\text{InformedTopology}_i$, is valid (with respect to p_k), within $\mathcal{O}(|N|)$ asynchronous rounds.*

Proof. Let $c \in R$ be a configuration achieved in Lemma 2 within $\mathcal{O}(|N|)$ asynchronous rounds. We show that within $\mathcal{O}(|N|)$ asynchronous rounds after c , the system reaches a configuration in which $\text{InformedTopology}_i$ is valid (with respect to p_k), see Definition 5.

In configuration c , all messages in transit on correct paths are valid, see Lemma 2. Thus, the only messages entering $\text{InformedTopology}_i$ are either valid or have passed through Byzantine nodes. Denote m_{barrier} to be the top message the queue (i.e., the last message entered into the queue) $\text{InformedTopology}_i$ in configuration c . Moreover, $\text{ValidInformation}_{i,k}$ includes all the messages in $\text{InformedTopology}_i$, that are between the queue's head and m_{barrier} .

We show that condition (1) of Definition 5 holds. There are $2f + 1$ vertex-disjoint paths between p_i and p_k . At most f nodes are Byzantine and thus, there are at least $f + 1$ vertex-disjoint paths between p_i and p_k that are correct. By Claim A within $\mathcal{O}(|C|)$ asynchronous rounds, a valid message, m_k , is received on all $f + 1$ (correct) vertex-disjoint paths. Message m_k is inserted to $\text{InformedTopology}_i$ after configuration c . Therefore, m_k is in front of m_{barrier} . Hence, the set $\text{Valid} = \{m_\ell = \langle \ell, \text{Neighborhood}_\ell, \text{VisitedPath}_\ell \rangle : m_\ell \text{ is valid} \} \subseteq \text{ValidInformation}_{i,k}$ contains at least $f + 1$ valid messages whose respective visited paths, VisitedPath_ℓ , are vertex-disjoint.

We show that condition (2) of Definition 5 holds. Any invalid messages, m_k , that is sent after configuration c , must go through a Byzantine node, see Lemma 2.

Claim. Suppose that message m is relayed through a Byzantine node after configuration c , then in any following configuration, while m is still in transit, there is a Byzantine node in the visitedPath.

Proof. Observe the first correct node p_k after the last Byzantine node b on m 's path. p_k is correct, thus it inserts b to the visited path. b is the last on the path and so the visited path must contain it until end of transit or passing through a different Byzantine. \square

Each such Byzantine node is recorded in the message path, see Claim A. Since there are at most f Byzantine nodes, there could be at most f such messages with vertex-disjoint paths. This completes the proof condition (2) and the lemma.

Lemma 6 shows that correct information gets confirmed, and requires Definition 6.

Definition 6 (Message confirmation). We say that message $m_i = \langle k, \text{Neighborhood}_k, \text{VisitedPath}_{k_i} \rangle$ is confirmed (by node p_i) when $\text{Neighborhood}_k \subseteq \text{ConfirmedTopology}_i$.

Lemma 4 (Eventually confirmed messages). Let R be a fair execution of Algorithm 1 that starts in an arbitrary configuration and $p_i, p_k \in C$ be any pair of correct nodes. Within $\mathcal{O}(|N|)$ asynchronous rounds, the system reaches a configuration after which the fact that message $m_i = \langle k, \text{Neighborhood}_k, \text{VisitedPath}_{k_i} \rangle$ is confirmed, implies that $\text{Neighborhood}_k = \text{indices}(N_\ell)$.

Proof. Let $c \in R$ be the first configuration in which $\text{InformedTopology}_i$ is a valid queue and node p_i completes a full iteration of the do forever loop that starts in line 1. By Lemma 3, the system reaches c within $\mathcal{O}(|N|)$ asynchronous rounds.

We know that in configuration c , the array Result_i satisfies that $\text{Result}_i[k] = \text{indices}(N_\ell)$. We go through the computation of Result in lines 2 to 4.

- *ComputeResults()*, line 2. Let $\text{Res}_i[k] = \text{indices}(N'_\ell)$ be *ComputeResults()*'s return value with respect to node p_k . We show that $\text{Res}_i[k] = \text{indices}(N_\ell)$. Moreover, we show that the neighborhood that will be found will be that which is represented in $\text{Valid} = \{m_\ell = \langle k, \text{Neighborhood}_k, \text{VisitedPath}_\ell \rangle : m_\ell \text{ is valid} \} \subseteq \text{ValidInformation}_{i,k}$.

We recall that the set $\{\text{VisitedPath}_\ell\}$ encodes at least $f + 1$ disjoint paths. Also in the prefix $\text{ValidInformation}_{i,k}$ one can not find $f + 1$ invalid messages with vertex-disjoint messages; See Definition 5.

The function must choose the message containing the neighborhood Neighborhood_k . Otherwise, we have chosen a different neighborhood for k , say $\text{Neighborhood}'_k \neq \text{Neighborhood}_k = \text{indices}(N_k)$. That is, at the time of checking line 19 with neighborhood $\text{Neighborhood}_\ell = \text{Neighborhood}'_k$, there were at least $f + 1$ vertex-disjoint paths in $\text{opinion}[\text{Neighborhood}_\ell]$. This is in contradiction to condition (2) of Definition 5. Moreover in line 20, it holds $\text{Count}[k] > f + 1$, since at least all the correct paths were counted.

- *RemoveContradictions()*, line 3. Let $\text{Res}_i = \text{ComputeResults}()$ and $\text{ResRemoveContradictions}_i = \text{RemoveContradictions}(\text{Res}_i)$ (line 3). We show that $\text{ResRemoveContradictions}_i[r] = \text{indices}(N_r)$. The function *RemoveContradictions()* modifies $\text{Res}_i[r]$ only in line 26 by nullifying it whenever $\text{Count}[r] \leq f$. We demonstrate that, for any correct path VisitedPath_k , there exists no p_ℓ for which $\text{PathContradictsNeighborhood}(p_\ell, \text{Res}_i[\ell], \text{VisitedPath}_k) = \text{true}$, which is the condition in line 24.

We explain that there is no node p_ℓ and a contradicting edge (p_j, p_ℓ) with the set $\text{Res}_i[\ell]$. By the assumption that VisitedPath_k is correct and that node $p_\ell \in \text{VisitedPath}_k$, we have that $p_\ell \in C$ is correct. Thus $\text{Res}_i[\ell] = \text{indices}(N_\ell)$, see previous item of this claim on *ComputeResults()*. VisitedPath_k is correct, and therefore (p_j, p_ℓ) must be in VisitedPath_k .

- *RemoveGarbage()*, line 4. This procedure does not modify $\text{Res}_i = \text{RemoveContradictions}(\text{ComputeResults}())$. We have shown that $\text{Result}_i[k] = \text{indices}(N_k)$. Thus, only the correct neighborhood is confirmed for every correct node p_k .

Lemma 5 shows that eventually there are no fake nodes.

Lemma 5 (Eventually no fake nodes). *Let R be a fair execution of Algorithm 1 that starts in an arbitrary configuration, $p_j \in N$ be any node, and $p_\ell \in P \setminus (N)$ be a node that is not included in the communication graph, G . Within $\mathcal{O}(|N|)$ asynchronous rounds, the system reaches a configuration after which $(p_j, p_\ell) \notin \text{ConfirmedTopology}_i$.*

Proof. Let $c \in R$ be the configuration reached within $\mathcal{O}(|N|)$ asynchronous rounds according to Lemma 4. For any correct node, $p_i \in C$, we show that in c , the execution of $\text{RemoveContradictions}()$ results in $\text{Count}_i[\ell] \leq f$ and nullifies $\text{Result}_i[\ell]$.

We start by showing that for every path p that relays a message which encodes the set $\text{Result}_i[\ell]$, and does not contain Byzantine nodes, a contradiction is found in $\text{RemoveContradictions}()$. Namely, the if conditions of line 24 holds.

Note that, p may not be a correct path even though it contains no Byzantine nodes. For example p may contain nodes p_z that are not even in the communication graph, i.e., $p_z \in P \setminus (N)$.

Let $p_r \in N$ be the first correct node on path p . Such a node exists, because p_i is correct and on the path p . Since p_r is correct, after the execution of $\text{ComputeResults}()$, we have that p_r 's neighborhood, N_r , is encoded in $\text{Result}_i[r]$, see Lemma 4.

Denote the last edge in the path (p_r, p_s) , where $p_s \in P \setminus (N)$. Note that node p_s is not a node in the system and since $\text{Result}_i[r]$ encodes N_r 's neighborhood, we have that $p_s \notin \text{Result}_i[r]$. Thus, the edge (p_r, p_s) is contradicting with the set $\text{Result}_i[r]$. Namely, by the condition in line 24, we have that line 25 must decrease $\text{Count}[\ell]$.

We note that immediately before the function $\text{RemoveContradictions}()$ returns, the integer $\text{Count}[\ell]$ may count only incorrect paths, which contain at least one Byzantine node. Since there are at most f Byzantine nodes, $\text{Count}[\ell] \leq f$ as needed.

Theorem 1 demonstrates the self-stabilization properties.

Theorem 1 (Self-stabilization). *Let R be a fair execution of Algorithm 1 that starts in an arbitrary configuration and $p_i \in C$ be a correct node. Within $\mathcal{O}(|N|)$ asynchronous rounds, the system reaches a safe configuration after which p_i 's output is always legal, see Definition 2.*

Proof. The systems reaches configuration $c \in R$ of Lemma 4 within $\mathcal{O}(|N|)$ asynchronous rounds. We show that c is a safe configuration by showing that the output is legal, we must show that $\text{ConfirmedTopology}_i$ encodes a graph $G_{\text{output}} = (V_{\text{out}}, E_{\text{out}})$, such that: (1) $C \subseteq V_{\text{out}}$, (2) $(E \cap (C \times C)) \subseteq E_{\text{out}}$, (3) $V_{\text{out}} \subseteq C \cup B \subseteq N$, and (4) $E_{\text{out}} \subseteq (E \cap (C \times C)) \cup (B \times (N)) \subseteq P \times N$.

For every correct node $p_k \in C$, we have that N_k is confirmed in c , see Lemma 4. Thus, $p_k \in V_{\text{out}}$ and condition (1) holds.

Let (p_j, p_k) be an edge in the communication graph between two correct nodes, we show $(p_j, p_k) \in E_{\text{out}}$. Since p_j is correct, it is inserted to $\text{ConfirmedTopology}_i$, see Lemma 4. Thus, $(p_j, p_k) \in \text{edges}(N_j) \wedge \text{edges}(N_j) \subseteq \text{ConfirmedTopology}_i$ in c , thus condition (2) holds as well.

There is no $p_\ell \in P \setminus (N)$ and node $p_j \in N$, such that $(p_\ell p_j) \in \text{ConfirmedTopology}_i$, see Lemma 5. Thus, $V_{\text{out}} \subseteq C \cup B \subseteq N$ and $E_{\text{out}} \subseteq (E \cap (C \times C)) \cup (B \times (N)) \subseteq P \times N$. I.e., conditions (3) and (4) hold in c .

B Implementation proposals for *getDisjointPaths()*

We consider the problem of relaying messages over the set *CorrectPaths* when only *ConfirmedTopology* is known, and propose three implementations to the function *getDisjointPaths()*. The value of *ConfirmedTopology* is a set of directed edges (p_i, p_j) . An undirected edge is approved if both (p_i, p_j) and (p_j, p_i) appear in *ConfirmedTopology*. Other edges in *ConfirmedTopology* are said to be suspicious. The arguments used here assume that the system is in a safe configuration with respect to Algorithm 1. For each of the proposed implementations, we show that $|Paths|$ is polynomial and *CorrectPaths* \subseteq *Paths*. Thus, the sender and the receiver can exchange messages using a polynomial number of paths and message send operations, because each path in *Paths* is of linear length.

The case of constant r and Δ . The sender and the receiver exchange messages by using all possible paths between them. This is feasible only when considering r -neighborhoods, rather than the entire connected component, where the neighborhood radius, r , and the node degree Δ are constants.

The case of constant f . This procedure entails sending a message on a path set, *Paths*, where $|Paths|$ is polynomial and *CorrectPaths* \subseteq *Paths*.

For each possible choice of f system nodes, p_1, p_2, \dots, p_f , the sender and the receiver compute a new graph $G(p_1, p_2, \dots, p_f)$ that is the result of removing p_1, p_2, \dots, p_f , from G_{out} , which is the graph defined by the discovered topology, *ConfirmedTopology*. Let $\mathcal{P}(p_1, p_2, \dots, p_f)$ be a set of $f + 1$ vertex-disjoint paths in $G(p_1, p_2, \dots, p_f)$ (or the empty set when $\mathcal{P}(p_1, p_2, \dots, p_f)$ does not exist) and $Paths = \bigcup_{p_1, p_2, \dots, p_f} \mathcal{P}(p_1, p_2, \dots, p_f)$. We show polynomial message cost by showing that $|Paths|$ is polynomial. We also show that for at least one choice of p_1, p_2, \dots, p_f , has a corresponding set $\mathcal{P}(p_1, p_2, \dots, p_f)$ that contains *CorrectPaths*.

First we show that this procedure only sends messages through a polynomial number of paths. There are $\mathcal{O}(n^f)$ possible choices of f system nodes. Thus, $\mathcal{O}(n^f)$ path sets are computed, and since f is a constant, this number is polynomial. Moreover, each such set contains at most $f + 1$ simple paths of linear length, because p_i only computes sets, $\mathcal{P}(p_1, p_2, \dots, p_f)$, of size $f + 1$. Thus, the sender and the receiver can exchange messages using a polynomial number of paths and message send operations.

We show that *CorrectPaths* \subseteq *Paths*. Consider the permutation choice, p_1, p_2, \dots, p_f , in which the set actually contains the set of Byzantine nodes in the system. Thus $G(p_1, p_2, \dots, p_f)$ contains only correct nodes. Furthermore, at least $f + 1$ paths that were present in G_{out} are still present in $G(p_1, p_2, \dots, p_f)$, since $G(p_1, p_2, \dots, p_f)$ was obtained from G_{out} by the removal of f (Byzantine) nodes, p_1, p_2, \dots, p_f . Hence, there are at least $f + 1$ correct vertex-disjoint paths in $G(p_1, p_2, \dots, p_f)$, in $\mathcal{P}(p_1, p_2, \dots, p_f)$ and in *Paths*.

The case of no Byzantine neighbors The procedure assumes that any Byzantine node has no directly connected Byzantine neighbor in the communication graph. Specifically, this polynomial cost solution considers the (extended) graph, G_{ext} , that includes all the edges in *confirmedTopology* and *suspicious edges*. Given three nodes, $p_i, p_j, p_k \in P$, we say that node p_i considers the undirected edge (p_k, p_j) suspicious, if the edge appears as a directed edge in *ConfirmedTopology_i* for only one direction, e.g., (p_j, p_k) .

The extended graph, G_{ext} , may contain fake edges that do not exist in the communication graph, but Byzantine nodes report on their existence. Nevertheless, G_{ext} includes all the correct paths of the communication graph, G . Therefore, the $2f + 1$ vertex-disjoint paths that exist in G also exist in G_{ext} and they can facilitate a polynomial cost solution for the message exchange task, as we next show.

Let $G' = (N, E_{G'})$ be the graph computed from *ConfirmedTopology* and its suspicious edges. We demonstrate that G' 's edges, $E_{G'}$, contains the edges, E_G , of the communication graph, G . Let us consider $e = (p_j, p_k) \in E_G$ and show that $e \in E_{G'}$. When both p_j and p_k are correct, the correctness of Algorithm 1 implies $e \in E_{G'}$. Suppose that p_j is correct and p_k is Byzantine, and consider the different cases in which p_k decides to report (or not to report) about e as part of its local neighborhood. Namely, either $e \in \text{ConfirmedTopology}$, or e is a suspicious edge, because p_i reports about e , and p_k decides to report, and respectively, not to report. Since $G \subseteq G'$, G' must contain $2f + 1$ vertex-disjoint paths between any sender p_s and receiver p_r , because G does. Moreover, the same argument implies that there may be at most f incorrect paths, which contain each at least one Byzantine node. Hence, there are at least $f + 1$ correct vertex-disjoint paths in $Paths$.

C Correctness of Algorithm 2

Definitions 7, 8 and 9 are needed for Claim C, Claim C and Lemma 6.

Definition 7 (Confirmation). *Given configuration c , we say that message m is confirmed (by the receiver) when $m \in \text{OutputMessageQueue}$.*

Definition 8 (Approve). *Given fair execution, R , of Algorithm 2, we say that message $m = \langle \text{Source}, \text{Destination}, \text{VisitedPath}, \text{IntendedPath}, \text{ARQLabel}, \text{DATA}, \text{Payload} \rangle$ is being approved (by the sender p_{Source}) during the first atomic step, a_{sender} , in which the sender executes line 18, where $\text{Source} = \text{sender ARQLabel} = m.\text{ARQLabel}$ and $\text{Payload} = m.\text{Payload}$, see line 17. Denote by c_{approved} the configuration that immediately follows a_{sender} . Given configuration c that appears after c_{approved} in R , we say that message m is approved (by the sender) in configuration c .*

Definition 9 (Clear-sender-receiver). *Given configuration c , we say that the sender is clear (with respect to the receiver), if the queue $\text{Confirmations}[\text{receiver}] = \emptyset$ in c . Moreover, the receiver is clear (with respect to the sender), if the queue $\text{ReceivedMessages}[\text{sender}] = \emptyset$ in c .*

Claim C shows that a message that is relayed on a correct path is received at the destination within $\mathcal{O}(|N|)$ asynchronous rounds. Moreover, the destination receives the message with correct visiting set.

Claim. Let R be a fair execution of Algorithm 2 that starts in a safe configuration, c , with respect to Algorithm 1. Let $p_{\text{source}}, p_{\text{dest}} \in C$ be pair of correct nodes. Let c_{send} be the configuration immediately following a step in which p_{source} sends message $M\text{sg}$ on a correct path $Path = p_{\text{source}}, p_1, p_2, \dots, p_{\text{dest}}$ from source, p_{source} , to destination, p_{dest} . Within $\mathcal{O}(|N|)$ asynchronous rounds, p_{dest} receives $M\text{sg}$ with a visiting set containing all nodes on $Path$ except p_{dest} .

Proof. Upon the arrival of message m to p_k (line 5), node p_i asserts that he is not the destination, p_{dest} , (line 6). Immediately after, p_i sends the message m to the next neighbor, p_{i+1} , see line 8. Since the same argument holds when p_j sends m to the next node in $path$, we have that within $|Path|$ asynchronous rounds, m is delivered to node p_{dest} .

Claim C says that when the sender repeatedly sends message Msg , for a duration of at least $\mathcal{O}(|N|)$ asynchronous rounds, the receiver eventually confirms message Msg .

Claim. Let R be a fair execution of Algorithm 2 that starts in a safe configuration, c , with respect to Algorithm 1. Let $p_s, p_r \in C$ be a pair of correct sending and receiving nodes. Suppose that, for a duration of at least $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, p_s 's steps include only the execution of the function *ByzantineFaultTolerantSend*(Msg) in the loop of line 4. Within that period, the system reaches configuration $c_{receive}$ in which p_r confirms Msg .

Proof. Denote c_{send} as the configuration immediately following the first step in which p_s sends message Msg in R , see line 26. Within $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, the first frame containing Msg arrives at p_r , see Claim C. Moreover, after another $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, every correct path relays message Msg at least $\mathcal{O}(capacity \cdot |N|)$ times. This is correct since every asynchronous round, p_s sends a new frame containing Msg on each of the $2f + 1$ vertex-disjoint paths. Moreover, by Claim C, the last frame sent on all $2f + 1$ paths arrives after another $\mathcal{O}(capacity \cdot |N|)$.

Assume, in the way of proof by contradiction, that Msg is not confirmed by p_r . This implies that the queues, $ReceivedMessages[p_s][*]$, in p_r containing messages sent from p_s were not cleared at least since c_{send} , see line 22. Thus, p_r contains $capacity \cdot n + 1$ indications of Msg on $f + 1$ vertex-disjoint paths. Denote c_{last} as the configuration immediately after the arrival of the $(capacity \cdot n + 1)$ -th frame of the $f + 1$ 'th path to relay $capacity \cdot n + 1$ frames containing Msg . Immediately after c_{last} , p_s must go through line 12, because the conditions in line 12 hold. Thus, a contradiction and Msg is confirmed within $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds.

Claim C says that when the receiver is sending acknowledgments about a message, that message eventually becomes approved. We note that Claim C considers acknowledgments sent from the receiver to the sender, rather than messages sent from the sender to the receiver, as in Claim C.

Claim. Let R be a fair execution of Algorithm 2 that starts in a safe configuration, c , with respect to Algorithm 1. Let $p_s, p_r \in C$ be a pair of correct sending and receiving nodes. Suppose that, for a duration of at least $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, p_r 's steps include only the execution of the function *ByzantineFaultTolerantSend*(Ack) in the loop of line 23. That is, p_r is sending acknowledgments on message Msg . Within that period, the system reaches configuration $c_{receive}$ in which p_s approves Msg , see Definition 8.

Proof. Denote c_{send} as the configuration immediately following the first step in which p_r sends acknowledgment Ack in R , see line 23. Within $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, the first frame containing Ack arrives at p_s , see Claim C. Moreover,

after another $\mathcal{O}(\text{capacity} \cdot |N|)$ asynchronous rounds, every correct path relays message *Ack* at least $\mathcal{O}(\text{capacity} \cdot |N|)$ times. This is correct since every asynchronous round, p_r sends a new frame containing *Ack* on each of the $2f + 1$ vertex-disjoint paths. Moreover, by Claim C, the last frame sent on all $2f + 1$ paths arrives after another $\mathcal{O}(\text{capacity} \cdot |N|)$.

The queues, $\text{Confirmations}[p_r][*]$ are cleared only when a message sent to p_r is approved, see line 2. Since, p_r is acknowledging the current message, *Msg*, by sending *Ack*, the only message that can be approved is *Msg*. This is true since each path, *Path*, may contain at most $\text{capacity} \cdot |N|$ acknowledgments for other messages in the path queues.

Assume, in the way of proof by contradiction, that *Msg* is not approved by p_s . By the arguments above, p_s 's queues, $\text{Confirmations}_s[p_r][*]$, which contains p_r 's acknowledgments that p_s received, were not cleared at least since c_{send} , see line 2. Thus, p_s contains $\text{capacity} \cdot n + 1$ indications of *Ack* on $f + 1$ vertex-disjoint paths. Denote c_{last} as the configuration immediately after the arrival of the $(\text{capacity} \cdot n + 1)$ -th frame of the $f + 1$ 'th path to relay $\text{capacity} \cdot n + 1$ frames containing *Ack*. Immediately after c_{last} , p_s must go through line 18, because the conditions in line 18 hold. Thus, a contradiction and *Msg* is approved within $\mathcal{O}(\text{capacity} \cdot |N|)$ asynchronous rounds.

Lemma 6 shows that the senders repeatedly fetch messages.

Lemma 6. *Let R be a fair execution of Algorithm 2 that starts in a safe configuration, c , with respect to Algorithm 1. Let $p_s, p_r \in C$ be pair of correct sending and receiving nodes. Moreover, c_ℓ is the configuration that immediately follows the ℓ -th time in R in which p_s fetches a message from the input queue. For every ℓ , the system reaches c_ℓ within $\mathcal{O}(\ell \cdot |N|)$ asynchronous rounds.*

Proof. By the code of Algorithm 2, on every iteration of the do forever loop (lines 2 to 4), a message is fetched in line 3. This do forever loop includes another loop in line 4. We prove the lemma by showing that the loop of line 4 is completed within $\mathcal{O}(|N|)$ asynchronous rounds.

The proof considers the case in which the sender, p_s , does not wait in line 4 for a long time before considering the case in which p_s does wait. We show that for the latter case, the receiver, p_r , confirms p_s 's current message. After confirming the message, the receiver, p_r , begins sending acknowledgments to the sender, p_s . The proof shows that after the acknowledgments are sent, p_s approves the message and fetches a new one. We show this by considering the case in which p_r repeatedly sends acknowledgments for a sufficient amount of time, and a case in which it does not.

Suppose that p_s does not wait in line 4 more than $\mathcal{O}(\text{capacity} \cdot |N|)$ asynchronous rounds. In this case, p_s starts the infinite loop again within $\mathcal{O}(\text{capacity} \cdot |N|)$ asynchronous rounds, and fetch a new message, see line 3. Thus, for the case in which p_s does not wait in line 4 more than $\mathcal{O}(\text{capacity} \cdot |N|)$ asynchronous rounds, the lemma is correct.

Suppose that p_s is executing line 4 and waits for acknowledgments on message *Msg* for more than $\mathcal{O}(\text{capacity} \cdot |N|)$ asynchronous rounds. Thus, p_s floods $2f + 1$ vertex-disjoint paths with the message *Msg*, see Figure 2. Eventually, the receiver, p_r , receives message *Msg* for $\mathcal{O}(\text{capacity} \cdot |N|)$ times on $f + 1$ vertex-disjoint paths and confirms

Msg , see Claim C. After confirming it, the receiver sends acknowledgments on $2f + 1$ vertex-disjoint paths until confirming a new message Msg_{new} . This is true because the condition in line 23 holds only when a new message is confirmed, see line 13.

Let us consider the case in which, during $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, message Msg_{new} does not arrive to the receiver. By Claim C, eventually the sender receives the acknowledgments for $capacity \cdot n + 1$ times on $f + 1$ vertex-disjoint paths. Claim C also says that the sender considers the message accepted by the receiver. In line 18, the sender assigns $Approved = \mathbf{true}$. Thus, the condition in line 4 holds and the sender fetches the next message, see line 3. Hence, the system reaches configuration c_{fetch} that immediately follows a step in which the sender, p_s , fetches the next message. Thus, for the case in which, during $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, message Msg_{new} does not arrive to the receiver, the lemma is correct.

We continue by considering the case in which, during $\mathcal{O}(capacity \cdot |N|)$ asynchronous rounds, message Msg_{new} does arrive to the receiver. Let c_{conf} be the configuration that immediately follows the step in which p_r confirms Msg . Since the receiver confirms Msg , we have that p_r is clear (with respect to the sender) in configuration c_{conf} , see Definition 9 and line 22.

If Msg_{new} was sent by the sender, it must have been fetched after c , and c_{fetch} is reached when message Msg_{new} is fetched. It may be the case however, that Msg_{new} was not sent by the sender. Message Msg_{new} was confirmed by $2f + 1$ vertex-disjoint paths. Since there are at most f Byzantine nodes, at least one of these paths, $Path$, must be correct. Moreover, in c_{conf} , the receiver is clear, thus the $capacity \cdot n + 1$ that p_r counts in $ReceivedMessages[p_s][*]$ have all been received after configuration c_{conf} . Note that the sender sends at least one of these messages, because at most $capacity \cdot n$ messages could be in the edges of $Path$ at any given configuration. Thus the sender sends Msg_{new} , which p_s fetches immediately before c_{fetch} . I.e., the system reaches c_{fetch} .

Theorem 1 says that, starting from the fourth (or even the third) message that the sender fetches, the receiver confirms the sender's messages. The proof of Theorem 1 is based on Lemma 7, which says that, in every sequence of four messages that the sender is fetching, the receiver confirms the fourth (or even the third) message.

Lemma 7. *Let R be a fair execution of Algorithm 2 that starts in a safe configuration, c_{start} , with respect to Algorithm 1. Let c_h be a configuration that immediately follows the h -th step in which the sender fetches the h -th input queue message, m_h . Within $\mathcal{O}(|N|)$ asynchronous rounds, the receiver confirms message m_4 .*

Proof. Claim. In c_2 , the sender is clear (with respect to the receiver), see Definition 9.

Proof. By definition, c_2 immediately follows atomic step a_2 , in which, after clearing the confirmation queue in line 2, the sender fetches message m_2 and sends it. \square

Claim. Between the configurations c_3 and c_4 , there is a configuration $c_{receiver-clear}$ in which the receiver is clear (with respect to the sender).

Proof. Suppose, without the loss of generality, that immediately after $c_{sender-clear}$, the sender is waiting for a message with label 1. By lemma 6, the sender eventually fetches the next message. The sender can only fetch a new message once

Approved is true, see line 4. Moreover, *Approved* is only set to *true* once the queue $Confirmations[receiver][*]$ contains $2f + 1$ flooded paths, see line 18. Thus, the sender counts $2f + 1$ vertex-disjoint paths that relayed acknowledgments with label 1. Moreover, the sender is clear in $c_{sender-clear}$. Hence, configuration $c_{sender-clear}$ contains no message in $Confirmations[receiver][*]$. Starting from $c_{sender-clear}$, the sender receives $capacity \cdot n + 1$ acknowledgments on $2f + 1$ vertex-disjoint paths for the current message with label 1. Note that at least one of these $2f + 1$ paths, $Path$, is correct, because there are f Byzantine. Since $|Path| \leq n$ and each edge on $Path$ may contain at most $capacity$ messages, we have that at least one of the acknowledgments that includes $Path$ as its visiting path, is sent by the receiver between $c_{sender-clear}$ and configuration $c_{receiver-send} \in R$. We show that $c_{receiver-send} = c_{receiver-clear}$.

This means that after $c_{sender-clear}$, the sender clears the confirmations queue, $Confirmations[receiver][*]$, and fetches the next message, assigning it the label 2, see lines 2 through line 4. By similar arguments, we know that the receiver sends at least one acknowledgment with label 2.

To conclude, there is a configuration $c \in R$ in which the receiver is sending acknowledgments with label 1, and then a configuration c' in which the receiver sends acknowledgments with label 2. Moreover, between two consecutive executions of line 23, the receiver has to go through line 22. Thus, the receiver cleared its message queues, $Confirmations[sender][*]$, immediately before configuration $c_{receiver-clear}$ and $c_{receiver-send} = c_{receiver-clear}$. \square

Let us consider configuration $c_{receiver-clear}$ from the end of proof of Claim C.

The next message to be sent after $c_{receiver-clear}$, is m_4 , the message fetched in c_4 , with label 0. Between $c_{receiver-clear}$ and c_4 , all messages sent by the sender have the label 2. By arguments stated above, the message, m , that is the next message to be confirmed after $c_{receiver-clear}$, must have been sent by the sender at least once since $c_{receiver-clear}$. The sender, sends only messages with label 0 and 2. Moreover, the last message to be confirmed had a label 2. Thus, $CurrentLabel = 2$, see line 21. Any sent message with label 2 is not inserted to the confirmations queue, $Confirmations[sender][*]$ between $c_{receiver-clear}$ and the configuration that immediately follows the next sender's fetch, see line 20. Thus, by line 3, the next message to be confirmed is a message with label 0, which must be m_4 .

Theorem 1 (Self-stabilization) *Let R be a fair execution of Algorithm 2 that starts in an arbitrary configuration. Within $\mathcal{O}(|N|)$ asynchronous rounds, the system reaches a safe configuration c after which: (1) for every step a_s^m where the sender sends m there is a corresponding step $a_r^m \in R$ where the receiver confirms message m , and (2) for every step a_r^m , there is a corresponding step, $a_s^m \in R$, that occurs before a_r^m and in which the sender sends m .*

Proof. Let c be the configuration that Claim C denote as c_4 , which the system reaches within $\mathcal{O}(|N|)$ asynchronous rounds, see Lemma 6. Let m_i be the i -th message fetched.

Suppose that $i \geq 4$. Lemma 7 considers the four consecutive messages m_{i-3}, \dots, m_i and says that the receiver confirms message m_i . Thus, condition (1) holds.

Condition (2) follows from arguments similar to the ones used in the proof of Claim C. Namely, for the case of $i \geq 5$, message m_{i-1} is con-

firmed, see lemma 7. Immediately after the receiver confirms m_{i-1} , it clears the queue $ReceivedMessages[sender][*]$, see lines 20 to 22. Thus, there exists a configuration $c_{receiver-clear}$ in which the receiver is clear (with respect to the sender) before c_i , see Definition 9. Moreover, a message is confirmed only if the queue $ReceivedMessages[sender][*]$ contains $2f + 1$ flooded paths, see line 12. These flooded paths implies that in configuration c_i , the queue $ReceivedMessages[sender][*]$ contains $capacity \cdot n + 1$ indications of m_i on $2f + 1$ node disjoint paths. Thus, m_i is confirmed only after a period that follows $c_{receiver-clear}$ and includes its reception at least $capacity \cdot n + 1$ times on each of the $2f + 1$ vertex-disjoint paths.

Recall that we assume that there are at most f Byzantine nodes in the system. At least one path, $Path$, of the above $2f + 1$ paths is correct. Moreover, $|Path| \leq n$ and each edge on $Path$ may contain at most $capacity$ messages. Thus, at least one of the $capacity \cdot n + 1$ message that were relayed on the correct path $Path$ was sent by the sender. This completes the correctness proof.