

Discovering Network Topology in the Presence of Byzantine Faults

Mikhail Nesterenko^{1*} and Sébastien Tixeuil^{2**}

¹ Computer Science Department, Kent State University Kent, OH, 44242, USA,
`mikhail@cs.kent.edu`

² LRI-CNRS UMR 8623 & INRIA Grand Large
Université Paris Sud, France, `tixeuil@lri.fr`

Abstract. We study the problem of Byzantine-robust topology discovery in an arbitrary asynchronous network. We formally state the weak and strong versions of the problem. The weak version requires that either each node discovers the topology of the network or at least one node detects the presence of a faulty node. The strong version requires that each node discovers the topology regardless of faults.

We focus on non-cryptographic solutions to these problems. We explore their bounds. We prove that the weak topology discovery problem is solvable only if the connectivity of the network exceeds the number of faults in the system. Similarly, we show that the strong version of the problem is solvable only if the network connectivity is more than twice the number of faults.

We present solutions to both versions of the problem. Our solutions match the established graph connectivity bounds. The programs are terminating, they do not require the individual nodes to know either the diameter or the size of the network. The message complexity of both programs is low polynomial with respect to the network size.

1 Introduction

In this paper, we investigate the problem of Byzantine-tolerant distributed topology discovery in an arbitrary network. Each node is only aware of its neighboring peers and it needs to learn the topology of the entire network.

Topology discovery is an essential problem in distributed computing (*e.g.* see [1]). It has direct applicability in practical systems. For example, link-state based routing protocols such as OSPF use topology discovery mechanisms to compute the routing tables. Recently, the problem has come to the fore with the introduction of ad hoc wireless sensor networks, such as Berkeley motes [2], where topology discovery is essential for routing decisions.

* This author was supported in part by DARPA contract OSU-RF#F33615-01-C-1901 and by NSF CAREER Award 0347485.

** This author was supported in part by the FNS grants FRAGILE and SR2I from ACI “Sécurité et Informatique”.

As reliability demands on distributed systems increase, the interest in developing robust topology discovery programs grows. One of the strongest fault models is *Byzantine* [3]: the faulty node behaves arbitrarily. This model encompasses rich set of fault scenarios. Moreover, Byzantine fault tolerance has security implications, as the behavior of an intruder can be modeled as Byzantine. One approach to deal with Byzantine faults is by enabling the nodes to use cryptographic operations such as digital signatures or certificates. This limits the power of a Byzantine node as a non-faulty node can verify the validity of received topology information and authenticate the sender across multiple hops. However, this option may not be available. For example, wireless sensors may not have the capacity to manipulate digital signatures. Another way to limit the power of a Byzantine process is to assume synchrony: all processes proceed in lock-step. Indeed, if a process is required to send a message with each pulse, a Byzantine process cannot refuse to send a message without being detected. However, the synchrony assumption may be too restrictive for practical systems.

Our contribution. In this study we explore the fundamental properties of topology discovery. We select the weakest practical programming model, establish the limits on the solutions and present the programs matching those limits.

Specifically, we consider arbitrary networks of arbitrary topology where up to fixed number of nodes k is faulty. The execution model is asynchronous. We are interested in solutions that do not use cryptographic primitives. The solutions should be terminating and the individual processes should not be aware of the network parameters such as network diameter or its total number of nodes.

We state two variants of the topology discovery problem: *weak* and *strong*. In the former — either each non-faulty node learns the topology of the network or one of them detects a fault; in the latter — each non-faulty node has to learn the topology of the network regardless of the presence of faults.

As negative results we show that any solution to the weak topology discovery problem can not ascertain the presence of an edge between two faulty nodes. Similarly, any solution to the strong variant can not determine the presence of an edge between a pair of nodes at least one of which is faulty. Moreover, the solution to the weak variant requires the network to be at least $(k+1)$ -connected. In case of the strong variant the network must be at least $(2k + 1)$ -connected.

The main contribution of this study are the algorithms that solve the two problems: *Detector* and *Explorer*. The algorithms match the respective lower bounds. To the best of our knowledge, these are the first asynchronous Byzantine-robust solutions to the topology discovery problem that do not use cryptographic operations. *Explorer* solves the stronger problem. However, *Detector* has better message complexity. *Detector* either determines topology or signals fault in $O(\delta n^3)$ messages where δ and n are the maximum neighborhood size and the number of nodes in the system respectively. *Explorer* finishes in $O(n^4)$ messages. We extend our algorithms to (a) discover a fixed number of routes instead of complete topology and (b) reliably propagate arbitrary information instead of topological data.

Related work. A number of researchers employ cryptographic operations to counter Byzantine faults. Avromopolus et al [4] consider the problem of secure routing. Therein see the references to other secure routing solutions that rely on cryptography. Perrig et al [5] survey robust routing methods in ad hoc sensor networks. The techniques covered there also assume that the processes are capable of cryptographic operations.

A naive approach of solving the topology discovery problem without cryptography would be to use a Byzantine-resilient broadcast [6–9]: each node advertises its neighborhood. However all existing solutions for arbitrary topology known to us require that the graph topology is *a priori* known to the nodes.

Let us survey the non-cryptography based approaches to Byzantine fault-tolerance. Most programs described in the literature [10–13] assume completely connected networks and can not be easily extended to deal with arbitrary topology. Dolev [7] considers Byzantine agreement on arbitrary graphs. He states that for agreement in the presence of up to k Byzantine nodes, it is necessary and sufficient that the network is $(2k + 1)$ -connected and the number of nodes in the system is at least $3k + 1$. However, his solution requires that the nodes are aware of the topology in advance. Also, this solution assumes the synchronous execution model. Recently, the problem of Byzantine-robust reliable broadcast has attracted attention [6, 8, 9]. However, in all cases the topology is assumed to be known. Bhandari and Vaidya [6] and Koo [8] assume two-dimensional grid. Pelc and Peleg [9] consider arbitrary topology but assume that each node knows the exact topology a priori. A notable class of algorithms tolerates Byzantine faults locally [14–16]. Yet, the emphasis of these algorithms is on containing the fault as close to its source as possible. This is only applicable to the problems where the information from remote nodes is unimportant such as vertex coloring, link coloring or dining philosophers. Thus, local containment approach is not applicable to topology discovery.

Masuzawa [17] considers the problem of topology discovery and update. However, Masuzawa is interested in designing a self-stabilizing solution to the problem and thus his fault model is not as general as Byzantine: he considers only transient and crash faults.

The rest of the paper is organized as follows. After stating our programming model and notation in Section 2, we formulate the topology discovery problems, as well as state the impossibility results in Section 3. We present *Detector* and *Explorer* in Sections 4 and 5 respectively. We discuss the composition of our programs and their extensions in Section 6 and conclude the paper in Section 7.

2 Notation, Definitions and Assumptions

Graphs. A distributed *system* (or *program*) consists of a set of processes and a *neighbor* relation between them. This relation is the system *topology*. The

topology forms a graph G . Denote n and e to be the number of nodes³ and edges in G respectively. Two processes are *neighbors* if there is an edge in G connecting them. A set P of neighbors of process p is *neighborhood* of p . In the sequel we use small letters to denote singleton variables and capital letters to denote sets. In particular, we use a small letter for a process and a matching capital one for this process' neighborhood. Since the topology is symmetric, if $q \in P$ then $p \in Q$. Denote δ to be the maximum number of nodes in a neighborhood.

A *node-cut* of a graph is the set of nodes U such that $G \setminus U$ is disconnected or trivial. A *node-connectivity* (or just *connectivity*) of a graph is the minimum cardinality of a node-cut of this graph. In this paper we make use of the following fact about graph connectivity that follows from Menger's theorem (see [18]): if a graph is k -connected (where k is some constant) then for every two vertices u and v there exists at least k internally node-disjoint paths connecting u and v in this graph.

Program model. A process contains a set of variables. When it is clear from the context, we refer to a variable *var* of process p as *var.p*. Every variable ranges over a fixed domain of values. For each variable, certain values are *initial*. Each pair of neighbor processes share a pair of special variables called *channels*. We denote *Ch.b.c* the channel from process b to process c . Process b is the *sender* and c is the *receiver*. The value for a channel variable is chosen from the domain of (potentially infinite) sequences of messages.

A *state* of the program is the assignment of a value to every variable of each process from its corresponding domain. A state is *initial* if every variable has initial value. Each process contains a set of actions. An action has the form $\langle name \rangle : \langle guard \rangle \longrightarrow \langle command \rangle$. A *guard* is a boolean predicate over the variables of the process. A *command* is sequence of assignment and branching statements. A guard may be a receive-statement that accesses the incoming channel. A command may contain a send-statement that modifies the outgoing channel. A parameter is used to define a set of actions as one parameterized action. For example, let j be a parameter ranging over values 2, 5 and 9; then a parameterized action *ac.j* defines the set of actions $ac.(j = 2) [] ac.(j = 5) [] ac.(j = 9)$. Either guard or command can contain quantified constructs [19] of the form: $(\langle quantifier \rangle \langle bound variables \rangle : \langle range \rangle : \langle term \rangle)$, where *range* and *term* are boolean constructs.

Semantics. An action of a process of the program is *enabled* in a certain state if its guard evaluates to **true**. An action containing receive-statement is enabled when appropriate message is at the head of the incoming channel. The execution of the command of an action updates variables of the process. The execution of an action containing receive-statement removes the received message from the head of the incoming channel and inserts the value the message contains into the specified variables. The execution of send-statement appends the specified message to the tail of the outgoing message.

³ We use terms *process* and *node* interchangeably.

A *computation* of the program is a maximal fair sequence of states of the program such that the first state s_0 is initial and for each state s_i the state s_{i+1} is obtained by executing the command of an action whose state is enabled in s_i . That is, we assume that the action execution is *atomic*. The maximality of a computation means that the computation is either infinite or it terminates in a state where none of the actions are enabled. The fairness means that if an action is enabled in all but finitely many states of an infinite computation then this action is executed infinitely often. That is, we assume *weak fairness* of action execution. Notice that we define the receive statement to appear as a standalone guard of an action. This means, that if a message of the appropriate type is at the head of the incoming channel, the receive action is enabled. Due to weak fairness assumption, this leads to *fair message receipt* assumption: each message in the channel is eventually received. Observe that our definition of a computation considers *asynchronous* computations.

To reason about program behavior we define boolean predicates on program states. A program *invariant* is a predicate that is **true** in every initial state of the program and if the predicate holds before the execution of the program action, it also holds afterwards. Notice that by this definition a program invariant holds in each state of every program computation.

Faults. Throughout a computation, a process may be either Byzantine (faulty) or non-faulty. A Byzantine process contains an action that assigns to each local variable an arbitrary value from its domain. This action is always enabled. Observe that this allows a faulty node to send arbitrary messages. We assume, however, that messages sent by such node conform to the format specified by the algorithm: each message carries the specified number of values, and the values are drawn from appropriate domains. This assumption is not difficult to implement as message syntax checking logic can be incorporated in receive-action of each process. We assume *oral record* [3] of message transmission: the receiver can always correctly identify the message sender. The channels are reliable: the messages are delivered in FIFO order and without loss or corruption. Throughout the paper we assume that the maximum number of faults in the system is bounded by some constant k .

Graph exploration. The processes discover the topology of the system by exchanging messages. Each message contains the identifier of the process and its neighborhood. Process p *explored* process q if p received a message with (q, Q) . When it is clear from the context, we omit the mention of p . An *explored* subgraph of a graph contains only explored processes. A Byzantine process may potentially circulate information about the processes that do not exist in the system altogether. A process is *fake* if it does not exist in the system, a process is *real* otherwise.

3 Topology Discovery Problem: Statement and Solution Bounds

Problem statement.

Definition 1 (Weak Topology Discovery Problem). A program is a solution to the weak topology discovery problem if each of the program's computation satisfies the following properties: *termination* — either all non-faulty processes determine the system topology or at least one process detects a fault; *safety* — for each non-faulty process, the determined topology is a subset of the actual system topology; *validity* — the fault is detected only if there are faulty processes in the system.

Definition 2 (Strong Topology Discovery Problem). A program is a solution to the strong topology discovery problem if each of the program's computations satisfies the following properties: *termination* — all non-faulty processes determine the system topology; *safety* — the determined topology is a subset of the actual system topology.

According to the safety property of both problem definitions each non-faulty process is only required to discover a subset of the actual system topology. However, the desired objective is for each node to discover as much of it as possible. The following definitions capture this idea. A solution to a topology discovery problem is *complete* if every non-faulty process always discovers the complete topology of the system. A solution to the problem is *node-complete* if every non-faulty process discovers all nodes of the system. A solution is *adjacent-edge complete* if every non-faulty node discovers each edge adjacent to at least one non-faulty node. A solution is *two-adjacent-edge complete* if every non-faulty node discovers each edge adjacent to two non-faulty nodes.

Solution bounds. The proofs for the theorems stated in this section are to be found elsewhere [20].

Theorem 1. There does not exist a complete solution to the weak topology discovery problem.

Theorem 2. There exists no node- and adjacent-edge complete solution to the weak topology problem if the connectivity of the graph is lower or equal to the total number of faults k .

Observe that for $(k+1)$ -connected graphs an adjacent-edge complete solution is also node complete.

Theorem 3. There does not exist an adjacent-edge complete solution to the strong topology discovery problem.

Theorem 4. There exists no node- and two-adjacent-edge complete solution to the strong topology problem if the connectivity of the graph is less than or equal to twice the total number of faults k .

4 Detector

Outline. *Detector* solves the weak topology discovery problem for system graphs whose connectivity exceeds the number of faulty nodes k . The algorithm leverages the connectivity of the graph. For each pair of nodes, the graph guarantees the presence of at least one path that does not include a faulty node. The topology data travels along every path of the graph. Hence, the process that collects information about another process can find the potential inconsistency between the information that proceeds along the path containing faulty nodes and the path containing only non-faulty ones.

Care is taken to detect the fake nodes whose information is introduced by faulty processes. Since the processes do not know the size of the system, a faulty process may potentially introduce an infinite number of fake nodes. However, the graph connectivity assumption is used to detect fake nodes. As faulty processes are the only source of information about fake nodes, all the paths from the real nodes to the fake ones have to contain a faulty node. Yet, the graph connectivity is assumed to be greater than k . If a fake node is ever introduced, one of the non-faulty processes eventually detects a graph with too few paths leading to the fake node.

Detailed Description. The program is shown in Figure 1. Each process p stores the identifiers of its immediate neighbors. They are kept in set P . Each process keeps the upper bound k on the number of faulty processes. Process p maintains the following variables. Boolean variable *detect* indicates if p discovers a fault in the system. Boolean variable *start* guards the execution of the action that sends p 's neighborhood information to its neighbors. Set *TOP* (for topology) stores the subgraph explored by p ; *TOP* contains tuples of the form: (*process identifier, its neighborhood*). In the initial state, *TOP* contains (p, P) .

Function **path_number** evaluates the topology of the subgraph stored in *TOP*. Recall that a node u is unexplored by p if for every tuple $(s, S) \in TOP$, s is not the same as u . That is u may appear in S only. We construct graph G' by adding an edge to every pair of unexplored processes present in *TOP*. We calculate the value of **path_number** as follows. If the information of *TOP* is inconsistent, that is:

$$\begin{aligned} &(\exists u, v, U, V : ((u, U) \in TOP) \wedge ((v, V) \in TOP) : \\ &(u \in V) \wedge (v \notin U)) \end{aligned}$$

then **path_number** returns 0. If there is exactly one explored node in *TOP*, **path_number** returns $k + 1$. Otherwise the function returns the minimum number of internally node disjoint paths between two explored nodes in G' . In the correctness proof for this program we show that unless there is a fake node, the **path_number** of G' is no smaller than the connectivity of G .

Processes exchange messages of the form (*process identifier, its neighborhood id set*). A process contains two actions: *init* and *accept*. Action *init* starts the propagation of p 's neighborhood throughout the system. Action *accept* receives

```

process  $p$ 
const
   $P$ : set of neighbor identifiers of  $p$ 
   $k$ : integer, upper bound on the number of faulty processes
parameter
   $q : P$ 
var
   $detect$  : boolean, initially false, signals fault
   $start$  : boolean, initially true, controls sending of  $p$ 's neighborhood info
   $TOP$  : set of tuples, initially  $\{(p, P)\}$ , (process ids, neighbor id set)
  received by  $p$ 

  * [
     $init$ :       $start \longrightarrow$ 
                 $start := \mathbf{false},$ 
                 $(\forall j : j \in P : \mathbf{send} (p, P) \mathbf{to} j)$ 
  ]
  ]
   $accept$ :    [
                receive  $(r, R)$  from  $q \longrightarrow$ 
                if  $(\exists s, S : (s, S) \in TOP : s = r \wedge S \neq R) \vee$ 
                 $(\mathbf{path\_number}(TOP \cup \{(r, R)\}) < k + 1)$ 
                then
                   $detect := \mathbf{true}$ 
                else
                  if  $(\nexists s, S : (s, S) \in TOP : s = r)$  then
                     $TOP := TOP \cup \{(r, R)\},$ 
                     $(\forall j : j \in P : \mathbf{send} (r, R) \mathbf{to} j)$ 
                ]
  ]

```

Fig. 1. Process of *Detector*

the neighborhood data of some process, records it, checks against other data already available for p and possibly further disseminates the data. If the data received from neighbor q about a process r contradicts what p already holds about r in TOP or if the newly arrived information implies that G is less than $(k + 1)$ -connected p indicates that it detected a fault by setting $detect$ to **true**. Alternatively, if p did not previously have the information about r , p updates TOP and sends the received information to all its neighbors.

Theorem 5. *Detector* is an adjacent-edge complete solution to the weak topology discovery problem in case the connectivity of system topology graph exceeds the number of faults.

A correctness proof of the theorem can be found elsewhere [21].

Efficiency evaluation. Since we consider an asynchronous model, the number of messages a Byzantine process can send in a computation is infinite. To evaluate the efficiency of *Detector* we assume that each process is familiar with the upper

bound on the number of processes in the system and this upper bound is in $O(n)$. A non-faulty process then detects a fault if the number of processes it explores exceeds this bound or if it receives more than one identical message from the same neighbor. We assume that the process stops and does not send or receive any more messages if it detects a fault.

In this case we can estimate the number of messages that are received by non-faulty processes before one of them detects a fault or before the computation terminates. To make the estimation fair, we assume that the unit is $\log(n)$ bits. Since it takes that many bits to assign unique process identifiers to n processes, we assume that one identifier is exactly one unit of information. A message in *Detector* carries up to $\delta + 1$ identifiers, where δ is the maximum number of nodes in the neighborhood of a process. Observe that a process can receive at most n messages from each incoming channel. Thus, the total number of messages that can be sent by *Detector* is $2en$, where e is the number of edges in the graph. The message complexity of the program is in $O(2en\delta)$. If e is proportional to n^2 , then the complexity of the program is in $O(\delta n^3)$.

5 Explorer

Outline. The main idea of *Explorer* is for each process to collect information about some node's neighborhood such that the information goes along more than twice as many paths as the maximum number of Byzantine nodes. While the paths are node-disjoint, the information is correct if it comes across the majority of the paths. In this case the recipient is in possession of confirmed information. It turns out that the topology information does not have to come directly from the source. Instead it can come from processes with confirmed information. The detailed description of *Explorer* follows.

To simplify the presentation, we describe and prove correct the version of *Explorer* that tolerates only one Byzantine fault. We describe how this version can be extended to tolerate multiple faults in the end of the section.

Description. Since we first describe the 1-fault tolerant version of *Explorer* we assume that the graph is 3-connected. The program is shown in Figure 2. Similar to *Detector*, each process p in *Explorer*, stores the ids of its immediate neighbors. Process p maintains the variable *start*, whose function is to guard the execution of the action that initiates the propagation of p 's own neighborhood. Unlike *Detector*, however, p maintains two sets that store the topology information of the network: *uTOP* and *cTOP*. Set *uTOP* stores the topology data that is unconfirmed; *cTOP* stores confirmed topology data. Set *uTOP* contains the tuples of neighborhood information that p received from other nodes. Besides the process id and the set of its neighbor ids, each such tuple contains a set of process identifiers, that relayed the information. We call it *visited set*. The tuples in *cTOP* do not require visited set.

Processes exchange messages where, along with the neighbor identifiers for a certain process, a visited set is propagated. A process contains two actions:

init and *accept*. The purpose of *init* is similar to that in the process of *Detector*. Action *accept* receives the neighborhood information of some process r , its neighborhood R which was relayed by nodes in set S . The information is received from p 's neighbor — q .

First, *accept* checks if the information about r is already confirmed. If so, the only manipulation is to record the received information in *uTOP*. Actually, this update of *uTOP* is not necessary for the correct operation of the program, but it makes the its proof of correctness easier to follow.

If the received information does not concern already confirmed process, *accept* checks if this information differs from what is already recorded in *uTOP* either in r or in R . In either case the information is broadcast to all neighbors of p . Before broadcasting p appends the sender — q to the visited set S .

If the information about r and R has already been received and recorded in *uTOP*, *accept* checks if the previously recorded information came along an internally node disjoint path. If so, the information about r is added to *cTOP*. In this case, this information is also broadcast to all p 's neighbors. Note, however, that p is now sure of the information it received. Hence, the visited set of nodes in the broadcast message is empty.

Theorem 6. *Explorer* is a two-adjacent-edge complete solution to the strong topology discovery problem in case of one fault and the system topology graph is at least 3-connected.

A correctness proof of the theorem be found elsewhere [21].

Modification to Handle $k > 1$ faults. Observe that *Explorer* confirms the topology information about a node's neighborhood, when it receives two messages carrying it over internally node disjoint paths. Thus, the program can handle a single Byzantine fault. The explorer can handle $k > 1$ faults, if it waits until it receives $k + 1$ messages before it confirms the topology info. All the messages have to travel along internally node disjoint paths. For the correctness of the algorithm, the topology graph has to be $(2k + 1)$ -connected.

Proposition 1. *Explorer* is a two-adjacent-edge complete solution to the strong topology discovery problem in case of k faults and the system topology graph is at least $(2k + 1)$ -connected.

Efficiency evaluation. Unlike *Detector*, *Explorer* does not quit when a fault is discovered. Thus, the number of messages a faulty node may send is arbitrary large. However, we can estimate the message complexity of *Explorer* in the absence of faults. Each message carries a process identifier, a neighborhood of this process and a visited set. The number of the identifiers in a neighborhood is no more than δ , and the number of identifiers in the visited set can be as large as n . Hence the message size is bounded by $\delta + n + 1$ which is in $O(n)$.

Notice, that for the neighborhood A of each process a , every process broadcasts a message twice: when it first receives the information, and when it confirms it. Thus, the total number of sent messages is $4e \cdot n$ and the overall message complexity of *Explorer* if no faults are detected is in $O(n^4)$.

```

process  $p$ 
const
     $P$ , set of neighbor identifiers of  $p$ 
parameter
     $q : P$ 
var
     $start$  : boolean, initially true, controls sending of  $p$ 's neighbor ids
     $cTOP$  : set of tuples, initially  $\{(p, P)\}$ ,
        (process id, neighbor id set) confirmed topology info
     $uTOP$  : set of tuples, initially  $\emptyset$ ,
        (process id, neighbor id set, visited id set)
        unconfirmed topology info
    * [
init:     $start \longrightarrow$ 
             $start := \text{false},$ 
             $(\forall j : j \in P : \text{send } (p, P, \emptyset) \text{ to } j)$ 
    ]
    ]
accept:  ]
            receive  $(r, R, S)$  from  $q \longrightarrow$ 
                if  $(\forall t, T : (t, T) \in cTOP : t \neq r)$  then
                    if  $(\forall t, T, U : (t, T, U) \in uTOP : t \neq r \vee T \neq R)$  then
                         $(\forall j : j \in P : \text{send } (r, R, S \cup \{q\}) \text{ to } j)$ 
                    elsif  $(\exists t, T, U : (t, T, U) \in uTOP :$ 
                         $t = r \wedge R = T \wedge ((U \cap (S \cup \{q\}))) \subset \{r\})$ 
                    then
                         $cTOP := cTOP \cup \{(r, R)\},$ 
                         $(\forall j : j \in P : \text{send } (r, R, \emptyset) \text{ to } j)$ 
                         $uTOP := uTOP \cup \{(r, R, S \cup \{q\})\}$ 
    ]

```

Fig. 2. Process of *Explorer*

6 Composition and Extensions

Composing *Detector* and *Explorer*. Observe that *Detector* has better message complexity than *Explorer* if the neighborhood size is bounded. Hence, if the incidence of faults is low, it is advantageous to run *Detector* and invoke *Explorer* only if a fault is detected. We assume that the processes can distinguish between message types of *Explorer* and *Detector*. In the combined program, a process running *Detector* switches to *Explorer* if it discovers a fault. Other processes follow suit, when they receive their first *Explorer* messages. They ignore *Detector* messages henceforth. A Byzantine process may potentially send an *Explorer* message as well, which leads to the whole system switching to *Explorer*. Observe that if there are no faults, the system will not invoke *Explorer*. Thus, the complexity of the combined program in the absence of faults is the same as that of *Detector*. Notice that even though *Detector* alone only needs $(k+1)$ -connectivity of the system topology, the combined program requires $(2k+1)$ -connectivity.

Message Termination. We have shown that *Detector* and *Explorer* comply with the functional termination properties of the topology discovery problem. That is, all processes eventually discover topology. However, the performance aspect of termination, viz. message termination, is also of interest. Usually an algorithm is said to be message terminating if all its computations contain a finite number of sent messages [22].

However, a Byzantine process may send messages indefinitely. To capture this, we weaken the definition of message termination. We consider a Byzantine-tolerant program *message terminating* if the system eventually arrives at a state where: (a) all channels are empty except for the outgoing channels of a faulty process; (b) all actions in non-faulty processes are disabled except for possibly the receive-actions of the incoming channels from Byzantine processes, these receive-actions do not update the variables of the process. That is, in a terminating program, each non-faulty process starts to eventually discard messages it receives from its Byzantine neighbors.

Making *Detector* terminating is fairly straightforward. As one process detects a fault, the process floods the announcement throughout the system. Since the topology graph for *Detector* is assumed $(k+1)$ -connected, every process receives such announcement. As the process learns of the detection, it stops processing or forwarding of the messages. Notice that the initiation of the flood by a Byzantine node itself, only accelerates the termination of *Detector* as the other processes quickly learn of the faulty node's existence.

The addition of termination to *Explorer* is more involved. To ensure termination, restrictions have to be placed on message processing and forwarding. However, the restrictions should be delicate as they may compromise the liveness properties of the program.

By the design of *Explorer*, each process may send at most one message about its own neighborhood to its neighbors. Hence, the subsequent messages can be ignored. However, a faulty process may send messages about neighborhoods of other processes. These processes may be real or fake. We discuss these cases separately.

Note that each process in *Explorer* can eventually obtain an estimate of the identities of the processes in the system and disregard fake process information. Indeed, a path to a fake node can only lead through faulty processes. Thus, if a process discovers that there may be at most k internally node disjoint paths between itself and a certain node, this node is fake. Therefore, the process may cease to process messages about the fake node's neighborhood. Notice, that since the system is $(2k+1)$ -connected, messages about real nodes will always be processed. Therefore, the liveness properties of *Explorer* are not affected.

As to the real processes, they can be either Byzantine or non-faulty. Recall that each non-faulty process of *Explorer* eventually confirms neighborhoods of all other non-faulty processes. After the neighborhood of a process is confirmed, further messages about it are ignored.

The last case is a Byzantine process u sending a message to its correct neighbor v about the neighborhood of another Byzantine process w . By the design of *Explorer*, v relays the message about w provided that the neighborhood information about w differs from what previously received about w . As we discussed above, eventually v estimates the identities of all real processes in the system. Therefore, there is a finite number of possible different neighborhoods of w that u can create. Hence, eventually they will be exhausted, and v starts ignoring further messages from u about w .

Thus, *Explorer* can be made terminating as well.

Other extensions. Observe that *Explorer* is designed to disseminate the information about the complete topology to all processes in the system. However, it may be desirable to just establish the routes from all processes in the system to one or a fixed number of distinguished ones. To accomplish this *Explorer* needs to be modified as follows. No, neighborhood information is propagated. Instead of the visited set, each message carries the propagation path of the message. That is the order of the relays is significant.

Only the distinguished processes initiate the message propagation. The other processes only relay the messages. Just as in the original *Explorer*, a process confirms a path to another process only if it receives $2k + 1$ internally node disjoint paths from the source or from other confirming nodes. Again, like in *Explorer*, such process rebroadcasts the message, but empties the propagation path. In the outcome of this program, for every distinguished process, each non-faulty process will contain paths to at least $2k + 1$ processes that lead to this distinguished node. Out of these paths, at least $k + 1$ ultimately lead to the distinguished node.

In *Explorer*, for each process the propagation of its neighborhood information is independent of the other neighborhoods. Thus, instead of topology, *Explorer* can be used for efficient fault-tolerant propagation of arbitrary information from the processes to the rest of the network.

7 Conclusion

In conclusion, we would like to outline a couple of interesting avenues of further research.

The existence of Byzantine-robust topology discovery solutions opens the question of theoretical limits of efficiency of such programs. The obvious lower bound on message complexity can be derived as follows. Every process must transmit its neighborhood to the rest of the nodes in the system. Transmitting information to every node requires at least n messages, so the overall message complexity is at least δn^2 . If k processes are Byzantine, they may not relay the messages of other nodes. Thus, to ensure that other nodes learn about its neighborhood, each process has to send at least $k + 1$ messages. Thus, the complexity of any Byzantine-robust solution to the topology discovery problem is at least in $\Omega(\delta n^2 k)$.

Observe that *Explorer* and *Detector* may not explicitly identify faulty nodes or the inconsistent view of the their immediate neighborhoods. We believe that

this can be accomplished using the technique used by Dolev [7]. In case there are $3k + 1$ non-faulty processes, they may exchange the topologies they collected to discover the inconsistencies. This approach, may potentially expedite termination of *Explorer* at the expense of greater message complexity: if a certain Byzantine node is discovered, the other processes may ignore its further messages.

References

1. Spinelli, J.M., Gallager, R.G.: Event-driven topology broadcast without sequence numbers. *IEEE trans. on commun.* **COM-37**, 5 (1989) 468–474
2. Hill, J., Culler, D.: Mica: A wireless platform for deeply embedded networks. *IEEE Micro* **22**(6) (2002) 12–24
3. Lamport, L., Shostak, R., Pease, M.: The byzantine generals problem. *ACM Transactions on Programming Languages and Systems* **4**(3) (1982) 382–401
4. Avramopoulos, I.C., Kobayashi, H., Wang, R., Krishnamurthy, A.: Highly secure and efficient routing. In: *Proceedings of INFOCOM: The Conference on Computer Communications, joint conference of the IEEE Computer and Communications Societies, Hong Kong* (2004)
5. Perrig, A., Stankovic, J., Wagner, D.: Security in wireless sensor networks. *Communications of the ACM* **47**(6) (2004) 53–57
6. Bhandari, V., Vaidya, N.H.: On reliable broadcast in a radio network. In: *Proceedings of the Twenty-Fourth Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC 2005)*, Las Vegas, Nevada (2005) to appear
7. Dolev, D.: The Byzantine generals strike again. *Journal of Algorithms* **3**(1) (1982) 14–30
8. Koo, C.Y.: Broadcast in radio networks tolerating byzantine adversarial behavior. In: *PODC '04: Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, New York, NY, USA, ACM Press (2004) 275–282
9. Pelc, A., Peleg, D.: Broadcasting with locally bounded byzantine faults. *Information Processing Letters* **93** (2005) 109–115
10. Attiya, H., Welch, J.: *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. McGraw-Hill Publishing Company, New York (1998) 6.
11. Malkhi, D., Reiter, M., Rodeh, O., Sella, Y.: Efficient update diffusion in byzantine environments. In: *The 20th IEEE Symposium on Reliable Distributed Systems (SRDS '01)*, Washington - Brussels - Tokyo, IEEE (2001) 90–98
12. Malkhi, D., Mansour, Y., Reiter, M.K.: Diffusion without false rumors: on propagating updates in a Byzantine environment. *Theoretical Computer Science* **299**(1–3) (2003) 289–306
13. Minsky, Y., Schneider, F.B.: Tolerating malicious gossip. *Distributed Computing* **16**(1) (2003) 49–68
14. Masuzawa, T., Tixeuil, S.: A self-stabilizing link-coloring protocol resilient to unbounded byzantine faults in arbitrary networks. *Technical Report 1396, Laboratoire de Recherche en Informatique* (2005)
15. Nesterenko, M., Arora, A.: Tolerance to unbounded byzantine faults. In: *Proceedings of 21st IEEE Symposium on Reliable Distributed Systems*. (2002) 22–29

16. Sakurai, Y., Ooshita, F., Masuzawa, T.: A self-stabilizing link-coloring protocol resilient to byzantine faults in tree networks. In: Proceedings of the 2004 International Conference on Principles of Distributed Systems (OPODIS'2004). Lecture Notes in Computer Science, Springer-Verlag (2004)
17. Masuzawa, T.: A fault-tolerant and self-stabilizing protocol for the topology problem. In: Proceedings of the Second Workshop on Self-Stabilizing Systems. (1995) 1.1–1.15
18. Yellen, J., Gross, J.L.: Graph Theory & Its Applications. CRC Press (1998) ISBN: 0–849–33982–0.
19. Dijkstra, E.W., Scholten, C.S.: Predicate Calculus and Program Semantics. Springer-Verlag, Berlin (1990)
20. Nesterenko, M., Tixeuil, S.: Bounds on topology discovery in the presence of byzantine faults. Technical Report TR-KSU-CS-2006-01, Dept. of Computer Science, Kent State University (2006) <http://www.cs.kent.edu/techreps/TR-KSU-CS-2006-01.pdf>.
21. Nesterenko, M., Tixeuil, S.: Discovering network topology in the presence of byzantine faults. Technical Report TR-KSU-CS-2005-01, Dept. of Computer Science, Kent State University (2005) <http://www.cs.kent.edu/techreps/TR-KSU-CS-2005-01.pdf>.
22. Dijkstra, E., Scholten, C.: Termination detection for diffusing computations. Information Processing Letters **11**(1) (1980) 1–4