

Reliable Message Transmission under Partial Knowledge

Aris Pagourtzis¹, Giorgos Panagiotakos², and Dimitris Sakavalas¹

¹ School of Electrical and Computer Engineering
National Technical University of Athens, 15780 Athens, Greece,
pagour@cs.ntua.gr, sakaval@corelab.ntua.gr

² Department of Informatics and Telecommunications, University of Athens
Athens, Greece.
gpanagiotakos@corelab.ntua.gr

Abstract. A fundamental primitive in distributed computing is *Reliable Message Transmission* (RMT), which refers to the task of correctly sending a message from a party to another, despite the presence of byzantine corruptions. In this work we address the problem in the general adversary model of Hirt and Maurer, which subsumes earlier models such as the global or local threshold adversaries. Regarding the topology knowledge, we employ the recently introduced *Partial Knowledge Model* [13], which encompasses both the full knowledge and the *ad hoc* model; the latter assumes knowledge of the local neighborhood only.

Our main contributions are: (a) A necessary and sufficient condition for achieving RMT in the partial knowledge model with a general adversary; in order to show sufficiency, we propose RMT-PKA, a protocol that solves RMT whenever this is possible, therefore it is a *unique* protocol (cf. [14]). To the best of our knowledge, this is the first unique protocol for RMT in against general adversaries in the partial knowledge model. (b) A study of efficiency in the case of the *ad hoc* network model: we show that either the \mathcal{Z} -CPA protocol [13] is fully polynomial or no unique fully polynomial protocol for RMT exists, thus introducing a new notion of uniqueness with respect to efficiency that we call *poly-time uniqueness*.

To obtain our results we introduce, among others, a *joint view* operation on adversary structures, a new notion of separator (RMT-cut), appropriate for RMT in unreliable networks, and a self-reducibility property of the RMT problem, which we show by means of a protocol composition. The latter plays a crucial role in proving the poly-time uniqueness of \mathcal{Z} -CPA.

Eligibility for best student paper award: Yes.

1 Introduction

Achieving reliable communication in unreliable networks is fundamental in distributed computing. Of course, if there is an authenticated channel between two parties then reliable communication between them is guaranteed. However, it is often the case that certain parties are only indirectly connected, and need to use intermediate parties as relays to propagate their message to the actual receiver. The *Reliable Message Transmission* problem (RMT) is the problem of achieving correct delivery of a message m from a *dealer* (sender) D to a receiver R even if some of the intermediate nodes are corrupted and do not relay the message as agreed. In this work we consider the worst case corruption scenario, in which the adversary is unbounded and may control several nodes and be able to make them deviate from the protocol arbitrarily by blocking,

rerouting, or even altering a message that they should normally relay intact to specific nodes. An adversary with this behavior is referred to as *Byzantine adversary*.

The RMT problem has been initially considered by Dolev [2] in the context of the closely related *Reliable Broadcast* (Byzantine Generals) problem, introduced by Lamport, Shostak and Pease [10]. In Reliable Broadcast the goal is to achieve correct delivery of the dealer’s D message to all parties in the network.

The problem of message transmission under Byzantine adversaries has been studied extensively in various settings: secure or reliable transmission, general or threshold adversary, perfect or unconditional security. Here we focus on perfectly reliable transmission under a general adversary and the partial knowledge model. Note that the general adversary model, introduced by Hirt and Maurer [6], subsumes both the global [10] and the local threshold adversary model [8]. Regarding the topology knowledge, the recently introduced *Partial Knowledge Model* [13] encompasses both the full knowledge and the *ad hoc* (unknown topology) models.

The motivation for partial knowledge considerations comes from large scale networks (e.g. the Internet) where topologically local estimation of the power of the adversary may be possible, while global estimation may be hard to obtain due to geographical or jurisdiction constraints. Additionally, proximity in social networks is often correlated with an increased amount of available information, further justifying the relevance of the model.

1.1 Related work

The RMT problem under a threshold Byzantine adversary, where a fixed upper bound t is set for the number of corrupted players was addressed in [3, 1], where additional secrecy restrictions were posed and in [15] where a probability of failure was allowed. Results for RMT in the general adversary model [6], were given in [9, 17, 16]. In general, very few studies have addressed RMT or related problems in the partial knowledge setting despite the fact that this direction was already proposed in 2002 by Kumar *et al.* [9].

The approach that we follow here stems from a line of work which addresses the Reliable Broadcast problem with an honest dealer in incomplete networks, initiated by Koo [8]. Koo studied the problem in *ad hoc* networks under the *t-locally bounded adversary model*, in which at most a certain number t of corruptions are allowed in the neighborhood of every node. This work was generalized by Pelc, Peleg in [14] who pointed out how full knowledge of the topology yields better solvability results. After a series of works ([7, 11, 18]) tight conditions for the solvability of the problem were obtained in the *ad hoc* case. Finally, in [13] the Partial Knowledge Model was introduced, in which the players only have partial knowledge of the topology and the adversary structure. Tight algorithms for the extreme cases of full topology knowledge and *ad hoc* setting were also obtained in [13] and the results were extended to the general adversary case as well. Trivially all the aforementioned results for Reliable Broadcast with an honest dealer can be adapted for the RMT problem.

1.2 Our results

We study the RMT problem under general adversaries. Our contribution is twofold:

(a) Feasibility of RMT in the Partial Knowledge model. We prove a necessary and sufficient condition for achieving RMT in this setting, and present an algorithm that achieves RMT whenever this condition is met. In terminology of [14, 13] this is a *unique* algorithm for the problem, in the sense that whenever any safe algorithm achieves RMT in a certain instance so does our algorithm. This settles an open question of [13] and is, to the best of our knowledge, the first algorithm with this property.

A key notion that we define and use is the *joint adversary structure* of (a set of) players which corresponds to the worst case adversary structure that conforms to each player’s initial knowledge; this notion is crucial in obtaining the tight condition mentioned above. We also make use of the concept of local pair-cut technique, introduced by Pelc and Peleg [14] in the context of Broadcast. This technique was later [13] extended in order to obtain characterizations of classes of graphs for which Broadcast is possible for various levels of topology knowledge and type of corruption distribution. However, an exact characterization for the partial knowledge setting was left as an open question. Here we answer this question by proposing an adequate pair-cut for the partial knowledge model together with a unique algorithm for RMT, the first unique algorithm for this specific model proposed. This new algorithm is quite general and encompasses earlier algorithms such as CPA [8], PPA and \mathcal{Z} -CPA [13] as special cases. A useful by-product of practical interest is that the new cut notion can be used to determine the exact subgraph in which RMT is possible.

(b) Efficiency of RMT in the *Ad Hoc* network model, where each node’s knowledge over the topology and the adversary structure is limited in its neighborhood.

We observe that a variation of \mathcal{Z} -CPA [13] is unique for RMT as well, the first such algorithm we have encountered, and examine whether and when it is fully polynomial. We show that no unique fully polynomial protocol for RMT exists if \mathcal{Z} -CPA is not fully polynomial, thus introducing a new meaningful notion of *poly-time uniqueness*. In particular, we prove that if \mathcal{Z} -CPA is not fully polynomial in any class of instances where RMT is solvable, then there exists a corresponding class of (solvable) simpler instances in which any protocol that achieves RMT cannot be fully polynomial. We obtain this result by showing that \mathcal{Z} -CPA yields a polynomial time self-reduction for the RMT problem. Therefore \mathcal{Z} -CPA, despite its simplicity and minimal propagation, proves to be at least as efficient (in the sense described above) as any other RMT protocol.

More intuitively, we enhance the uniqueness property of \mathcal{Z} -CPA by implicitly stating that, not only one cannot achieve better solvability by employing more complex propagation schemes but one cannot even achieve significantly lower complexity in this way. This restriction seems to be inherent in the *ad hoc* network setting where players’ knowledge strictly relies on the information received by their neighbors.

1.3 Model and definitions

In this work we address the problem of Perfectly Reliable Message Transmission, hereafter simply referred as Reliable Message Transmission (RMT) under the influence of a general Byzantine adversary. In our model the players have partial knowledge of the network topology and of the adversary structure.

We assume a synchronous network represented by a graph G consisting of the player (node) set $V(G)$ and edge set $E(G)$ which represents authenticated channels between players. The problem definition follows.

Reliable Message Transmission. We assume the existence of a designated player D , called the *dealer*, who wants to propagate a certain value $x_D \in X$, where X is the initial message space, to a designated player R , called the receiver. We say that a distributed protocol achieves (or solves) RMT if by the end of the protocol the receiver R has *decided on* x_D , i.e. if it has been able to deduce that x_D is the value originally sent by the dealer and output it as the correct value.

The Adversary Model. The *general adversary model* was introduced by Hirt and Maurer in [6]. In this work they study the security of multiparty computation protocols with respect to an *adversary structure*, that is, a family of subsets of the players; the adversary is able to corrupt one of these subsets. More formally, a structure \mathcal{Z} for the set of players V is a monotone family of subsets of V , i.e. $\mathcal{Z} \subseteq 2^V$, where all subsets of a set Z are in \mathcal{Z} if $Z \in \mathcal{Z}$. In this work we obtain our results w.r.t. a general byzantine adversary, i.e., a general adversary which can make all the corrupted players deviate arbitrarily from the given protocol.

We assume that the instances of the problem contain the adversary structure \mathcal{Z} and say that an RMT protocol is *resilient* for an instance \mathcal{I} if it achieves RMT on instance \mathcal{I} for any possible corruption set. An RMT protocol which never causes the receiver R to decide on an incorrect value is called *safe*. Finally we say that an RMT protocol is *unique* if it solves RMT in all instances where RMT is solvable.

The Partial Knowledge Model [13]. In this setting each player v only has knowledge of the topology of a certain connected subgraph G_v of G which includes v . Namely if we consider the family \mathcal{G} of connected subgraphs of G we use the *view function* $\gamma : V \rightarrow \mathcal{G}$, where $\gamma(v)$ represents the subgraph over which player v has knowledge of the topology. We extend the domain of γ by allowing as input a set $S \subseteq G$. The output will correspond to the joint view of nodes in S . The extensively studied *ad hoc* model can be seen as a special case of the Partial Knowledge Model, where we assume that the topology knowledge of each player is limited to its own neighborhood, i.e., $\forall v \in V(G), \gamma(v) = \mathcal{N}(v)$.

Considering the partial knowledge model under the existence of a general adversary, we assume that given the actual adversary structure \mathcal{Z} each player v only knows the possible corruption sets in his view $\mathcal{Z}_v = \{A \cap V(\gamma(v)) \mid A \in \mathcal{Z}\}$ (*local adversary structure*).

Runs and Views. Given a run (execution) e of a distributed protocol, the *view*(v, e, k) of player v consists of the messages exchanged by v and its neighbors until round k . For simplification we will write *view*(v, e) to refer to all the messages exchanged by v and its neighbors until the end of the run e . With *view*(v, e, k) $|_A$ (and *view*(v, e) $|_A$) we will denote the corresponding messages exchanged by v and the set $A \subseteq \mathcal{N}(v)$. The decision of a player v in run e will be denoted by *decision* $_e$ (v); it is in fact completely determined by player's v view on run e . We will simply write *decision*(v) whenever the run is implied by the context.

In our study we will often make use of node-cuts (separators) which separate the receiver R from the dealer, hence, node-cuts that do not include the dealer. From here on we will simply use the term *cut* to denote such a separator.

2 Partial knowledge against a general adversary

Considering two players who have partial knowledge of the adversary, it would be useful to define an operation to calculate their joint knowledge about the adversary. Let $\mathcal{E}, \mathcal{F}, \mathcal{H}$ be adversary structures and A, B, C be sets of nodes. Let $\mathcal{E}^A = \{Z \cap A \mid Z \in \mathcal{E}\}$ denote the restriction of the adversary structure \mathcal{E} to the set of nodes A . The joint adversary structure from two restricted adversary structures can be obtained through the \oplus operator. We are interested on properties of this operation on different structures because the adversary can pretend having a different structure than the real one.

Definition 1. Let $\mathbb{T}^A = 2^{2^A}$ denote the space of adversary structures on the set of nodes A . Then operation \oplus is a function of the form $\oplus : \mathbb{T}^A \times \mathbb{T}^B \rightarrow \mathbb{T}^{(A \cup B)}$, for any A, B and is defined as follows:

$$\mathcal{E}^A \oplus \mathcal{F}^B = \{Z_1 \cup Z_2 \mid (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{F}^B) \wedge (Z_1 \cap B \subseteq Z_2) \wedge (Z_2 \cap A \subseteq Z_1)\}$$

Fact. An equivalent definition is $\mathcal{E}^A \oplus \mathcal{F}^B = \{Z_1 \cup Z_2 \mid (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{F}^B) \wedge (Z_1 \cap B = Z_2 \cap A)\}$

We show some algebraic properties of this operation in the Appendix. The next theorem shows the importance of the \oplus operation in this work.

Lemma 1. For any adversary structures \mathcal{E}, \mathcal{F} , node sets A, B and $\mathcal{H} = \mathcal{E}^A \oplus \mathcal{F}^B$, it holds that $\mathcal{H}^A = \mathcal{E}^A$ and $\mathcal{H}^B = \mathcal{F}^B$.

Theorem 1. For any adversary structures \mathcal{E}, \mathcal{F} , node sets A, B and $\mathcal{H} = \mathcal{E}^A \oplus \mathcal{F}^B$, it holds that $\forall \mathcal{H}' \in \mathbb{T}^{A \cup B} : \text{if } \mathcal{H}'^A = \mathcal{E}^A \text{ and } \mathcal{H}'^B = \mathcal{F}^B \text{ then } \mathcal{H}' \subseteq \mathcal{H}$.

Proof. Suppose that there existed some \mathcal{H}' s.t. $\exists Z \in \mathcal{H}' : Z \notin \mathcal{H}$. For Z we have $Z_1 = Z \cap A \in \mathcal{E}^A$ and $Z_2 = Z \cap B \in \mathcal{F}^B$. Also $Z_1 \cap B = (Z \cap A) \cap B \subseteq Z \cap B = Z_2$ and symmetrically $Z_2 \cap A \subseteq Z_1$. But then from definition $Z \in \mathcal{H}$ which is a contradiction and no such \mathcal{H}' exists.

Corollary 1. For any adversary structure \mathcal{E} and node sets A, B : $\mathcal{E}^{(A \cup B)} \subseteq \mathcal{E}^A \oplus \mathcal{E}^B$.

What Theorem 1 tells us is that the \oplus operation gives the maximal possible adversary structure that is indistinguishable between two agents that know \mathcal{E}^A and \mathcal{F}^B respectively, i.e., it coincides with their knowledge of the adversary structures on sets A and B respectively. Recall that $\mathcal{Z}_u = \mathcal{Z}^{V(\gamma(u))}$. The difference is that we use \mathcal{Z}_u to denote the knowledge of a player and $\mathcal{Z}^{V(\gamma(u))}$ to denote a restriction of the adversary structure. For a given adversary structure \mathcal{Z} and a view function γ let

$$\mathcal{Z}_B = \bigoplus_{v \in B} \mathcal{Z}^{V(\gamma(v))}$$

Then \mathcal{Z}_B exactly captures the maximal adversary structure possible, restricted in $\gamma(B)$, relative to the initial knowledge of players in B . Also notice that using corollary 1 we

get $\mathcal{Z}^{V(\gamma(B))} \subseteq \mathcal{Z}_B$. The interpretation of this inequality in our setting, is that what nodes in B conceive as the worst case adversary structure indistinguishable to them, it always contains the actual adversary structure in their scenario. For the rest of this work we will use $\mathcal{Z}^{\gamma(v)}$ and $\mathcal{Z}^{V(\gamma(v))}$ interchangeably.

3 Reliable message transmission in the partial knowledge model

In RMT we want the dealer D to send a message to some player R (the receiver) in the network. We assume that the dealer knows the id of player R . We denote an instance of the problem by the tuple $(G, \mathcal{Z}, \gamma, D, R)$. To analyze feasibility of RMT we introduce the notion of RMT-cut.

Definition 2. Let $(G, \mathcal{Z}, \gamma, D, R)$ be an RMT instance and $C = C_1 \cup C_2$ be a cut in G , partitioning $V \setminus C$ in two sets $A, B \neq \emptyset$ where $D \in A$ and $R \in B$. C is a RMT-cut iff $C_1 \in \mathcal{Z}$ and $C_2 \in \mathcal{Z}_B$.

Theorem 2 (Necessity). Let $(G, \mathcal{Z}, \gamma, D, R)$ be an RMT instance. If there exists a RMT-cut in G then no safe and resilient RMT algorithm exists for $(G, \mathcal{Z}, \gamma, D, R)$.

The proof adapts arguments from [14, 13] making use of the \oplus operation and is deferred to the Appendix.

3.1 The RMT-PKA protocol

We present the *RMT-Partial Knowledge Algorithm* (RMT-PKA), an RMT protocol which succeeds whenever the condition of Theorem 2 is met, rendering it a tight condition on when RMT is possible. To prove this we need to define what a valid scenario is for a set of messages that a node receives. In the definition below we refer to two types of messages, namely *type 1* and *type 2*, as defined in the description of the protocol.

Definition 3. Let M be a subset of the messages of type 1 and 2 that the receiver node R receives at some round of the protocol on $(G, \mathcal{Z}, \gamma, D, R)$. We say that M corresponds to a valid scenario for R if

- $\forall m_1, m_2 \in M$ of type 1, their first part is the same. That is, all messages of type 1 agree on the value sent to R .
- $\forall m_1, m_2 \in M$ of type 2, their first part is the same when they refer to the same node. That is, all messages of type 2 that refer to the same node v should contain exactly the same information, except possibly for the propagation path.

Notice that a valid scenario uniquely determines a dealer's value x_M (or \perp if no message of type 1 exists in M), possibly not the correct one (x_D). Additionally, a valid scenario uniquely determines a part of a possible world conforming to the messages in M : $(G_M, \mathcal{Z}_M, x_M)$. So \mathcal{Z}_M for example is the joint adversary structure of all the structures mentioned on the type 2 messages in M . In the same way G_M is the joint view of all the nodes mentioned on the first part of the type 2 messages in M .

Theorem 3. Let $(G, \mathcal{Z}, \gamma, D, R)$ be an RMT instance. If an RMT-cut exists, then RMT-PKA does not make the receiver decide on an incorrect value.

Proof. We will show that if an *RMT-cut* exists then R does not decide on any value. Suppose that R decided on some value $x' \neq x_D$. Then $\exists C \subseteq V(G) : \text{decision}(R|C) = x'$.

Let C_1 be the set of corrupted nodes, $C_1 \in \mathcal{Z}$. Notice that $C \cup C_1$ is a separator between R and D , or else R would receive the value x_D from a path of honest nodes not in C . Let $B' = \{v \mid v \in G \setminus C_1 \wedge \exists \text{ path } p \text{ in } G \setminus (C \cup C_1) \text{ from } v \text{ to } R\}$. Since C_1 is the actual corruption set, $(C_1 \cap V(\gamma(B'))) \in \mathcal{Z}_{B'}$. So $\text{decision}(R|C) = \perp$, which is a contradiction.

Protocol 1: *RMT-PKA*

Input (for each node v): dealer's label D , $\gamma(v)$, \mathcal{Z}_v .

Message format: type 1 : pair (x, p) or type 2 : pair $((u, \gamma(u), \mathcal{Z}_u), p)$, where $x \in X$ (message space), u the id of some node, $\gamma(u)$ is the view of node u , \mathcal{Z}_u is the adversary structure of node u , and p is a path of G (message's propagation trail).

Code for D : send messages $(\text{value} : x_D, \{D\})$ and $((D, \gamma(D), \mathcal{Z}_D), \{D\})$ to all neighbors and terminate.

Code for $v \notin \{D, R\}$: send message $((v, \gamma(v), \mathcal{Z}_v), \{v\})$ to all neighbors.

upon reception of type 1 or type 2 message (a, p) from node u do:

if $(v \in p) \vee (\text{tail}(p) \neq u)$ then discard the message else send $(a, p||v)$ ³ to all neighbours.

Code for R : upon reception of (x, p) from node u do:

if $\text{decision}(R) \neq \perp$ then decide on $\text{decision}(R)$ and terminate.

function $\text{decision}(v)$

if $\exists C \subseteq V(G)$ s.t. $\text{decision}(v|C) = x$ then return x
else return \perp .

function $\text{decision}(v|C)$

M : the set of all messages with valid format that have been received via paths that do not pass through nodes in C .

if M not valid (Def. 3) or $C \notin \mathcal{Z}_M$ then return \emptyset .

if $\exists C' \subset V(G_M) \setminus C$, s.t. C' is a cut between D, R on $G_M \setminus C$ ⁴, B' is the connected component that u lies in and $(C' \cap V(\gamma(B'))) \in (\mathcal{Z}_M)_{B'}$ then return \perp .

else return x .

Theorem 4 (Sufficiency). *Let $(G, \mathcal{Z}, \gamma, D, R)$ be an *RMT* instance. If no *RMT-cut* exists, then *RMT-PKA* achieves reliable message transmission.*

Proof. Suppose that no *RMT-cut* exists in G and *RMT-PKA* does not achieve reliable message transmission. Then R has not decided on x_D .

³ By $p||v$ we denote the path consisting of path p and node v , with the last node of p connected to v .

⁴ For simplicity, with $G \setminus C$ we denote the node induced subgraph of G on the node set $V(G) \setminus C$.

Let C_1 be the set of corrupted nodes, $C_1 \in \mathcal{Z}$. We will show that R should have decided on x_D and we will arrive at a contradiction.

First we show that $decision(R|C_1) = x_D$. Node R has received messages from $G \setminus C_1$, and the value sent from these paths is x_D . Node R only uses information that he has received from paths in $G \setminus C_1$ e.g. only correct information. Now since no *RMT*-cut exists in the graph, for every cut C_2 of $G \setminus C_1$ where B' is the connected component that R lies in, $C_2 \notin \mathcal{Z}_{B'}$ should hold. Otherwise $C_1 \cup C_2$ would be an *RMT*-cut. So $decision(R|C_1) = x_D$.

Now suppose that there exists a set C' s.t. $decision(R|C') \notin \{x_D, \emptyset\}$. This means that the view that R makes from the messages that come from paths not crossing C' is valid. Additionally $C' \cup C_1$ should form a cut on G between D and R , or else the scenario would not be valid.

Let $B' = \{v \mid v \in G \setminus C_1 \wedge \exists \text{ path } p \text{ in } G \setminus (C' \cup C_1) \text{ from } v \text{ to } R\}$. Since C' is a possible adversary set for R , and nodes in B' are considered honest and can communicate through a fully honest path with R , $C' \cap \gamma(B') \in \mathcal{Z}_{B'}$. Otherwise the scenario would not be valid, since C' could no be a possible corruption set.

So $C_1 \in \mathcal{Z}$ and $C' \cap \gamma(B') \in \mathcal{Z}_{B'}$ and $C_1 \cup C'$ is a *RMT*-cut, a contradiction, since there is no *RMT*-cut in G . Therefore no set C' exists s.t. $decision(R|C') \notin \{x_D, \emptyset\}$.

It follows that $decision(R) = x_D$, also a contradiction. So *RMT*-PKA achieves reliable message transmission if no *RMT*-cut exists in G .

Note on termination of RMT-PKA. Due to the uncertainty of nodes about the topology, the adversary can present them with a fake view of a huge graph and make them terminate whenever he wants. To avoid this, a strategy must be employed on the order that the receiver checks the information she gets from other players.

First, to restrict the power of the adversary, we allow every player to send at most b bytes at every round (where b is some constant). So if some player wants to send a message that exceeds this size, she splits it in parts of size b and sends it in successive rounds. Every player receives messages in a round robin fashion, that is in every round she receives a message only from one of her neighbors. Players forward messages, prioritizing messages that they received earlier. This way we can guarantee first, that R will get sufficient information to decide until some round t_0 independently of the behavior of the adversary, and secondly, that at time t_0 there is an upper bound on the information that R has received, namely $bt_0N(R)$.

With these modifications a simple strategy to achieve termination is to make R check the older scenarios first. So R first checks all possible scenarios using the information that she received until round 1, then until round 2, and so on. But by the way messages are forwarded as we argued before, there is a limited number of scenarios that have to be checked, namely the maximum number of scenarios that can be generated from $bt_0N(R)$ bytes. This way when *RMT* is possible, there exists an upper bound on the time the protocol takes to terminate, independent of the behavior of the adversary.

Corollary 2. (*RMT-PKA Protocol Uniqueness*) *Given an RMT instance $(G, \mathcal{Z}, \gamma, D, R)$, if there exists any safe and resilient RMT algorithm for this instance, then RMT-PKA also achieves reliable message transmission on this instance.*

4 RMT in *ad hoc* networks

In this section we consider the Reliable Message Transmission problem (*RMT*) in *ad hoc* networks. In the closely related problem of Reliable Broadcast the receiver is not a single node but instead the whole set $V(G)$. Reliable Broadcast in *ad hoc* networks under the influence of a general Byzantine adversary was initially studied in [13] where an algorithm for this model was presented and proven unique. The results can trivially be adapted to the case of the RMT problem.

4.1 *Ad hoc* RMT

An instance of the RMT problem in the *ad hoc setting* consists of a tuple (G, \mathcal{Z}, D, R) as explained in previous sections. In the closely related problem of Reliable Broadcast with an honest dealer [8, 14, 13], the notion of \mathcal{Z} -pp cut (definition given in the Appendix) was introduced in [13] and it was proved that a necessary and sufficient condition for the solvability of the problem is that a \mathcal{Z} -pp cut does not exist in the instance. Furthermore, the protocol \mathcal{Z} -CPA (Certified Propagation Algorithm) was given and proved that it achieves Broadcast in every instance where Broadcast is possible, i.e., it is *unique*.

Since in the RMT problem we are only concerned about the decision of the receiver R , we slightly modify the definition of the \mathcal{Z} -pp cut in order to capture an analogous cut (*RMT \mathcal{Z} -pp cut*) between the dealer D and the receiver R ,

Definition 4 (RMT \mathcal{Z} -pp cut). *Let C be a cut of G partitioning $V \setminus C$ into sets $A, B \neq \emptyset$ s.t. $D \in A$ and $R \in B$. C is an RMT \mathcal{Z} -pp cut if there exists a partition $C = C_1 \cup C_2$ with $C_1 \in \mathcal{Z}$ and $\forall u \in B, \mathcal{N}(u) \cap C_2 \in \mathcal{Z}_u$.*

The \mathcal{Z} -CPA algorithm can be trivially adapted for solving the RMT problem. In this algorithm the dealer first sends its initial value x_D to all its neighbors and terminates. After that the actions of any player v are defined as follows.

\mathcal{Z} -CPA code for v

1. If $v \in \mathcal{N}(D)$ then upon reception of x_D from the dealer, decide on x_D .
2. If $v \notin \mathcal{N}(D)$ then upon receiving the same value x from all neighbors in a set $N \subseteq \mathcal{N}(v)$ s.t. $N \notin \mathcal{Z}_v$, decide on value x .
3. If $v = R$ and decided on x then output decision x and terminate, else if $v \neq R$ and decided on x , send x to all neighbors $\mathcal{N}(v)$ and terminate.

Note that \mathcal{Z} -CPA is safe, in a sense that, never causes any honest player to decide on an incorrect value. Following an analysis identical to that of [13], where \mathcal{Z} -CPA is proven unique among safe Broadcast algorithms, we prove the uniqueness of \mathcal{Z} -CPA (modified as explained) among safe *RMT* algorithms. The following theorems are completely analogous with those of [13]; for completeness, the proofs are given in the appendix.

Theorem 5 (Sufficient Condition). *Given an RMT instance (G, \mathcal{Z}, D, R) , if no RMT \mathcal{Z} -pp cut exists on G , then \mathcal{Z} -CPA achieves RMT in (G, \mathcal{Z}, D, R) .*

Theorem 6 (Necessary Condition). *Given an RMT instance (G, \mathcal{Z}, D, R) , if an RMT \mathcal{Z} -pp cut exists on G then no safe RMT algorithm exists for (G, \mathcal{Z}, D, R) .*

Thus, \mathcal{Z} -CPA, the first algorithm we have encountered for RMT in generic topology *ad hoc* networks against general adversaries, proves to be unique.

5 Protocol uniqueness with respect to efficiency

Up to now we have seen that \mathcal{Z} -CPA is unique among the safe *ad hoc* RMT algorithms. In terms of efficiency it is interesting to study whether \mathcal{Z} -CPA is also the more efficient one among unique RMT algorithms w.r.t. polynomial time. We will measure protocol complexity with respect to the size of the graph $|G| = n$ only, because we are mainly interested in protocols that are *fully polynomial* (of polynomial round, bit and local computations complexity) regardless of the size of the adversary structure description. Note however that, if the adversary structure is given explicitly, \mathcal{Z} -CPA is trivially *fully polynomial* w.r.t. the total size of the input.

Definition 5 (Poly-time Uniqueness). *We call a protocol \mathcal{A} poly-time unique for problem Π if it is unique (with respect to feasibility) and the existence of a unique fully polynomial protocol for Π implies that \mathcal{A} is also fully polynomial for Π .*

In other words, either \mathcal{A} is fully polynomial on all solvable instances or no fully polynomial protocol that solves Π on all solvable instances exists.

In the main theorem of this section we prove that \mathcal{Z} -CPA is poly-time unique for the RMT problem, and thus show that \mathcal{Z} -CPA is at least as efficient (in the sense described above) as any other RMT protocol. To show that, we build a self reduction for RMT based on \mathcal{Z} -CPA.

Under this approach, we need to take into account the complexity of membership check for \mathcal{Z} , due to the second decision rule of \mathcal{Z} -CPA. Indeed, we may observe that \mathcal{Z} -CPA is fully polynomial if the membership check for \mathcal{Z} can be performed in polynomial time w.r.t. $|G|$. In the sequel, we essentially show that if a unique fully polynomial RMT algorithm exists, it must be able to answer the membership question for \mathcal{Z} in polynomial time w.r.t. $|G|$ and therefore can be used as a subroutine to make \mathcal{Z} -CPA fully polynomial.

5.1 Self-reducibility of RMT

Consider the family of instances \mathcal{G} where achieving RMT is possible. By Theorems 5,6:

$$\mathcal{G} = \{(G, \mathcal{Z}, D, R) \mid \nexists RMT \mathcal{Z}\text{-pp cut in } G\}$$

Also consider the family of *basic* instances $\mathcal{G}' \subseteq \mathcal{G}$ which contains the tuples (G, \mathcal{Z}, D, R) where G is of the form shown in Figure 1 and RMT is solvable. More specifically, G contains the two distinguished nodes D, R and a “middle set” which we call $A(G)$. The only edges appearing are those which connect each player in the set $A(G)$ with the dealer D and the receiver-node R and in the resulting graph there does not exist a *RMT \mathcal{Z} -pp cut*. Finally for any $\mathcal{G}_1 \subseteq \mathcal{G}$ we define the family of instances $\mathcal{I}(\mathcal{G}_1) \subseteq \mathcal{G}'$ which consists of all the instances $(G', \mathcal{Z}', D', R') \in \mathcal{G}'$ such that graph’s G' middle-set $A(G')$ is a subset of a neighborhood of a node v in a graph contained in family \mathcal{G}_1 , as a part of the instance tuple (G, \mathcal{Z}, D, R) , and $\mathcal{Z}' = \mathcal{Z}_v$ ⁵. More precisely,

$$\mathcal{I}(\mathcal{G}_1) = \{(G', \mathcal{Z}', D', R') \in \mathcal{G}' \mid \exists (G, \mathcal{Z}, D, R) \in \mathcal{G}_1, \exists v \in V(G) \setminus \{D\} \text{ s.t. } A(G') \subseteq \mathcal{N}(v), \mathcal{Z}' = \mathcal{Z}_v\}$$

⁵ In this point we slightly abuse the terminology, for ease of exposition, and use $\mathcal{Z}' = \mathcal{Z}_v$ instead of $\mathcal{Z}' = \{S \cap A(G') \mid S \in \mathcal{Z}_v\}$. The second statement is more accurate in the case where $A(G') \subsetneq \mathcal{N}(v)$ because we defined \mathcal{Z} as a subset of the powerset of the nodes in the instance. This however does not affect our study because we can add the extra nodes $\mathcal{N}(v) \setminus A(G')$ in our instance $(G', \mathcal{Z}', D', R')$ as isolated nodes.




Fig. 1. Family of instances \mathcal{G}' . No RMT \mathcal{Z} -pp cut exists.

Intuitively the above family consists of the decomposition of every graph G in the \mathcal{G}_1 family into “small” graphs of the family \mathcal{G}' whose middle sets appear in G as (partial) neighborhoods of nodes, the adversary structures are subsets of the original structure, and the RMT problem is solvable.

We next show that the RMT problem in any family of instances $\mathcal{G}_1 \subseteq \mathcal{G}$ (denoted $RMT|_{\mathcal{G}_1}$), also referred to as the RMT problem with *promise set* \mathcal{G}_1 (cf. [5]), reduces in polynomial time w.r.t. the size of the graph n to the $RMT|_{\mathcal{I}(\mathcal{G}_1)}$ problem. That is, if there exists an algorithm for solving RMT in $\mathcal{I}(\mathcal{G}_1)$ in fully polynomial time it can be used, as a subroutine of \mathcal{Z} -CPA, to solve RMT in \mathcal{G}_1 in fully polynomial time.

Theorem 7. *If there exists a fully polynomial (in n) algorithm Π for solving $RMT|_{\mathcal{I}(\mathcal{G}_1)}$ then there exist a fully polynomial algorithm (in n) that solves $RMT|_{\mathcal{G}_1}$.*

Proof. We will use \mathcal{Z} -CPA to solve $RMT|_{\mathcal{G}_1}$. \mathcal{Z} -CPA has been proven unique, i.e., solves RMT in all instances where it is solvable, hence also for the family of instances \mathcal{G}_1 that we consider in this theorem.

We will show that \mathcal{Z} -CPA with protocol Π as a subroutine yields a fully polynomial algorithm for $RMT|_{\mathcal{G}_1}$. Namely, the decision rule of \mathcal{Z} -CPA which consists of a membership check for \mathcal{Z}_v will be answered through simulations of protocol Π in time $poly(n)$. Since the subroutine protocol Π will only be used in the local computations phase of \mathcal{Z} -CPA, the round and bit complexity of \mathcal{Z} -CPA will be maintained in the resulting algorithm.

First, from the description of \mathcal{Z} -CPA observe that the *round complexity* is linear in n because at least one new player decides in every round and each player terminates after decision. Thus the receiver R will decide in at most n rounds. Second, one can see that the *bit complexity* of \mathcal{Z} -CPA is also of order $poly(n)$ due to the fact that each player sends one message to all of its neighbors. For deducing the latter we can reasonably assume that the messages sent by honest players are of size $poly(n)$ or, to drop any such assumption, consider the space X of the messages exchanged as a part of the input of size n . It thus remains to show that in \mathcal{Z} -CPA, the *local computations complexity*, can be of order $poly(n)$ if we use Π as subroutine.

For an arbitrary run e of \mathcal{Z} -CPA in some instance of \mathcal{G}_1 , we can define $\mathcal{D}(i)$ to be the set of players that decide in round i of \mathcal{Z} -CPA. Moreover since run e is on an instance in the family $\mathcal{G}_1 \subseteq \mathcal{G}$, i.e., the RMT problem is solvable, it should be the case that $\exists i \in \{1, \dots, n\}, R \in \mathcal{D}(i)$. Observe that the function \mathcal{D} is well defined as we can assume that we use an arbitrary algorithm, e.g. exhaustive search, to answer the membership check for \mathcal{Z}_v (possibly in exponential time).

We next show that if we use Π as a subroutine for the local computations of the run e of \mathcal{Z} -CPA, we can achieve RMT in time $\text{poly}(n)$. Namely, we show by induction that for every round i , each player $v \in \mathcal{D}(i)$ will decide in $\text{poly}(n)$. Since $\exists i \in \{1, \dots, n\}, R \in \mathcal{D}(i)$ RMT will be achieved.

For round $i = 1$ all $v \in \mathcal{N}(D)$ receive the dealer's value x_D from the dealer and trivially decide on it in $\text{poly}(n)$ time.

Assume that, for every round $i \leq k$ every $v \in \mathcal{D}(i)$ decides in $\text{poly}(n)$ -time. Considering any $v \in \mathcal{D}(k+1)$ and the \mathcal{Z} -CPA message propagation, the latter means that by the end of round k , v will have received sufficient information $\text{view}(v, e, k)$ to decide, from players in $\bigcup_{i=1, \dots, k} \mathcal{D}(i)$, in $\text{poly}(n)$ -time, i.e., v will have received the same value x from all its neighbors in a set $N \subseteq \mathcal{N}(v)$ s.t. $N \notin \mathcal{Z}_v$. All valid messages exchanged in \mathcal{Z} -CPA consist of a single value $x \in X$ which corresponds to a possible dealer's value, and each player transmits only once to all its neighbors. Messages of different form, which we call *erroneous*, can be recognized by the recipient in $\text{poly}(n)$ time since $|X| = \text{poly}(n)$. Given $\text{view}(v, e, k)$, player v , in $\text{poly}(n)$ -time, can create a partition of

its neighborhood $\mathcal{N}(v) = \bigcup_{i=0}^{m+1} A_i$ such that

$$A_0 = \{u \in \mathcal{N}(v) \mid u \text{ sent nothing}\}$$

$$A_i = \{u \in \mathcal{N}(v) \mid u \text{ sent value } a_i \in X\}, \quad i = \{1, \dots, m\}$$

$$A_{m+1} = \{u \in \mathcal{N}(v) \mid u \text{ sent erroneous messages}\}$$

Since sets A_0, A_{m+1} do not affect our study we let $A = \bigcup_{i=\{1, \dots, m\}} A_i$. Denote with $H, Z \subseteq V$ the sets of actual honest and corrupted players of run e . Also consider the sets of honest and corrupted neighbors of v , $H_v = H \cap \mathcal{N}(v)$ and $Z_v = Z \cap \mathcal{N}(v)$ respectively. Given $\text{view}(v, e, k)$, observe that

$$\exists! h \in \{1, \dots, m\} \text{ s.t. } H_v \setminus A_0 \subseteq A_h$$

else there exists an honest player which sends an incorrect value, a contradiction because \mathcal{Z} -CPA is safe. Subsequently $Z_v \supseteq A \setminus A_h$. Note that all $u \in A_h$ transmit the correct value a_h (regardless of whether they are honest or not) and all $u \in A \setminus A_h$ transmit false values. Since, by assumption, $\text{view}(v, e, k)$ is sufficient for v to decide through \mathcal{Z} -CPA, it holds that $A_h \notin \mathcal{Z}_v$ due to the decision rule of \mathcal{Z} -CPA. Moreover $\forall i \in \{1, \dots, m\} \setminus \{h\}$ it holds that $A_i \in \mathcal{Z}_v$ since $A \setminus A_h \subseteq Z_v$. Consequently

$$\exists! h \in \{1, \dots, m\} \text{ s.t. } A_h \notin \mathcal{Z}_v \text{ and } A \setminus A_h \in \mathcal{Z}_v \quad (1)$$

We next show how player v can decide which is the actual value of h in $\text{poly}(n)$ time using the protocol Π , and thus decide on the correct value a_h .

For $l = 1, \dots, m$, we define the following runs of Π that can be simulated by v .

- Run e_0^l is on the instance $(G, \mathcal{Z}_v, D, v) \in \mathcal{G}'$ with $V(G) = A \cup \{D\} \cup \{v\}$, dealer's value $x_D = 0$, and corruption set $Z_v = A \setminus A_l$; in each round, all players in Z_v send the messages that send in the respective round of run e_1^l (where $A \setminus A_l$ is a set of honest players which runs Π). The latter means that v exchanges with Z_v messages that consist the $\text{view}(e_1^l, v)|_{A \setminus A_l}$.
- Run e_1^l , is on the same graph G , with dealer's value $x_D = 1$, and corruption set $Z_v = A_l$; Analogously with e_0 player v exchanges with Z_v the messages $\text{view}(e_0^l, v)|_{A_l}$.




Fig. 2. Runs e_0 and e_1 .

Player v simulates run e_1^l in order to determine the behavior of the corrupted players in e_0^l . Observe that for every l exactly one of e_0^l, e_1^l is not in the family of instances $\mathcal{I}(\mathcal{G}_1)$ (due to the selection of the corruption set) and thus the local computations complexity might not be polynomial. Since protocol Π is fully polynomial in $\mathcal{I}(\mathcal{G}_1)$, it means that there is an explicit bound B on the local computations complexity of Π in the family $\mathcal{I}(\mathcal{G}_1)$. Assuming that arbitrary player v knows such a bound⁶ we modify the above runs such that if the local computations complexity of a player w in a round i of e_0^l or e_1^l exceeds the bound B then v halts the simulation of the round i local computations of w and sends nothing on behalf of player w in round i . Such a modification of the run is necessary to obtain the desired result.

Player v runs the following protocol in order to decide on the value of the dealer of run e .

(*Decision Protocol.*) Player v simulates, in parallel, $2m = \text{poly}(n)$ runs $(e_0^l, e_1^l)_{l \in \{1, \dots, m\}}$ and halts all parallel simulations with decision a_l if run e_0^l terminates with $\text{decision}(v) = 0$.

We next show that v terminates run e_0^l with $\text{decision}(v) = 0$ if and only if $A_l \notin \mathcal{Z}_v$. More concretely

$$A_l \notin \mathcal{Z}_v \Leftrightarrow \text{decision}_{e_0^l}(v) = 0$$

“ \Rightarrow ”: $A_l \notin \mathcal{Z}_v \Rightarrow Z_v = A \setminus A_l \in \mathcal{Z}_v$. Since by assumption Π solves $\text{RMT}|_{\mathcal{I}(\mathcal{G}_1)}$, for any adversarial behavior, that of Z_v in e_0^l included, v will decide on the correct value $x_D = 0$, i.e., $\text{decision}_{e_0^l}(v) = 0$.

“ \Leftarrow ”: Let $A_l \in \mathcal{Z}_v$ and $\text{decision}_{e_0^l}(v) = 0$. This by equation (1) means that $Z_v = A \setminus A_l \notin \mathcal{Z}_v$. Observe now that the run e_0^l is not a valid run for the instance (G, \mathcal{Z}_v, D, v) because the adversarial behavior of $Z_v \notin \mathcal{Z}_v$ is not valid for the adversary structure \mathcal{Z}_v . But the view of v is the same as the valid run e_1^l in which $x_D = 1$ and $Z_v = A_l \in \mathcal{Z}_v$. Since Π solves $\text{RMT}|_{\mathcal{I}(\mathcal{G}_1)}$, for any adversarial behavior, that of Z_v in e_1^l included, v will decide on the correct value $x_D = 1$ in the run e_1^l i.e., $\text{decision}_{e_1^l}(v) = 1$. But since the decision

⁶ Although this assumption is natural and often used, it is possible to avoid it if we consider family $\mathcal{I}(\mathcal{G}_1)$ consisting of directed graphs (with edges from dealer to $A(G)$ and from $A(G)$ to v). In this case the view of all players in $A_l, A \setminus A_l$ would be the same as that of some run in $\mathcal{I}(\mathcal{G}_1)$ and thus their local computations complexity would be polynomial.

is a function of the view and player v receives exactly the same messages in runs e_0^l, e_1^l , it holds that $\text{decision}_{e_0^l}(v) = 1$, a contradiction.

The latter shows that the decision of player v in run e , which is acquired through the *Decision Protocol*, is correct and uniquely defined. Moreover all parallel simulations halt when the simulated run $e_0^l, l = h$ of Π terminates. Thus we have to show that run e_0^h can be simulated in polynomial time.

The problem is that run e_1^h , which is simulated to determine the behavior of the corrupted players in e_0^l , is not a run of *RMT* in the family $\mathcal{I}(\mathcal{G}_1)$ due to the selection of the corruption set. Therefore we lose guarantee of full polynomiality in that run. Non-polynomiality of the round complexity is not an obstacle since the simulations are done in parallel. Local computations' polynomial complexity of e_1^h is ensured by the fact that we halt any local computations that exceed the explicit bound B previously mentioned. Finally it is easy to see that the bit complexity of the simulated runs is polynomial if the round and local computations complexity is polynomial. Thus the simulated run e_0^l remains fully polynomial.

Therefore it follows that v will decide in run e in polynomial time because the simulation of a fully polynomial protocol can be done in polynomial time.

6 Open questions

Regarding the partial knowledge model, *RMT-PKA* makes players exchange information about the topology. Although topology discovery was not our motive, techniques used here (e.g. the \oplus operation) may be applicable to that problem under a Byzantine adversary ([12],[4]). A comparison with the techniques used in this field might give further insight on how to efficiently extract information from maliciously crafted topological data.

The unique protocol proposed for the partial knowledge model only answers the feasibility question. A natural question is whether and when we can devise a unique and also efficient algorithm for this setting. The techniques used so far in the bibliography to reduce the communication complexity [9] do not seem to be directly applicable to this model. So exploring this direction might give new insights on message delivery in partially known graphs.

It would also be interesting to argue about uniqueness with respect to efficiency for *RMT* in the partial knowledge model by extending our analysis of the *ad hoc* case.

Finally, it is possible to define a stronger type of poly-time uniqueness: we call a protocol \mathcal{A} *strongly poly-time unique* for problem Π if the existence of any fully-polynomial protocol for a class of instances \mathcal{I} implies that \mathcal{A} is also fully polynomial for all instances in \mathcal{I} . We conjecture that \mathcal{Z} -CPA is in fact strongly poly-time unique for *RMT* in the *ad hoc* model.

References

1. Yvo Desmedt and Yongge Wang. Perfectly secure message transmission revisited. In LarsR. Knudsen, editor, *Advances in Cryptology EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 502–517. Springer Berlin Heidelberg, 2002.
2. Danny Dolev. The byzantine generals strike again. *J. Algorithms*, 3(1):14–30, 1982.

3. Danny Dolev, Cynthia Dwork, Orli Waarts, and Moti Yung. Perfectly secure message transmission. *J. ACM*, 40(1):17–47, January 1993.
4. Shlomi Dolev, Omri Liba, and Elad Michael Schiller. Self-stabilizing byzantine resilient topology discovery and message delivery - (extended abstract). In Vincent Gramoli and Rachid Guerraoui, editors, *NETYS*, volume 7853 of *Lecture Notes in Computer Science*, pages 42–57. Springer, 2013.
5. Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, New York, NY, USA, 1 edition, 2008.
6. Martin Hirt and Ueli M. Maurer. Complete characterization of adversaries tolerable in secure multi-party computation (extended abstract). In James E. Burns and Hagit Attiya, editors, *PODC*, pages 25–34. ACM, 1997.
7. Akira Ichimura and Maiko Shigeno. A new parameter for a broadcast algorithm with locally bounded byzantine faults. *Inf. Process. Lett.*, 110(12-13):514–517, 2010.
8. Chiu-Yuen Koo. Broadcast in radio networks tolerating byzantine adversarial behavior. In Soma Chaudhuri and Shay Kutten, editors, *PODC*, pages 275–282. ACM, 2004.
9. M V N Ashwin Kumar, Pranava R. Goundan, K Srinathan, and C. Pandu Rangan. On perfectly secure communication over arbitrary networks. In *Proceedings of the twenty-first annual symposium on Principles of distributed computing*, PODC '02, pages 193–202, New York, NY, USA, 2002. ACM.
10. Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
11. Chris Litsas, Aris Pagourtzis, and Dimitris Sakavalas. A graph parameter that matches the resilience of the certified propagation algorithm. In Jacek Cichon, Maciej Gebala, and Marek Klonowski, editors, *ADHOC-NOW*, volume 7960 of *Lecture Notes in Computer Science*, pages 269–280. Springer, 2013.
12. Mikhail Nesterenko and Sébastien Tixeuil. Discovering network topology in the presence of byzantine faults. *IEEE Trans. Parallel Distrib. Syst.*, 20(12):1777–1789, 2009.
13. Aris Pagourtzis, Giorgos Panagiotakos, and Dimitris Sakavalas. Reliable broadcast with respect to topology knowledge. In Fabian Kuhn, editor, *Distributed Computing - 28th International Symposium, DISC 2014, Austin, TX, USA, October 12-15, 2014. Proceedings*, volume 8784 of *Lecture Notes in Computer Science*, pages 107–121. Springer, 2014.
14. Andrzej Pelc and David Peleg. Broadcasting with locally bounded byzantine faults. *Inf. Process. Lett.*, 93(3):109–115, 2005.
15. Bhavani Shankar, Prasant Gopal, Kannan Srinathan, and C. Pandu Rangan. Unconditionally reliable message transmission in directed networks. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '08, pages 1048–1055, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
16. Kannan Srinathan, Arpita Patra, Ashish Choudhary, and C. Pandu Rangan. Unconditionally secure message transmission in arbitrary directed synchronous networks tolerating generalized mixed adversary. In *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security*, ASIACCS '09, pages 171–182, New York, NY, USA, 2009. ACM.
17. Kannan Srinathan and C. Pandu Rangan. Possibility and complexity of probabilistic reliable communication in directed networks. In Eric Ruppert and Dahlia Malkhi, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing, PODC 2006, Denver, CO, USA, July 23-26, 2006*, pages 265–274. ACM, 2006.
18. Lewis Tseng, Nitin Vaidya, and Vartika Bhandari. Broadcast using certified propagation algorithm in presence of byzantine faults. *Information Processing Letters*, 115(4):512 – 514, 2015.

Appendix

A The \oplus operation

Theorem 8. *Operator \oplus is commutative.*

Proof. A binary operation $*$ is called commutative if $a * b = b * a$. For any adversary structures \mathcal{E}, \mathcal{F} and node sets A, B :

$$\begin{aligned}\mathcal{E}^A \oplus \mathcal{F}^B &= \{Z_1 \cup Z_2 \mid (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{F}^B) \wedge (Z_1 \cap B \subseteq Z_2) \wedge (Z_2 \cap A \subseteq Z_1)\} \\ &= \{Z_2 \cup Z_1 \mid (Z_2 \in \mathcal{F}^B) \wedge (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \cap A \subseteq Z_1) \wedge (Z_1 \cap B \subseteq Z_2)\} \\ &= \mathcal{F}^B \oplus \mathcal{E}^A\end{aligned}$$

So operator \oplus is commutative.

To prove that \oplus is also associative we will need the following lemma.

Lemma 2. *For any node sets A, B, C it holds that*

$$\begin{aligned}(Z_1 \cap B \subseteq Z_2) \wedge (Z_2 \cap A \subseteq Z_1) \wedge (Z_1 \cup Z_2 \cap C \subseteq Z_3) \wedge (Z_3 \cap A \cup B \subseteq Z_1 \cup Z_2) \\ \Leftrightarrow \\ (Z_2 \cap C \subseteq Z_3) \wedge (Z_3 \cap B \subseteq Z_2) \wedge (Z_2 \cup Z_3 \cap A \subseteq Z_1) \wedge (Z_1 \cap B \cup C \subseteq Z_2 \cup Z_3)\end{aligned}$$

Proof. First we prove the \Rightarrow direction. From $(Z_1 \cup Z_2 \cap C \subseteq Z_3)$ it follows that:

$$\begin{aligned}(Z_1 \cup Z_2) \cap C \subseteq Z_3 &\Rightarrow (Z_1 \cap C) \cup (Z_2 \cap C) \subseteq Z_3 \\ &\Rightarrow (Z_1 \cap C) \subseteq Z_3 \wedge (Z_2 \cap C) \subseteq Z_3\end{aligned}$$

From $(Z_3 \cap (A \cup B) \subseteq Z_1 \cup Z_2)$ it follows that:

$$\begin{aligned}Z_3 \cap (A \cup B) \subseteq Z_1 \cup Z_2 &\Rightarrow (Z_3 \cap A) \cup (Z_3 \cap B) \subseteq Z_1 \cup Z_2 \\ &\Rightarrow (Z_3 \cap B) \subseteq Z_1 \cup Z_2 \\ &\Rightarrow (Z_3 \cap B) \cap B \subseteq (Z_1 \cup Z_2) \cap B \\ &\Rightarrow (Z_3 \cap B) \subseteq (Z_1 \cap B) \cup (Z_2 \cap B) \\ &\Rightarrow (Z_3 \cap B) \subseteq (Z_2 \cap B) \\ &\Rightarrow (Z_3 \cap B) \subseteq Z_2\end{aligned}$$

$$\begin{aligned}Z_3 \cap (A \cup B) \subseteq Z_1 \cup Z_2 &\Rightarrow (Z_3 \cap A) \cup (Z_3 \cap B) \subseteq Z_1 \cup Z_2 \\ &\Rightarrow (Z_3 \cap A) \subseteq Z_1 \cup Z_2 \\ &\Rightarrow (Z_3 \cap A) \subseteq Z_1 \cup Z_2 \\ &\Rightarrow (Z_3 \cap A) \cap A \subseteq (Z_1 \cup Z_2) \cap A \\ &\Rightarrow (Z_3 \cap A) \subseteq (Z_1 \cap A) \cup (Z_2 \cap A) \\ &\Rightarrow (Z_3 \cap A) \subseteq (Z_2 \cap A) \\ &\Rightarrow (Z_3 \cap A) \subseteq Z_2\end{aligned}$$

Also :

$$\begin{aligned} (Z_2 \cup Z_3) \cap A &\subseteq (Z_2 \cap A) \cup (Z_3 \cap A) \\ &\subseteq Z_1 \cup Z_1 \\ &\subseteq Z_1 \end{aligned}$$

And

$$\begin{aligned} (Z_1 \cap (B \cup C)) &\subseteq (Z_1 \cap B) \cup (Z_1 \cap C) \\ &\subseteq Z_2 \cup Z_3 \end{aligned}$$

The proof for the \Rightarrow direction is complete. The other direction follows from symmetry.

Theorem 9. *Operator \oplus is associative.*

Proof. A binary operation $*$ is called associative if $(a * b) * c = a * (b * c)$ for any well defined a, b, c . For any adversary structures $\mathcal{E}, \mathcal{F}, \mathcal{H}$ and node sets A, B, C :

$$\begin{aligned} (\mathcal{E}^A \oplus \mathcal{F}^B) \oplus \mathcal{H}^C &= \{Z_1 \cup Z_2 | (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{F}^B) \wedge (Z_1 \cap B \subseteq Z_2) \wedge (Z_2 \cap A \subseteq Z_1)\} \oplus \mathcal{H}^C \\ &= \{Z_1 \cup Z_2 \cup Z_3 | (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{F}^B) \wedge (Z_3 \in \mathcal{H}^C) \wedge (Z_1 \cap B \subseteq Z_2) \\ &\quad \wedge (Z_2 \cap A \subseteq Z_1) \wedge (Z_1 \cup Z_2 \cap C \subseteq Z_3) \wedge (Z_3 \cap A \cup B \subseteq Z_1 \cup Z_2)\} \end{aligned}$$

$$\begin{aligned} \mathcal{E}^A \oplus (\mathcal{F}^B \oplus \mathcal{H}^C) &= \mathcal{E}^A \oplus \{Z_2 \cup Z_3 | (Z_2 \in \mathcal{F}^B) \wedge (Z_3 \in \mathcal{H}^C) \wedge (Z_2 \cap C \subseteq Z_3) \wedge (Z_3 \cap B \subseteq Z_2)\} \\ &= \{Z_1 \cup Z_2 \cup Z_3 | (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{F}^B) \wedge (Z_3 \in \mathcal{H}^C) \wedge (Z_2 \cap C \subseteq Z_3) \\ &\quad \wedge (Z_3 \cap B \subseteq Z_2) \wedge (Z_2 \cup Z_3 \cap A \subseteq Z_1) \wedge (Z_1 \cap B \cup C \subseteq Z_2 \cup Z_3)\} \end{aligned}$$

But from lemma 2 it follows that:

$$\mathcal{E}^A \oplus (\mathcal{F}^B \oplus \mathcal{H}^C) = (\mathcal{E}^A \oplus \mathcal{F}^B) \oplus \mathcal{H}^C$$

So operator \oplus is associative.

Theorem 10. *Operation \oplus is idempotent.*

Proof. Given some operation $*$ we say that it is idempotent iff $a * a = a$ for any possible a .

$$\begin{aligned} \mathcal{E}^A \oplus \mathcal{E}^A &= \{Z_1 \cup Z_2 | (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{E}^A) \wedge (Z_1 \cap A \subseteq Z_2) \wedge (Z_2 \cap A \subseteq Z_1)\} \\ &= \{Z_1 \cup Z_2 | (Z_1 \in \mathcal{E}^A) \wedge (Z_2 \in \mathcal{E}^A) \wedge (Z_1 = Z_2)\} \\ &= \{Z_1 | (Z_1 \in \mathcal{E}^A)\} \\ &= \mathcal{E}^A \end{aligned}$$

So operation \oplus is idempotent.

Theorem 11. *Let V be a finite set and $S = \{(\mathcal{E}, A) | \mathcal{E} \subseteq 2^A \wedge A \subseteq V\}$. Then $\langle S, \oplus \rangle$ is a semilattice.*

Proof. A set L with some operations $*$ is a semilattice if the operation $*$ is commutative, associative and idempotent. From the previous theorem all these properties hold for the \oplus operation and the set S .

B Proof of Theorem 2

Proof. Let $C = C_1 \cup C_2$ be the RMT-cut which partitions $V \setminus C$ in sets $A, B \neq \emptyset$ s.t. $D \in A$ and $R \in B$. Then there exists a different scenario where $\mathcal{Z}' = \mathcal{Z}_B$ and all other parameters are the same. From lemma 1, nodes in B have the same initial knowledge in both scenarios, since $\mathcal{Z}_B = \mathcal{Z}'_B$.

Suppose R could decide correctly with \mathcal{Z} being the actual adversary structure. Then using a standard argument employed in [14, 13], an attack on the safeness of the algorithm would be possible in the same setting with \mathcal{Z}' being the actual adversary structure. The details of the proof are similar and are based on the difficulty of the honest players in B to distinguish which scenario they participate in, with respect to the actual adversary structure: the one with \mathcal{Z} or the one with \mathcal{Z}' .

C Definition of \mathcal{Z} -pp cut

Definition 6 (\mathcal{Z} -partial pair cut). Let C be a cut of G partitioning $V \setminus C$ into sets $A, B \neq \emptyset$ s.t. $D \in A$. C is a \mathcal{Z} -partial pair cut (\mathcal{Z} -pp cut) if there exists a partition $C = C_1 \cup C_2$ with $C_1 \in \mathcal{Z}$ and $\forall u \in B, N(u) \cap C_2 \in \mathcal{Z}_u$.

D Proof of Theorem 5

Proof. Suppose that \mathcal{Z} -CPA does not achieve RMT in (G, \mathcal{Z}, D, r) . Then we can split the graph in 3 parts: A being the honest decided nodes, B being the honest undecided nodes with $R \in B$ and C being the corrupted nodes. Now since every node in B is undecided we have that $\forall u \in B : N(u) \cap A \in \mathcal{Z}_u$ (otherwise u would have decided). But then $C \cup A$ is an RMT \mathcal{Z} -pp cut which is a contradiction. Hence, \mathcal{Z} -CPA achieves RMT in (G, \mathcal{Z}, D, R) .

E Proof of Theorem 6

Proof. Let $C = C_1 \cup C_2$ be the RMT \mathcal{Z} -pp cut which partitions $V \setminus C$ in sets $A, B \neq \emptyset$ s.t. $D \in A$ and $R \in B$. Let $\mathcal{Z}' = \{\bigcup_{u \in B} Z \cap N(u) : Z \in \mathcal{Z}\} \cup \{C_2\}$. We have that $\mathcal{Z}'_u = \{Z \cap N(u) : Z \in \mathcal{Z}'\} \cup \{C_2 \cap N(u)\} = \{(\bigcup_{v \in B} Z \cap N(v)) \cap N(u) : Z \in \mathcal{Z}\} \cup \{C_2 \cap N(u)\} = \{Z \cap N(u) : Z \in \mathcal{Z}\} \cup \{C_2 \cap N(u)\}$ but since $\forall u \in B : N(u) \cap C_2 \in \mathcal{Z}_u$, for every node u in B : $\mathcal{Z}_u = \mathcal{Z}'_u$. So far we have established that (a) nodes in B cannot tell whether \mathcal{Z} or \mathcal{Z}' is the adversary structure since $\forall u \in B : \mathcal{Z}_u = \mathcal{Z}'_u$ and (b) C_2 is an admissible corruption set in \mathcal{Z}' .

Suppose that there exists a safe algorithm \mathcal{A} which achieves RMT in instance (G, \mathcal{Z}, D, r) . We consider the following runs e and e' of \mathcal{A} :

- Run e is on the instance (G, \mathcal{Z}, D, R) , with dealer's value $x_D = 0$, and corruption set C_1 ; in each round, all players in C_1 perform the actions that perform in the respective round of run e' (where C_1 is a set of honest players).
- Run e' is on the the instance (G, \mathcal{Z}', D, R) , with dealer's value $x_D = 1$, and corruption set C_2 ; in each round, all players in C_2 perform the actions that perform in the respective round of run e (where C_2 is a set of honest players).

Note that C_1, C_2 are admissible corruption sets in instances $(G, \mathcal{Z}, D, R), (G, \mathcal{Z}', D, R)$ respectively. Since $C_1 \cup C_2$ is a cut which separates D from R in both $(G, \mathcal{Z}, D, R), (G, \mathcal{Z}', D, R)$ and the actions of every node of this cut are identical in both runs e, e' , the messages that R receives are the same in both runs, i.e., $view(v, e) = view(v, e')$. Therefore the decision of $R \in B$ must be identical in both runs. Since, by assumption, algorithm \mathcal{A} achieves RMT in instance (G, \mathcal{Z}, D, R) , R must decide on the dealer's message 0 in run e on (G, \mathcal{Z}, D, R) , and must do the same in run e' on (G, \mathcal{Z}', D, R) . However, in run e' the dealer's message is 1. Therefore \mathcal{A} makes R decide on an incorrect message in (G, \mathcal{Z}', D, R) . This contradicts the assumption that \mathcal{A} is safe.