Reliable Message Transmission under Partial Knowledge and General Adversaries

Aris Pagourtzis[1,2] Giorgos Panagiotakos[3] Dimitris Sakavalas[1]

[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, Kings College London, UK
[3] Department of Informatics and Telecommunications,
University of Athens, 15784 Athens, Greece.
g.panagiotakos@di.uoa.gr

Abstract

A fundamental communication 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 [6], 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, yielding a characterization of the minimal level of knowledge which renders the problem solvable in a certain network. In order to show the sufficiency of the condition, we propose the RMT-Partial Knowledge Algorithm (RMT-PKA), an algorithm that solves RMT whenever the condition is met. This implies that RMT-PKA achieves reliable message transmission in every instance where this is possible, therefore it is a *unique* algorithm [14]. To the best of our knowledge, this is the first unique protocol for RMT 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, allowing participants to combine their local knowledge in a sound manner, 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.

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, full or local topology knowledge. Here we focus on perfectly reliable transmission under a general adversary and the partial knowledge model. In the general adversary model, introduced by Hirt and Maurer [6], the adversary may corrupt any player-set among a given family of all possible corruption sets (*adversary structure*); it subsumes both the global [10] and the local threshold adversary model [8]. For instance, the global threshold model, which assumes that the adversary can corrupt at most t players, corresponds to the family of sets with cardinality at most t. Regarding the topology knowledge, the recently introduced *Partial Knowledge Model* [13] assumes that each player only has knowledge over some arbitrary subgraph including itself and the intersection of this subgraph with the adversary structure; it 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.

The strength of this work lies in the combination of these two quite general models (general adversary and partial knowledge), forming the most general setting we have encountered so far within the synchronous deterministic 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], where 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 of specific topology 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. A simple, yet powerful Reliable Broadcast protocol called *Certified Propagation Algorithm* (CPA) was proposed in this work; CPA is based on the idea that if a set of $t + 1$ neighbors of v provides the same information to v then the information is valid because at least one of them is honest. This work was extended in the context of generic networks by Pelc, Peleg in [14] who also pointed out how full knowledge of the topology yields better solvability results. After a series of works ([7, 11, 18]) tight conditions for the correctness of CPA were obtained in the *ad hoc* case. Observe that all of these aforementioned works only considered the t-locally bounded adversary model and did not not provide tight conditions for the solvability of the problem. Finally, in [13] the Partial Knowledge Model was introduced, in which the players only have partial knowledge of the topology and the adversary structure. In [13] both the t-locally bounded adversary model and the general adversary model were considered and tight conditions

for the solvability of the problem along with matching algorithms for the extreme cases of full topology knowledge and *ad hoc* setting were proposed. Trivially all the aforementioned results for Reliable Broadcast with an honest dealer can be adapted for the RMT problem. However, it was left as an open problem in [13] to determine a necessary and sufficient condition (tight) for the most general case of the partial knowledge model. Moreover these previous studies have focused on feasibility and not efficiency and no complexity studies have been conducted in this context. The latter two issues appeared to be most challenging and are both considered and answered in this work.

1.2 Our results

We study the RMT problem under partial knowledge and 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 RMT-PKA, an algorithm that achieves RMT whenever this condition is met. In terminology of [14] (formally defined in [13]) this is a *unique* algorithm for the problem, in the sense that whenever any algorithm achieves RMT in a certain instance so does RMT-PKA. This settles an open question of [13] and is, to the best of our knowledge, the first algorithm with this property. It is worth mentioning that RMT-PKA can achieve RMT with the minimal amount of player's knowledge that renders the problem solvable.

A key algorithmic tool that we define and use is the *joint view operation* which computes the *joint adversary structure* of (a set of) players, i.e., the worst case adversary structure that conforms to each player's initial knowledge. This operation is crucial in obtaining the tight condition mentioned above since it provides a way to safely utilize the maximal valid information from all the messages exchanged. 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 quite general model.

This new algorithm 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 in a network design phase. A remarkable property of our algorithm is its *safety*: even when RMT is not possible the receiver will never make an incorrect decision despite the increased adversary's attack capabilities, which include reporting fictitious topology and false local knowledge among others.

(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 propose an adaptation of $\mathcal{Z}$-CPA [13] appropriate for RMT. We prove that this protocol is unique for RMT in the *Ad Hoc* model and is the first such algorithm that we know of. We examine whether and when this algorithm 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 there exist classes of instances where RMT is solvable, such that if $\mathcal{Z}$-CPA is not fully polynomial in any of these classes then no RMT protocol can be fully polynomial in the same class. 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 solvabililty 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 $G = (V, E)$ consisting of the player (node) set $V(G)$ and edge set $E(G)$ which represents undirected authenticated channels between players. The set of neighbors of a player v is denoted with $\mathcal{N}(v)$. 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. The problem definition follows.

Reliable Message Transmission. We assume the existence of a designated player $D \in V$, 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 output the value x_D originally sent by the dealer.

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.

The Partial Knowledge Model [13]. In this setting each player v only has knowledge of the topology of a certain subgraph G_v of G which includes v. Namely if we consider the family $\mathcal{G}$ of subgraphs of G we use the *view function* $\gamma : V(G) \to \mathcal{G}$, 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 $S \subseteq V(G)$. The output will correspond to the *joint view* of nodes in S. More specifically, if $\gamma(v) = G_v = (V_v, E_v)$ then $\gamma(S) = G_S = (\bigcup_{v \in S} V_v, \bigcup_{v \in S} E_v)$. 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*).

We denote an instance of the problem by the tuple $\mathcal{I} = (G, \mathcal{Z}, \gamma, D, R)$. We next define some useful protocol properties.

We 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 and any admissible behavior of the corrupted players. We say that an RMT protocol is *safe* if it never causes the receiver R to decide on an incorrect value in any instance.

Definition 1 (Uniqueness of Algorithm)**.** *Let $\mathcal{A}$ be a family of algorithms. An algorithm A is unique (for RMT) among algorithms in $\mathcal{A}$ if the existence of an algorithm of family $\mathcal{A}$ which achieves RMT in an instance $\mathcal{I}$ implies that A also achieves RMT in $\mathcal{I}$.*

A unique algorithm A among $\mathcal{A}$, naturally defines the class of instances in which the problem is solvable by $\mathcal{A}$-algorithms, namely the ones that A achieves RMT in.

2 Partial knowledge and general adversaries

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. For an adversary structure $\mathcal{E}$ and a node set A let $\mathcal{E}^A = \{Z \cap A \mid Z \in \mathcal{E}\}$ denote the restriction of $\mathcal{E}$ to the set A. The joint adversary structure from two restricted adversary structures can be obtained through the $\oplus$ operator. We define the operation on two possibly different structures $\mathcal{E}, \mathcal{F}$ so that the operation is well defined even if a corrupted player provides a different structure than the real one to an honest player.

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 2. *Let $\mathbb{T}^A = 2^{2^A}$ denote the space of adversary structures on a set of nodes A. For any node sets A, B and adversary structures $\mathcal{E}, \mathcal{F}$, the operation $\oplus : \mathbb{T}^A \times \mathbb{T}^B \to \mathbb{T}^{(A \cup B)}$, is defined as follows:*

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

Informally, the $\mathcal{E}^A \oplus \mathcal{F}^B$ operation unites possible corruption sets from $\mathcal{E}^A$ and $\mathcal{F}^B$ that 'agree' on $A \cap B$. In the Appendix, we show that the $\oplus$ operation is commutative, associative and idempotent. The next theorem shows a property of the $\oplus$ operation which is important for our study.

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 $\forall \mathcal{H}' \in \mathbb{T}^{A \cup B}$: if $\mathcal{H}'^A = \mathcal{E}^A$ and $\mathcal{H}'^B = \mathcal{F}^B$ 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 $Z_1 \cap B = Z \cap A \cap B = Z_2 \cap A$. But then, definition 2 implies $Z \in \mathcal{H}$, a contradiction. □

Corollary 2. *For any adversary structure $\mathcal{Z}$ and node sets A, B: $\mathcal{Z}^{(A \cup B)} \subseteq \mathcal{Z}^A \oplus \mathcal{Z}^B$.*

What Corollary 2 tells us is that the $\oplus$ operation gives the maximal (w.r.t inclusion) possible adversary structure that is indistinguishable by two agents that know $\mathcal{Z}^A$ and $\mathcal{Z}^B$ respectively, i.e., it coincides with their knowledge of the adversary structures on sets A and B respectively. Recall that $\mathcal{Z}_u = \mathcal{Z}^{V(\gamma(u))}$. We will prefer to use $\mathcal{Z}_u$ to denote the local adversary structure of player u and $\mathcal{Z}^{V(\gamma(u))}$ to denote the corresponding restriction of the adversary structure. This allows us to define the combined knowledge of a set of nodes B about the adversary structure $\mathcal{Z}$ as follows. For a given adversary structure $\mathcal{Z}$, a view function γ and a node set B let

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

Note that $\mathcal{Z}_B$ 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 2 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.

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 3 (RMT-cut)**.** *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 $R \in B'$. Let $B \subseteq B'$ be the node set of the connected component that R lies in. Then C is a* RMT-cut *iff $C_1 \in \mathcal{Z}$ and $C_2 \cap V(\gamma(B)) \in \mathcal{Z}_B$.*

Theorem 3 (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 combines ideas from [14, 13] with the $\oplus$ operation and is deferred to the Appendix.

3.1 The RMT Partial Knowledge Algorithm (RMT-PKA)

We next present the *RMT Partial Knowledge Algorithm* (RMT-PKA), an RMT protocol which succeeds whenever the condition of Theorem 3 (in fact, its negation) is met, rendering it a tight condition on when RMT is possible. To prove this we provide some supplementary notions.

In Protocol 1 there are two type of messages exchanged. *Type 1 messages* are used to propagate the dealer's value and are of the form (x, p) where $x \in X$ and p is a path. *Type 2 messages* of the form $((v, \gamma(v), \mathcal{Z}_v), p)$ are used for every node v to propagate its initial information $\gamma(v), \mathcal{Z}_v$ throughout the graph. Let M denote 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 will say that $value(M) = x$ if and only if all the type 1 messages of M report the same dealer value x, i.e., for every such message (y, p), it holds that $y = x$, for some $x \in X$. Observe that M may consist of messages which contain contradictory information. We next define the form of a message set M which contains no contradictory information in our setting (a valid set M).

Definition 4 (Valid set M)**.** *A set M of both type 1 and type 2 messages* corresponds to a valid scenario*, or more simply is* valid*, if*

- *$\exists x \in X$ s.t. $value(M) = x$. That is, all type 1 messages relay the same x as dealer's value.*
- *$\forall m_1, m_2 \in M$ of type 2, their first component is the same when they refer to the same node. That is, if $m_1 = ((v, \gamma(v), \mathcal{Z}_v), p)$ and $m_2 = (((v', \gamma'(v), \mathcal{Z}'_v), p')$, then $v = v'$ implies that $\gamma(v) = \gamma'(v)$ and $\mathcal{Z}_v = \mathcal{Z}'_v$.*

For every valid M we can define the pair (G_M, x_M) where $x_M = value(M)$. To define G_M let V_M be the set of nodes u for which the information $\gamma(u), \mathcal{Z}_u$ is included in M, namely $V_M = \{v \mid ((v, \gamma(v), \mathcal{Z}_v), p) \in M \text{ for some path } p\}$. Then, G_M is the node induced subgraph of graph $\gamma(V_M)$ on node set V_M. Therefore, a valid message set M uniquely determines the pair (G_M, x_M). We next propose two notions that we use to check if a valid set M contains correct information.

Definition 5 (full message set)**.** *A* full *message set M is a valid set M that contains all the $D - R$ paths which appear in G_M as part of type 1 messages.*

Definition 6 (Adversary cover of set M)**.** *A set $C \subseteq V_M$ is an* adversary cover *of message set M if C has the following property: C is a cut between D, R on G_M and if B is the node set of the connected component that R lies in, it holds that $(C \cap V(\gamma(B))) \in \mathcal{Z}_B$.*

Protocol 1: *RMT-PKA*

Input (for each node v): dealer's label D, $\gamma(v)$, $\mathcal{Z}_v$. Additional input for D : value $x_D \in X$ (message space).
Message format: *type 1*: pair (x, p) or *type 2*: pair $((u, \gamma(u), \mathcal{Z}_u), p)$, where $x \in X$, u the id of some node, $\gamma(u)$ is the view of node u, $\mathcal{Z}_u$ is the local adversary structure of node u, and p is a path of G (message's propagation trail).

Code for D: send messages $(x_D, \{D\})$ 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` $(v \in p) \vee (tail(p) \neq u)$[1] `then discard` (a, p) `else send` $(a, p||v)$ [2] to all neighbours.

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

`if decision` $\neq \perp$ `then output decision` and terminate.

Subroutine `decision`

(* *dealer propagation rule* *)

`if` $R \in \mathcal{N}(D)$ and R receives $(x_D, \{D\})$ `then return` x_D.

(* *full message set propagation rule* *)

`if` R receives a full set M with $value(M) = x$ and $\nexists$ an adversary cover for M

`then return` x `else return` $\perp$.

We next show the somewhat counterintuitive safety property of RMT-PKA, i.e., that the receiver will never decide on an incorrect value despite the increased adversary's attack capabilities, which includes reporting fictitious nodes and false local knowledge.

Theorem 4 (RMT-PKA Safety)**.** *RMT-PKA is safe.*

Proof. It is trivial to see that the receiver R will not decide on an incorrect dealer value by using the dealer propagation rule (case $R \in \mathcal{N}(D)$) due to the dealer's presumed honesty.

The hard part is to prove that R will not decide on any value $x \neq x_D$ by using the full message set propagation rule (case $R \notin \mathcal{N}(D)$). Let $T \in \mathcal{Z}$ be any admissible corruption set and consider the run e_T of RMT-PKA where T is the actual corruption set. Assume that at some round of e_T, R receives a full message set M' with $value(M') = x \neq x_D$. Since all $D - R$ paths of $G_{M'}$ propagate an incorrect value x it means that $C = T \cap V_{M'}$ forms a $D - R$ cut in graph $G_{M'}$, otherwise there would be a $D - R$ path in $G_{M'}$ consisting only of honest nodes and propagating x_D, a contradiction because $value(M') = x$. Since $C \in \mathcal{Z}$, it holds by definition that $C \cap V(\gamma(S)) \in \mathcal{Z}_S$, $\forall S \subseteq V(G)$. Therefore if B is the connected component that R lies in under the partition that C imposes in $G_{M'}$, it holds that $C \cap V(\gamma(B)) \in \mathcal{Z}_B$ due to the fact that B only contains honest nodes; more specifically, B does not contain any corrupted nodes due to the definition of C. Moreover, the adversary cannot introduce any fictitious nodes in B because T has to be a cut between R and every nonexistent node claimed by the adversary. The latter observations about B imply that R can correctly compute $\mathcal{Z}_B$. Thus M' has an adversary cover and R will not decide in value $x \neq x_D$ due to the full message set propagation rule. □

Theorem 5 (Sufficiency)**.** *Let $(G, \mathcal{Z}, \gamma, D, R)$ be an RMT instance. If no RMT-cut exists, then RMT-PKA achieves reliable message transmission.*

[1] We use $tail(p)$ to denote the last node of path p. Checking whether $tail(p) \neq u$ we ensure that at least one corrupted node will be included in a faulty propagation path.

[2] By $p||v$ we denote the concatenation of path p with node v.

The proof (see Appendix) combines techniques from [13] (correctness of the Path Propagation Agorithm) with the novel notions of full message set M, adversary cover of M and corresponding graph G_M.

Corollary 6 (Uniqueness)**.** *RMT-PKA is unique among safe algorithms, i.e., given an RMT instance $(G, \mathcal{Z}, \gamma, D, R)$, if there exists any safe RMT algorithm which is resilient for this instance, then RMT-PKA also achieves reliable message transmission on this instance.*

RMT under minimal knowledge. Observe that the non-existence of an RMT-cut proves to be a necessary and sufficient condition for achieving RMT safely (with a safe algorithm). Equivalently we observe that the condition describes the minimal amount of initial knowledge needed to achieve RMT. Namely, we can define a natural partial ordering of the view functions s.t. for a certain graph $G = (V, E)$ and adversary structure $\mathcal{Z}$ it holds that $\gamma' < \gamma$ if and only if $\forall v \in V, \gamma'(v)$ is a subgraph of $\gamma(v)$. Then a minimal amount of initial knowledge which is needed to achieve RMT corresponds to a minimal function γ, with respect to the above partial ordering, such that there does not exist an RMT-cut in G.

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 7 (*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 7 (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 8 (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. Observe that, if the adversary structure is given explicitly, $\mathcal{Z}$-CPA is trivially *fully polynomial* w.r.t. the total size of the input. However $\mathcal{Z}$ can be of exponential size w.r.t. $|G|$. Measuring the complexity of $\mathcal{Z}$-CPA is not straightforward since the computations included in the protocol are not explicitly defined. In fact, $\mathcal{Z}$-CPA is actually a *protocol scheme* that refers to a "functionally specified" subroutine rather than to an actual (implementation of such a) subroutine. We will use the following notions to facilitate our study on distributed protocol schemes.

Definition 8 (Protocol scheme)**.** *A* protocol scheme $\mathcal{A}$ *is a family of protocols which contains calls to a subroutine X for solving a problem S. The computation of X is not specified, that is, X is used as a black box. Therefore, for every algorithm B which solves problem S a different member (protocol) $\mathcal{A}_B$ of $\mathcal{A}$ is defined; that is, $\mathcal{A}_B$ implements subroutine X through algorithm B.*

Essentially, if protocol scheme $\mathcal{A}$ solves problem Q, then $\mathcal{A}$ is in fact a reduction from Q to S in the distributed setting.

We will say that a *protocol scheme $\mathcal{A}$ is fully polynomial* if there exists an algorithm B, solving S, for which $\mathcal{A}_B$ is fully polynomial.

Observe that $\mathcal{Z}$-CPA is a protocol scheme which contains the membership check subroutine ($N \notin \mathcal{Z}_v$) appearing in its second rule. Regarding the efficiency of $\mathcal{Z}$-CPA, one can easily observe that it is of polynomial round and bit complexity (details appear in the proof of Theorem 9). To argue about the local computations complexity of the scheme we need to take into account the complexity of the membership check subroutine. Indeed, the $\mathcal{Z}$-CPA scheme is fully polynomial if there exists an algorithm B through which the membership check can be performed in polynomial time w.r.t. the size of the input graph $|G|$. We next introduce the property of *poly-time uniqueness*; a protocol scheme that is poly-time unique for a problem is, in a sense, optimally efficient w.r.t. a polynomial factor.

Definition 9 (Poly-time Uniqueness)**.** *We call a protocol scheme $\mathcal{A}$* poly-time unique *for problem $\mathcal{P}$ if it is unique (with respect to feasibility) and the existence of a unique fully polynomial protocol for $\mathcal{P}$ implies that $\mathcal{A}$ is also fully polynomial for $\mathcal{P}$.*

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 terms of reducibility, the concept of poly-time uniqueness of a protocol scheme $\mathcal{A}$ implies that $\mathcal{A}$ can be used as a self reduction for the given problem, as will be clear in the following.

We believe that this concept could be of more general interest, since it can be used to argue about optimality of protocol schemes and identify subproblems that are crucial for solving the original problem.

In the main theorem of this section we prove that the $\mathcal{Z}$-CPA scheme is poly-time unique for the RMT problem, and thus show that the $\mathcal{Z}$-CPA scheme is at least as efficient, up to a polynomial factor, as any other RMT protocol scheme. To show that, we build a self reduction for RMT based on $\mathcal{Z}$-CPA. We essentially show that if a unique fully polynomial RMT algorithm exists, it must be able to answer the membership check 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 7,8:

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

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

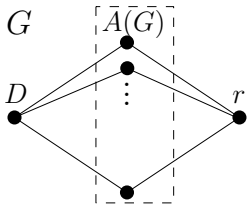

Figure 1: Family of instances $\mathcal{G}'$. No RMT $\mathcal{Z}$-pp cut exists.

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 $\mathcal{Z}' = \mathcal{Z}_v$ [3]. More precisely,

$$\mathcal{I}(\mathcal{G}_1) = \{(G', \mathcal{Z}', D', R') \in \mathcal{G}' \mid \exists (G, \mathcal{Z}, D, R) \in \mathcal{G}_1,\ \exists v \in V(G) \setminus \{D\}, A(G') \subseteq \mathcal{N}(v), \mathcal{Z}' = \mathcal{Z}_v\}$$

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

[3] 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.

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. For convenience, we will use the following notation regarding the executions (runs) of the algorithms and the views of the players.

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 $decision_e(v)$; for deterministic protocols, considered in this work, the $decision_e(v)$ function 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.

Theorem 9. *Let $\mathcal{G}_1$ be a set of instances where RMT is solvable. 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 $RMT|_{\mathcal{G}_1}$.*

The main idea is to use the $\mathcal{Z}$-CPA with protocol Π as a subroutine in order to obtain a fully polynomial algorithm for $RMT|_{\mathcal{G}_1}$. In particular, 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. The complete proof is deferred to the Appendix.

Based on the definition of poly-time uniqueness and Theorem 9, we obtain the following corollary on $\mathcal{Z}$-CPA.

Corollary 10. *Protocol scheme $\mathcal{Z}$-CPA is poly-time unique for RMT.*

6 Conclusions and open questions

Regarding the partial knowledge model, the RMT-PKA protocol employs topology information exchange between players. 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.

We have shown that RMT-PKA protocol is unique for the partial knowledge model; this only addresses the feasibility issue. A natural question is whether and when we can devise a unique and also efficient algorithm for this setting. The techniques used so far to reduce the communication complexity (e.g. [9]) do not seem to be directly applicable to this model. So, exploring this direction further is particularly meaningful.

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.

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

Observe that an equivalent definition for the The $\oplus$ operation is the following,

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

Theorem 11. *Operator $\oplus$ is commutative.*

Proof. For any adversary structures $\mathcal{E}, \mathcal{F}$ and node sets A, B:

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

So operator $\oplus$ is commutative. □

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

Lemma 12. *For any node sets A, B, C it holds that*

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

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 13. *Operator $\oplus$ is associative.*

Proof. For any adversary structures $\mathcal{E}, \mathcal{F}, \mathcal{H}$ and node sets A, B, C:

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

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

But from lemma 12 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 14. *Operation $\oplus$ is idempotent.*

Proof.

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

So operation $\oplus$ is idempotent. □

Theorem 15. *Let V be a finite set and $S = \{(\mathcal{E}, A) | \mathcal{E} \subseteq 2^A \wedge A \subseteq V\}$. Then $< S, \oplus >$ 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 3

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 $R \in B$. Without loss of generality assume that B is connected. If it is not, then by adding all nodes, that do not belong to the connect component of R, to A, a different RMT-cut is built with the desired properties. Consider the instance where $\mathcal{Z}' = \mathcal{Z}_B$ and all other parameters are the same as the previous scenario. Then, all nodes in B have the same initial knowledge in both instances, 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 safety of the algorithm would be possible in the same setting with $\mathcal{Z}'$ being the actual adversary structure. The basic idea is that we can construct two indistinguishable scenarios, where the value that the dealer broadcasts is different. 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 Proof of Theorem 5

Proof. Observe that if $R \in \mathcal{N}(D)$ then R trivially decides on x_D due to the *dealer propagation rule*, since the dealer is honest. Assuming that no RMT-cut exists, we will show that if $R \notin \mathcal{N}(D)$ then R will decide on x_D due to the *full message set propagation rule*.

Let $T \in \mathcal{Z}$ be any admissible corruption set and consider the run e_T of RMT-PKA where T is the actual corruption set. Let P be the set of all paths connecting D with R and are composed entirely by nodes in $V(G) \setminus T$ (honest nodes). Observe that $P \neq \emptyset$, otherwise T is a cut separating D from R which is trivially a RMT-cut, a contradiction.

Since paths in P are entirely composed by honest nodes, it should be clear by the protocol that by round $|V(G)|$, R will have obtained x_D through all paths in P by receiving the corresponding type 1 messages M_1. Furthermore, by round $|V(G)|$, R will have received type 2 messages set M_2 which includes information for all the nodes connected with R via paths that do not pass through nodes in T. This includes all nodes of paths in P. Consequently, R will have received the full message set $M = M_1 \cup M_2$ with $value(M) = x_D$.

We next show that there is no adversary cover for M and thus R will decide on x_D through the full message set propagation rule on M. Assume that there exists an adversary cover C for M. This, by definition means that C is a cut between D, R on G_M and if B is the node set of the the connected component that R lies in, it holds that and $(C \cap V(\gamma(B))) \in \mathcal{Z}_B$ (observe that R can compute $\mathcal{Z}_B$ using the information contained in M_2 as defined in the previous paragraph). Then obviously $T \cup C$ is a cut in G separating D from R, since every path of G that connects D with R contains at least a node in $T \cup C$. Let the cut $T \cup C$ partition $V(G) \setminus \{T \cup C\}$ in the sets A, B s.t. $D \in A$. Then clearly $T \cup C$ is an RMT cut by definition, a contradiction. Thus there is no adversary cover for M and R will decide on x_D.

Moreover, since RMT-PKA is safe, the receiver will not decide on any other value different from x_D.

□

D Definition of $\mathcal{Z}$-pp cut

Definition 10 ($\mathcal{Z}$-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 $\forall u \in B,\ \mathcal{N}(u) \cap C_2 \in \mathcal{Z}_u$.*

E Proof of Theorem 7

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 $(G, \mathcal{Z}, D, R)$. □

F Proof of Theorem 8

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. □

G Proof of Theorem 9

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 $\exists i \in \{1, \ldots, n\}, R \in \mathcal{D}(i)$. 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 $poly(n)$. Namely, we show by induction that for every round i, each player $v \in \mathcal{D}(i)$ will decide in $poly(n)$. Since $\exists i \in \{1, \ldots, n\}, R \in \mathcal{D}(i)$ 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 $poly(n)$ time.

Assume that, for every round $i \leq k$ every $v \in \mathcal{D}(i)$ decides in $poly(n)$-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 $view(v, e, k)$ to decide, from players in $\bigcup_{i=1,\ldots,k} \mathcal{D}(i)$, in $poly(n)$-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 $poly(n)$ time since $|X| = poly(n)$. Given $view(v, e, k)$, player v, in $poly(n)$-time, can create a partition of its neighborhood $\mathcal{N}(v) = \bigcup_{i=0}^{m+1} A_i$ such that

$$\begin{aligned} A_0 =& \{u \in \mathcal{N}(v) \mid \text{u sent nothing}\} \\ A_i =& \{u \in \mathcal{N}(v) \mid \text{u sent value } a_i \in X\},\ i = \{1, \ldots m\} \\ A_{m+1} =& \{u \in \mathcal{N}(v) \mid \text{u sent erroneous messages}\} \end{aligned}$$

Since sets A_0, A_{m+1} do not affect our study we let $A = \bigcup_{i=\{1,\ldots m\}} A_i$. 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 $view(v, e, k)$, observe that

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

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, $view(v, e, k)$ 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 $\forall i \in \{1, \ldots, m\} \setminus \{h\}$ it holds that $A_i \in \mathcal{Z}_v$ since $A \setminus A_h \subseteq Z_v$. Consequently

$$\exists!\ h \in \{1, \ldots, m\} \text{ s.t. } A_h \notin \mathcal{Z}_v \text{ and } A \setminus A_h \in \mathcal{Z}_v \tag{1}$$

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

For $l = 1, \ldots m$, 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 $view(e_1^l, v)|_{A \setminus A_l}$.

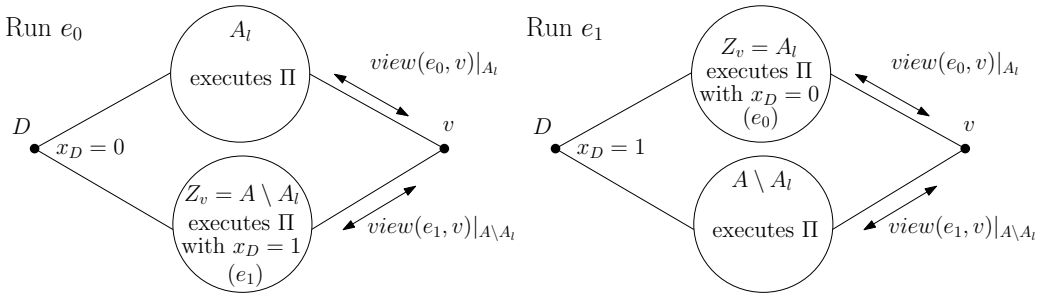

Figure 2: Runs e_0 and e_1.

- 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 $view(e_0^l, v)|_{A_l}$.

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 bound [4] 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, $2m = poly(n)$ runs $(e_0^l, e_1^l)_{l \in \{1,\ldots m\}}$ and halts all parallel simulations with decision a_l if run e_0^l terminates with $decision(v) = 0$.

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

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

"⇒": $A_l \notin \mathcal{Z}_v \Rightarrow Z_v = A \setminus A_l \in \mathcal{Z}_v$. Since by assumption Π solves $RMT|_{\mathcal{I}(\mathcal{G}_1)}$, 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., $decision_{e_0^l}(v) = 0$.

"⇐": Let $A_l \in \mathcal{Z}_v$ and $decision_{e_0^l}(v) = 0$. 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 $RMT|_{\mathcal{I}(\mathcal{G}_1)}$, 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., $decision_{e_1^l}(v) = 1$. But since the decision 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 $decision_{e_0^l}(v) = 1$, 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 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.

[4] 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 $\mathcal{I}(G_1)$ and thus their local computations complexity would be polynomial.

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. □