Reliable communication via semilattice properties of partial knowledge

Aris Pagourtzis[1] Giorgos Panagiotakos[2] 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 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 [5], which subsumes earlier models such as the global or local threshold adversaries. Regarding the topology knowledge, we employ the recently introduced *Partial Knowledge Model* [12], which encompasses both the full knowledge and the *ad hoc* model; the latter assumes knowledge of the local neighborhood only.

Our main contribution is a tight condition for achieving RMT in the partial knowledge model under a general adversary. A key algorithmic tool that we define and use is the *joint view operation* which imposes a semilattice structure on the partial knowledge possessed by different parties. In this context, we prove that the worst possible adversary structure, conforming with the initial knowledge of a set of parties, can be expressed as the supremum of the parties' knowledge under the semilattice partial order. The new operation allows for the definition of an appropriate network separator notion that yields a necessary condition for achieving RMT. In order to show the sufficiency of the condition, we propose the RMT Partial Knowledge Algorithm (RMT-PKA), an algorithm that also employs the joint view operation to solve 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 [13]. To the best of our knowledge, this is the first unique protocol for RMT against general adversaries in the partial knowledge model. Due to the generality of the model, our results provide, for any level of topology knowledge and any adversary structure, an exact characterization of instances where RMT is possible and an algorithm to achieve RMT on such instances.

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 [9]. 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 [5], the adversary may corrupt any player-set among a given family of all possible corruption sets (*adversary structure*); it subsumes both the global [9] and the local threshold adversary model [7]. 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* [12] 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 [5], where given in [8, 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.* [8].

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 [7]. 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 [13] who also pointed out how full knowledge of the topology yields better solvability results. After a series of works ([6, 10, 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 [12] the Partial Knowledge Model was introduced, in which the players only have partial knowledge of the topology and the adversary structure. In [12] 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 [12] 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 concerns the 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 [13] (formally defined in [12]) 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 [12] 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. This new algorithm encompasses earlier algorithms such as CPA [7], PPA and $\mathcal{Z}$-CPA [12] as special cases. 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.

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 show that this operation actually implies a semilattice structure on the partial knowledge that players may have. In semilattice terminology, the joint adversary structure is the supremum of the adversary structures known by the involved players under the induced partial order.

To obtain our result we also generalize earlier pair-cut techniques, introduced by Pelc and Peleg [13] and extended in [12] in the context of Broadcast. This technique was used in [12] to obtain characterizations of classes of graphs for which Broadcast is possible for various levels of topology knowledge and types of corruption distribution; however, an exact characterization for the partial knowledge setting was left as an open question. Here we address this question by proposing a new type of pair-cut appropriate for the partial knowledge model, coupled with a proof that RMT-PKA works exactly whenever no such pair-cut exists. This, as already mentioned, implies a tight solvability condition for RMT in the quite general model of partial knowledge with general adversaries. A useful by-product of practical interest is that the new cut notion can be used, in a network design phase, in order to determine the exact subgraph in which RMT is possible.

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 [5]. 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, an adversary 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 [12]. 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., $\forall v \in V(G),\ \gamma(v) = \mathcal{N}(v)$.

In order to capture partial knowledge in this setting we need to define the restriction of some structure to an a set of nodes.

Definition 1. *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.*

Hence, we assume that given the actual adversary structure $\mathcal{Z}$ each player v only knows the possible corruption sets under his view $\mathcal{Z}_v$, which is equal to $\mathcal{Z}^{V(\gamma(v))}$ (the *local adversary structure*).

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

We say that an RMT protocol is *resilient* for an instance I if it achieves RMT on instance 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 2 (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 I implies that A also achieves RMT in 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 The algebraic structure of partial knowledge

In this section we delve into the algebraic structure of the knowledge of players regarding the adversary. We do this by first defining an operation used to calculate their joint knowledge. The

operation takes into account potentially different adversarial structures, so that it is well defined even if a corrupted player provides a different structure than the real one to some honest player.

Definition 3. *Let V be a finite node set; let also $\mathbb{T} = \{(\mathcal{E}, A) \mid \mathcal{E} \subseteq 2^A, A \subseteq V, \mathcal{E}$ is monotone$\}$ denote the space all pairs consisting of a monotone family of subsets of a node set along with that node set. The operation $\oplus : \mathbb{T} \times \mathbb{T} \to \mathbb{T}$, is defined as follows:*

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

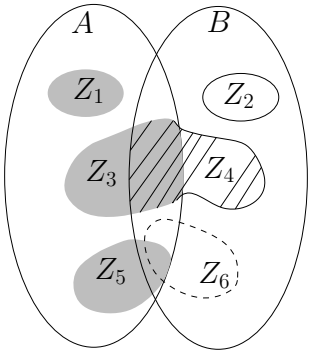

Figure 1: Example of the $\oplus$ operation in the pairs $(\mathcal{E}, A), (\mathcal{F}, B)$: $Z_1 \cup Z_2$ and $Z_3 \cup Z_4$ belong to $(\mathcal{E}, A) \oplus (\mathcal{F}, B)$ but $Z_5 \cup Z_6$ does not, neither does $Z_1 \cup Z_4$.

Informally, $(\mathcal{E}, A) \oplus (\mathcal{F}, B)$ unites possible corruption sets from $\mathcal{E}$ and $\mathcal{F}$ that 'agree' on $A \cap B$ (see Figure 1). The following theorem offers further insight on the algebraic properties of this operation, by revealing a semilattice structure on the space of partial knowledge obtained by the players. The semilattice structure is shown by proving the commutativity, associativity and idempotence properties of operation $\oplus$ (see [14]). The proof is deferred to the Appendix.

Theorem 1. *$\langle \mathbb{T}, \oplus \rangle$ is a semilattice.*

From semilattice theory, it is well known that the algebraic definition of the join-semilattice [1] $\langle \mathbb{T}, \oplus \rangle$ implies a binary relation $\geq$ that partially orders $\mathbb{T}$ in the following way: for all elements $x, y \in \mathbb{T}$, $x \geq y$ if and only if $x = x \oplus y$. This binary relation provides the equivalent order theoretic definition of the same semilattice $\langle \mathbb{T}, \geq \rangle$.

The following theorem reveals the binary relation implied by the $\oplus$ operation.

Theorem 2. *The partial ordering " $\geq$ " induced by the $\oplus$ operation on $\mathbb{T}$ satisfies the following: for $(\mathcal{E}, A), (\mathcal{F}, B) \in \mathbb{T}$, $(\mathcal{E}, A) \geq (\mathcal{F}, B)$ if and only if $(B \subseteq A) \wedge (\mathcal{E}^B \subseteq \mathcal{F})$.*

Proof. By the relation of the algebraic and order theoretic definitions of a semilattice we have to show that

$$(\mathcal{E}, A) \oplus (\mathcal{F}, B) = (\mathcal{E}, A) \Leftrightarrow (B \subseteq A) \wedge (\mathcal{E}^B \subseteq \mathcal{F})$$

" $\Leftarrow$ "

Observe that $(B \subseteq A)$ implies that $A \cup B$ is equal to A. Therefore let $(\mathcal{E}, A) \oplus (\mathcal{F}, B) = (\mathcal{H}, A)$. From $\mathcal{E}^B \subseteq \mathcal{F}$ and the monotonicity property it follows that:

$$\forall Z_1 \in \mathcal{E}, \exists Z_2 \in \mathcal{F} : (Z_1 \cap B = Z_2 \cap A) \wedge (Z_1 \cup Z_2 = Z_1) \Rightarrow \mathcal{E} \subseteq \mathcal{H}$$

Notice that, due to monotonicity even if $Z_1 \cap B = \emptyset$ the relation holds for $Z_2 = \emptyset$. It remains to show that $\mathcal{H} \subseteq \mathcal{E}$. From the definition of the $\oplus$ operation it follows that:

$$\forall Z \in \mathcal{H}, \exists Z_1 \in \mathcal{E}, Z_2 \in \mathcal{F} : (Z_1 \cap B = Z_2 \cap A) \wedge (Z_1 \cup Z_2 = Z) \Rightarrow \mathcal{H} \subseteq \mathcal{E}$$

[1] We make the convention $\langle \mathbb{T}, \oplus \rangle$ is a join-semilattice because the implied partial order $\geq$ captures our case more intuitively. The notion of meet-semilattice can be used as well by inversing the ordering.

where the implication follows by the fact that $Z_1 \cup Z_2 = Z_1$. Thus, $\mathcal{H} = \mathcal{E}$.
" $\Rightarrow$ "

By Definition 3 it is implied that $A \cup B = A$ or equivalently $B \subseteq A$. Thus, it suffices to show that $\mathcal{E}^B \subseteq F$. For the sake of contradiction, assume that $\mathcal{E}^B \not\subseteq \mathcal{F}$. It holds that:

$$\begin{aligned}\mathcal{E}^B \not\subseteq \mathcal{F} &\Rightarrow \exists Z_1 \in \mathcal{E}^B : Z_1 \notin \mathcal{F} \\ &\Rightarrow \exists Z_1 \in \mathcal{E}, \forall Z_2 \in \mathcal{F} : Z_1 \cap B = Z_2 \cap A \\ &\Rightarrow (\mathcal{E}, A) \oplus (\mathcal{F}, B) \neq (\mathcal{E}, A)\end{aligned}$$

which leads to a contradiction. Thus, the theorem follows. □

The semilattice structure guarantees that every non-empty finite subset of $\langle \mathbb{T}, \geq \rangle$ has a supremum with respect to the " $\geq$ " relation (also called a *join*). Moreover it holds that for $(\mathcal{E}, A), (\mathcal{F}, B) \in \mathbb{T}$, $\sup\{(\mathcal{E}, A), (\mathcal{F}, B)\} = (\mathcal{E}, A) \oplus (\mathcal{F}, B)$. The latter implies a property of the $\oplus$ operation which is important in our study. Namely,

Corollary 3. *Let $\langle \mathbb{T}, \geq \rangle$ be a semilattice as defined above. For any $z \in \mathbb{T}$ it holds that if $x, y \leq z$, then $x \oplus y \leq z$.*

Proof. The join of x, y is their least upper bound. Thus, since z is an upper bound of x, y, it must also be greater or equal to their join, i.e. $x \oplus y$. The Corollary follows. □

Returning to our problem after this short detour, notice that for any adversary structure $\mathcal{Z}$ it holds that $(\mathcal{Z}^A, A), (\mathcal{Z}^B, B) \leq (\mathcal{Z}^{A \cup B}, A \cup B)$. We immediately get by Corollary 3 the following corollary.

Corollary 4. *For any adversary structure $\mathcal{Z}$ and node sets A, B:*

$$\text{if } (\mathcal{H}, A \cup B) = (\mathcal{Z}^A, A) \oplus (\mathcal{Z}^B, B) \text{ then } \mathcal{Z}^{(A \cup B)} \subseteq \mathcal{H}$$

What Corollary 4 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.

Now recall that $\mathcal{Z}_u = \mathcal{Z}^{V(\gamma(u))}$. 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

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

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 4 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, always contains the actual adversary structure in their scenario.

3 A tight condition for RMT

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 4 (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 $C_2 \cap V(\gamma(B)) \in \mathcal{Z}_B$.*

The necessary condition proof adapts techniques and ideas from [13, 12] to the partial knowledge with general adversary setting.

Theorem 5 (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 $(G, \mathcal{Z}, \gamma, D, R)$.*

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 to A all nodes that do not belong to the connected component of R, an RMT-cut with the desired property is obtained. Consider a second instance where $\mathcal{Z}' = \mathcal{Z}_B$ and all other parameters are the same as in the original instance. Recall that $\mathcal{Z}_B$ is defined using the $\oplus$ operator and exactly captures (by Corollary 4) the worst case adversary structure possible, restricted to $V(\gamma(B))$, relative to the initial knowledge of players in B. Hence, all nodes in B have the same initial knowledge in both instances, since $\mathcal{Z}_B = \mathcal{Z}'_B$.

The proof is by contradiction. Suppose that there exists a safe algorithm $\mathcal{A}$ which is resilient for $(G, \mathcal{Z}, \gamma, D, R)$. We consider the following executions σ and σ' of $\mathcal{A}$:

- Execution σ is on instance $(G, \mathcal{Z}, \gamma, D, R)$, with dealer's value $x_D = 0$, and corruption set C_1; in each round, each corrupted player in C_1 performs the actions that its corresponding player performs in the respective round of execution σ' (where C_1 consists of honest players only).
- Execution σ' is on instance $(G, \mathcal{Z}', \gamma, D, R)$, with dealer's value $x_D = 1$, and corruption set C_2; in each round, each corrupted player in C_2 performs the actions that its corresponding player performs in the respective round of execution σ (where C_2 consists of honest players only).

Note that C_1, C_2 are admissible corruption sets in scenarios σ, σ' respectively since they belong to $\mathcal{Z}$ and $\mathcal{Z}'$ (resp.) It is easy to see that $C_1 \cup C_2$ is a cut which separates D from B in both instances and that actions of every node of this cut are identical in both executions σ, σ'. Consequently, the actions of any honest node $w \in B$ must be identical in both executions. Since, by assumption, algorithm $\mathcal{A}$ is resilient on $(G, \mathcal{Z}, \gamma, D, R)$, R must decide on the dealer's message 0 in execution σ, and must do the same in execution σ'. However, in execution σ' the dealer's message is 1. Therefore $\mathcal{A}$ makes R decide on an incorrect message in $(G, \mathcal{Z}', \gamma, D, R)$. This contradicts the assumption that $\mathcal{A}$ is safe. □

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 5 (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 RMT-PKA there are two types 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 5 (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 $x_M = value(M)$; we assume that $x_M = \perp$ if no type 1 messages are included in 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 $V_M = \{v \mid ((v, \gamma(v), \mathcal{Z}_v), p) \in M$ 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 6 (full message set)**.** *A* full *message set* M *received by* R*, is a valid set* M*, with* $value(M) \neq \perp$*, that contains all the* $D \rightsquigarrow R$ *paths which appear in* G_M *as part of type 1 messages.*

Next we define the notion of *adversary cover* of a full message set M. If such a cut exists, then there is a scenario where all propagated values might be false.

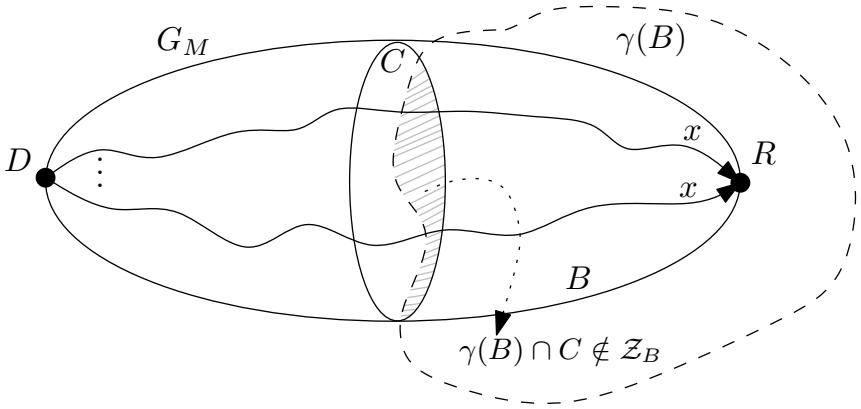

Figure 2: Node set C is an adversary cover of message set M, if it disconnects D, R in G_M and "looks" corruptible under the joint knowledge of B, which represents the node set of the connected component that R lies in.

Definition 7 (Adversary cover of full message set M)**.** *A set* $C \subseteq V_M$ *is an* adversary cover *of full message set* M *if* C *has the following property:* C *is a cut between* D *and* R *on* G_M *and if* B *is the node set of the connected component that* R *lies in, it holds that* $(C \cap V(\gamma(B))) \in \mathcal{Z}_B$.

A graphical representation of the adversary cover is depicted in Figure 2. With the predicate `nocover`(M) we will denote the non existence of an adversary cover of M.

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 6 (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 $D \rightsquigarrow R$

RMT Partial Knowledge Algorithm (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).
Type 1 message format: pair (x, p)
Type 2 message format: 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` $(v \in p) \vee (tail(p) \neq u)$[2] `then discard` (a, p) `else send` $(a, p||v)$ [3] to all neighbours.

Code for R: Initialize $M_R \leftarrow \emptyset$
upon reception of type 1 or type 2 message (x, p) from node u `do:`

`if` $(v \in p) \vee (tail(p) \neq u)$ `then discard` (x, p) `else` $M_R \leftarrow M_R \cup (a, p)$
`if` (x, p) is a type 1 message then
  $lastmsg \leftarrow (x, p)$
  `if decision`$(M_R, lastmsg) = x$ `then output` x and terminate.

`function decision`$(M_R, lastmsg)$

`if` $R \in \mathcal{N}(D)$ `then`
  `if` $lastmsg = (x_D, \{D\})$ `then return` x_D
  `else return` $\perp$.
`for all` valid $M \subseteq M_R$ with $value(M) = value(lastmsg)$ `do`
  compute graph G_M
  $M_1 \leftarrow$ type 1 messages of M
  $\mathcal{P}_1 \leftarrow$ set of all paths p with $(x, p) \in M_1$
  $\mathcal{P}_{D,R} \leftarrow$ set of all $D \rightsquigarrow R$ paths of G_M
  `if` $(\mathcal{P}_{D,R} \subseteq \mathcal{P}_1) \wedge$ `nocover`(M) `then` ▷ *full message set with no*
    `return` $value(lastmsg)$ `else return` $\perp$. ▷ *adversary cover*

`function nocover`(M)

$check \leftarrow$ `true`
`for all` $C \subseteq V_M$ `do`
  `if` C is a (D, R) cut on G_M `then`
    $B \leftarrow$ connected component of R in $G_M \setminus C$
    $(\mathcal{Z}_B, V(\gamma(B))) \leftarrow \bigoplus_{v \in B} (\mathcal{Z}_v, V(\gamma(v)))$ ▷ *joint adversary structure*
    `if` $(C \cap V(\gamma(B))) \in \mathcal{Z}_B$ `then` $check \leftarrow$`false`
`return` $check$

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 $D \rightsquigarrow R$ 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. □

The sufficiency proof combines techniques from [12] (correctness of the Path Propagation Agorithm) with the novel notions of full message set M, adversary cover of M and corresponding graph G_M.

Theorem 7 (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. 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.

□

Corollary 8 (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 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 ([11],[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. [8]) do not seem to be directly applicable to this model. So, exploring this direction further is particularly meaningful.

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] 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.

[6] 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.

[7] Chiu-Yuen Koo. Broadcast in radio networks tolerating byzantine adversarial behavior. In Soma Chaudhuri and Shay Kutten, editors, *PODC*, pages 275–282. ACM, 2004.

[8] 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.

[9] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[10] 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.

[11] Mikhail Nesterenko and Sébastien Tixeuil. Discovering network topology in the presence of byzantine faults. *IEEE Trans. Parallel Distrib. Syst.*, 20(12):1777–1789, 2009.

[12] 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.

[13] Andrzej Pelc and David Peleg. Broadcasting with locally bounded byzantine faults. *Inf. Process. Lett.*, 93(3):109–115, 2005.

[14] Steven Roman. *Lattices and ordered sets*. Springer Science & Business Media, 2008.

[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 Proof of Theorem 1 (Algebraic Properties of the $\oplus$ operation)

A set L with some operation $*$ is a semilattice if the operation $*$ is commutative, associative and idempotent. In the following theorems we prove that all these properties hold for the $\oplus$ operation on set $\mathbb{T}$.

Theorem 9. *Operator $\oplus$ is commutative.*

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

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

□

Theorem 10. *Operation $\oplus$ is idempotent.*

Proof.

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

□

Theorem 11. *Operation $\oplus$ is associative.*

Proof. We will prove that $((\mathcal{E}, A) \oplus (\mathcal{F}, B)) \oplus (\mathcal{H}, C) = (\mathcal{E}, A) \oplus ((\mathcal{F}, B) \oplus (\mathcal{H}, C))$. We have that

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

and

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

therefore it suffices to show the following equivalence

$$\overbrace{(Z_1 \cap B = Z_2 \cap A)}^{(1)} \wedge \overbrace{((Z_1 \cup Z_2) \cap C = Z_3 \cap (A \cup B))}^{(2)}$$
$$\Leftrightarrow$$
$$\underbrace{(Z_2 \cap C = Z_3 \cap B)}_{(3)} \wedge \underbrace{((Z_2 \cup Z_3) \cap A = Z_3 \cap (B \cup C))}_{(4)}$$

First we prove the "⇒" direction.

$$\begin{aligned}(2) \Rightarrow & (Z_1 \cup Z_2) \cap C \cap B = Z_3 \cap (A \cup B) \cap B \\ \Rightarrow & (Z_1 \cap B \cap C) \cup (Z_2 \cap B \cap C) = Z_3 \cap B \\ \overset{(1)}{\Rightarrow} & (Z_2 \cap A \cap C) \cup (Z_2 \cap C) = Z_3 \cap B \\ \Rightarrow & Z_2 \cap C = Z_3 \cap B \end{aligned} \tag{5}$$

which proves that (3) holds. In the same way we obtain that

$$\begin{aligned}(2) \Rightarrow & (Z_1 \cup Z_2) \cap C \cap A = Z_3 \cap (A \cup B) \cap A \\ \Rightarrow & (Z_1 \cap A \cap C) \cup (Z_2 \cap A \cap C) = Z_3 \cap A \\ \overset{(1)}{\Rightarrow} & (Z_1 \cap C) \cup (Z_1 \cap B \cap C) = Z_3 \cap A \\ \Rightarrow & Z_1 \cap C = Z_3 \cap A \end{aligned} \tag{6}$$

Finally we prove the validity of equation (4),

$$\begin{aligned}(Z_2 \cup Z_3) \cap A = (Z_2 \cap A) \cup (Z_3 \cap A) \overset{(1),(6)}{=} (Z_1 \cap B) \cup (Z_1 \cap C) = \\ = Z_1 \cap (B \cup C)\end{aligned}$$

The proof for the "⇒" is complete, the other direction follows from symmetry. □