

Reliable Broadcast with Respect to Topology Knowledge

Aris Pagourtzis

Giorgos Panagiotakos

Dimitris Sakavalas

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

Abstract

We study the Reliable Broadcast problem in incomplete networks against a Byzantine adversary. We examine the problem under the *locally bounded adversary model* of Koo (2004) and the *general adversary model* of Hirt and Maurer (1997) and explore the tradeoff between the level of topology knowledge and the solvability of the problem.

We refine the local pair-cut technique of Pelc and Peleg (2005) in order to obtain impossibility results for every level of topology knowledge and any type of corruption distribution. On the positive side we devise protocols that match the obtained bounds and thus, exactly characterize the classes of graphs in which Reliable Broadcast is possible.

Among others, we show that Koo's Certified Propagation Algorithm (CPA) is *unique* against locally bounded adversaries in *ad hoc* networks, that is, it can tolerate as many local corruptions as any other non-faulty algorithm; this settles an open question posed by Pelc and Peleg. We also provide an adaptation of CPA against general adversaries and show its uniqueness in this case too. To the best of our knowledge this is the first optimal algorithm for Reliable Broadcast in generic topology *ad hoc* networks against general adversaries.

1 Introduction

A fundamental problem in distributed networks is Reliable Broadcast (Byzantine Generals), in which the goal is to distribute a message correctly despite the presence of Byzantine faults. That is, an adversary 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. In general, agreement problems have been primarily studied under the threshold adversary model, where a fixed upper bound t is set for the number of corrupted players and broadcast can be achieved if and only if $t < n/3$, where n is the total number of players. The Broadcast problem has been extensively studied in complete networks under the threshold adversary model mainly in the period from 1982, when it was introduced by Lamport, Shostak and Pease [11], to 1998, when Garay and Moses [5] presented the first fully polynomial Broadcast protocol optimal in resilience and round complexity.

The case of Reliable Broadcast under a threshold adversary in incomplete networks has been studied to a much lesser extent, in a study initiated in [1, 2, 10], mostly through protocols for Secure Message Transmission which, combined with a Broadcast protocol for complete networks, yield Broadcast protocols for incomplete networks. Naturally, connectivity constraints are required to hold in addition to the $n/3$ bound. Namely, at most $t < c/2$ corruptions can be tolerated, where c is network connectivity, and this bound is tight[1].

In the case of an honest dealer, particularly meaningful in wireless networks, the impossibility threshold of $n/3$ does not hold; for example, in complete networks with an honest dealer the problem becomes trivial regardless of the number of corrupted players. However, in incomplete networks the situation is different. A small number of traitors (corrupted players) may manage to block the entire protocol if they control a critical part of the network, e.g. if they form a separator of the graph. It therefore makes sense to define criteria (or parameters) depending on the structure of the graph, in order to bound the number or restrict the distribution of traitors that can be tolerated.

An approach in this direction is to consider topological restrictions on the adversary's corruption capacity. We will first focus on local restrictions, the importance of which comes, among others, from the fact that they may be used to derive criteria which can be employed in *ad hoc* networks. Such a paradigm is the *t-locally bounded adversary model*, introduced in [9], in which at most a certain number t of corruptions are allowed in the neighborhood of every node.

The locally bounded adversarial model is particularly meaningful in real-life applications and systems. For example, in social networks it is more likely for an agent to have a quite accurate estimation of the maximum number of malicious agents that may appear in its neighborhood, than having such information, as well as knowledge of connectivity, for the whole network. In fact, this scenario applies to all kinds of networks, where each node is assumed to be able to estimate the number of traitors in its close neighborhood. It is also natural for these traitor bounds to vary among different parts of the network. Motivated by such considerations, in this work we will introduce a generalization of the *t-locally bounded model*.

1.1 Related Work

Considering *t-locally bounded adversaries*, Koo [9] proposed a simple, yet powerful protocol, namely the *Certified Propagation Algorithm* (CPA) (a name coined by Pelc and Peleg in [14]), and applied it to networks of specific topology. CPA is based on the idea that a set of $t + 1$ neighbors of a node always contain an honest one. Pelc and Peleg [14] considered the *t-locally bounded model* in generic graphs and gave a sufficient topological condition for CPA to achieve Broadcast. They also provided an upper bound on the number of corrupted players t that can be locally tolerated in order to achieve Broadcast by any protocol, in terms of an appropriate graph parameter; they left the deduction of tighter bounds as an open problem. To this end, Ichimura and Shigeno [8] proposed an efficiently computable graph parameter which implies a more tight, but not exact, characterization of the class of graphs on which CPA achieves Broadcast. It had remained open until very recently to derive a tight parameter revealing the maximum number of traitors that can be locally tolerated by CPA in a graph G with dealer D . Such a parameter is implicit in the work of Tseng *et al.* [16], who gave a

necessary and sufficient condition for CPA Broadcast. Finally, in [12] such a graph parameter was presented explicitly, together with an efficient 2-approximation algorithm for computing its value.

A more general approach regarding the adversary structure was initiated by Hirt and Maurer in [7] where they studied the security of multiparty computation protocols with respect to an *adversary structure*, i.e. a family of sets of players, such that the adversary may entirely corrupt any set in the family. This line of work has yielded results on Broadcast against a general adversary in complete networks [4] but, to the best of our knowledge, the case of Broadcast against general adversaries in incomplete networks has not been studied as such.¹ A study on the related problem of Iterative Approximate Byzantine Consensus against general adversaries can be seen in [15] where a similar model for the *ad hoc* case is considered.

1.2 Our Results

In this work we study the tradeoff between the level of topology knowledge and the solvability of the problem, under various adversary models.

We first consider a natural generalization of the t -locally bounded model, namely the *non-uniform t -locally bounded model* which subsumes the (uniform) model studied so far. The new model allows for a varying bound on the number of corruptions in each player's neighborhood. We address the issue of locally resilient Broadcast in the non-uniform model. We present a new necessary and sufficient condition for CPA to be t -locally resilient by extending the notion of *local pair cut* of Pelc and Peleg [14] to the notion of *partial local pair cut*. Note that although equivalent conditions exist [16, 12], the simplicity of the new condition allows to settle the open question of CPA Uniqueness [14] in the affirmative: we show that if any *safe* (non-faulty) algorithm achieves Broadcast in an *ad hoc* network then so does CPA. We next prove that computing the validity of the condition is NP-hard and observe that the latter negative result also has a positive aspect, namely that a polynomially bounded adversary is unable to design an optimal attack unless $P = NP$.

We next shift focus on networks of known topology and devise an optimal resilience protocol, which we call *Path Propagation Algorithm* (PPA). Using PPA we prove that a topological condition which was shown in [14] to be necessary for the existence of a Broadcast algorithm is also sufficient. Thus, we manage to exactly characterize the class of networks for which there exists a solution to the Broadcast problem. On the downside, we prove that it is NP-hard to compute an essential decision rule of PPA, rendering the algorithm inefficient. However, we are able to provide an indication that probably no efficient protocol of optimal resilience exists, by showing that efficient algorithms which behave exactly as PPA w.r.t. decision do not exist if $P \neq NP$.

We then take one step further, by considering a hybrid between *ad hoc* and known topology networks: each node knows a part of the network, namely a connected subgraph containing itself. We propose a protocol for this setting as well, namely the *Generalized Path Propagation Algorithm* (GPPA). We use GPPA to show that this *partial knowledge* model allows for Broadcast algorithms of increased resilience.

Finally, we study the general adversary model and show that an appropriate adaptation of CPA is unique against general adversaries in *ad hoc* networks. To the best of our knowledge this is the first algorithm for Reliable Broadcast in generic topology *ad hoc* networks against a general adversary. We show an analogous result for known topology networks, which however can be obtained implicitly from [10] as mentioned above.

We conclude by discussing how to extend our results to the case of a corrupted dealer by simulating Broadcast protocols for complete networks.

A central tool in our work is a refinement of the local pair-cut technique of Pelc and Peleg [14] which proves to be adequate for the exact (in most cases) characterization of the class of graphs for which Broadcast is possible for any level of topology knowledge and type of corruption distribution. A useful by-product of practical interest is that the refined cuts can be used to determine the exact subgraph in which Broadcast is possible.

¹Some related results are implicit in [10], but in the problem studied there, namely Secure Message Transmission, additional secrecy requirements are set which are out of the scope of our study.

For clarity we have chosen to present our results for the t -local model first (Sections 3,4,5), for which proofs and protocols are somewhat simpler and more intuitive, and then for the more involved general adversary model (Section 6).

2 Problem and Model Definition

In this paper we address the problem of *Reliable Broadcast with an honest dealer* in generic (incomplete) networks. As we will see in Section 7, this case essentially captures the difficulty of the general problem, where even the dealer may be corrupted. The problem definition follows.

Reliable Broadcast with Honest Dealer. The network is represented by a graph $G = (V, E)$, where V is the set of players, and E represents authenticated channels between players. We assume the existence of a designated honest player, called the *dealer*, who wants to broadcast a certain value $x_D \in X$, where X is the initial input space, to all players. We say that a distributed protocol achieves Reliable Broadcast if by the end of the protocol every honest player has *decided on* x_D , i.e. if it has been able to deduce that x_D is the value originally sent by the dealer and output it as its own decision.

The problem is trivial in complete networks; we will consider the case of incomplete networks here. For brevity we will refer to the problem as the Broadcast problem.

We will now formally define the adversary model by generalizing the notions originally developed in [9, 14]. We will also define basic notions and terminology that we will use throughout the paper. We refer to the participants of the protocol by using the terms *node* and *player* interchangeably.

Corruption function. Taking into account that each player might be able to estimate her own upper bound on the corruptions of its neighborhood, as discussed earlier, we introduce a model in which the maximum number of corruptions in each player's neighborhood may vary from player to player. We thus generalize the standard t -locally bounded model [9] in which a uniform upper bound on the number of local corruptions was assumed. Here we consider $t : V \rightarrow \mathbb{N}$ to be a *corruption function* over the set of players V .

Non-Uniform t -Locally Bounded Adversary Model. The network is represented by a graph $G = (V, E)$. One player $D \in V$ is the dealer (sender). A corruption function $t : V \rightarrow \mathbb{N}$ is also given, implying that an adversary may corrupt at most $t(u)$ nodes in the neighborhood $\mathcal{N}(u)$ of each node $u \in V$. The family of t -local sets plays an important role in our study since it coincides with the family of admissible corruption sets.

Definition 1 (t -local set). *Given a graph $G = (V, E)$ and a function $t : V \rightarrow \mathbb{N}$ a t -local set is a set $C \subseteq V$ for which $\forall u \in V$, $|\mathcal{N}(u) \cap C| \leq t(u)$. For $V' \subseteq V$ a t -local w.r.t. V' set is a set $C \subseteq V$ for which $\forall u \in V'$, $|\mathcal{N}(u) \cap C| \leq t(u)$.*

Uniform vs Non-Uniform Model. Obviously the original t -locally bounded model corresponds to the special case of t being a constant function. Hereafter we will refer to the original t -locally bounded model as the *Uniform Model* as opposed to the *Non-Uniform Model* which we introduce here.

In our study we will often make use of node-cuts which separate some players 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 node-cut. The notion of t -local pair cut was introduced in [14] and is crucial in defining the bounds for which correct dissemination of information in a network is possible.

Definition 2 (t -local pair cut). *Given a graph $G = (V, E)$ and a function $t : V \rightarrow \mathbb{N}$, a pair of t -local sets C_1, C_2 s.t. $C_1 \cup C_2$ is a cut of G is called a t -local pair cut.*

The next definition extends the notion of t -local pair cut and is particularly useful in describing capability of achieving Broadcast in networks of unknown topology (*ad hoc* networks) where each player's knowledge of the topology is limited in its own neighborhood.

Definition 3 (*t*-partial local 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 *t*-partial local pair cut (*t*-plp cut) if there exists a partition $C = C_1 \cup C_2$ where C_1 is *t*-local and C_2 is *t*-local w.r.t. B .

In the uniform model the *Local Pair Connectivity* ($\text{LPC}(G, D)$) [14] parameter of a graph G with dealer D , was defined to be the minimum integer t s.t. G has a *t*-local pair cut. To define the corresponding notion in the non-uniform model we need to define a (partial) order among corruption functions. Nevertheless, for reasoning about our results it suffices to consider the following decision problem:

Definition 4 (pLPC). Given a graph G , a dealer D and a corruption function t determine whether there exists a *t*-plp cut in G .

Definition 5 (*t*-locally resilient algorithm). An algorithm which achieves Broadcast for any *t*-local corruption set in graph G with dealer D is called *t*-locally resilient for (G, D) .

Definition 6 (safe / *t*-locally safe algorithm). A Broadcast algorithm which never causes an honest node to decide on an incorrect value, is called safe.

A Broadcast algorithm which never causes an honest node to decide on an incorrect value under any *t*-local corruption set, is called *t*-locally safe.

3 Ad Hoc Networks

3.1 Certified Propagation Algorithm (CPA)

The Certified Propagation algorithm [9] uses only local information and thus is particularly suitable for *ad hoc* networks. CPA is probably the only Broadcast algorithm known up to now for the *t*-locally bounded model, which does not require knowledge of the network topology. Protocol 2, presented in the Appendix, is a modification of the original CPA that can be employed under the generalized corruption model introduced here. Namely a node v , upon reception of $t(v) + 1$ messages with the same value x from $t(v) + 1$ distinct neighbors, decides on x , sends it to all neighbors and terminates. It can easily be proven by induction that CPA is a *t*-locally safe Broadcast algorithm.

3.2 CPA Uniqueness in Ad Hoc Networks

Based on the above definitions we can now prove the *CPA uniqueness conjecture* for *ad hoc* networks, which was posed as an open problem in [14]. The conjecture states that no algorithm can locally tolerate more corrupted nodes than CPA in networks of unknown topology.

We consider only the class of *t*-locally safe Broadcast algorithms. We assume the *ad hoc* network model, as described e.g. in [14]. In particular we assume that nodes know only their own labels, the labels of their neighbors and the label of the dealer. We call a distributed Broadcast algorithm that operates under these assumptions an *ad hoc Broadcast algorithm*.

Theorem 1 (Sufficient Condition). Given a graph G , a corruption function t and a dealer D , if no *t*-plp cut exists, then CPA is *t*-locally resilient for (G, D) .

Proof. Suppose that no *t*-plp cut exists in G . Let T be the corruption set; clearly $T \cup N(D)$ is a cut on G as defined before (i.e. not including node D). Since T is *t*-local and $T \cup N(D)$ is not a *t*-plp cut there must exist $u_1 \in V \setminus (T \cup N(D) \cup D)$ s.t. $|N(u_1) \cap (N(D) \setminus T)| \geq t(u_1) + 1$. Since u_1 is honest it will decide on the dealer's value x_D . Let us now use the same argument inductively to show that every honest node will eventually decide on the correct value x_D through CPA. Let $C_k = (N(D) \setminus T) \cup \{u_1, u_2, \dots, u_{k-1}\}$ be the set of the honest nodes that have decided until a certain round of the protocol. Then $C_k \cup T$ is a cut. Since T is *t*-local, by the same argument as before there exists a node u_k s.t. $|C_k \cap N(u_k)| \geq t(u_k) + 1$ and u_k will decide on x_D . Eventually all honest players will decide on x_D . Thus CPA is *t*-locally resilient in G . \square




Figure 1: Graphs G and G'

Theorem 2 (Necessary Condition). *Let \mathcal{A} be a t -locally safe ad hoc Broadcast algorithm. Given a graph G , a corruption function t and a dealer D , if a t -plp cut exists, then \mathcal{A} is not t -locally resilient in (G, D) .*

Proof. Assume that there exists a t -plp cut $C = T \cup H$ in graph G with dealer D with T being the t -local set of the partition and H the t -local w.r.t. to B set (Figure 1). Let G' be a graph that results from G if we remove some edges that connect nodes in $A \cup T \cup H$ with nodes in H so that the set H becomes t -local in G' (e.g. we can remove all edges that connect nodes in $A \cup T \cup H$ with nodes in H). Note that the existence of a set of edges that guarantees such a property is implied by the fact that H is t -local w.r.t. B .

The proof is by contradiction. Suppose that there exists a t -locally safe Broadcast algorithm \mathcal{A} which is t -locally resilient in graph G with dealer D . We consider the following executions σ and σ' of \mathcal{A} :

- Execution σ is on the graph G with dealer D , with dealer's value $x_D = 0$, and corruption set T ; in each round, all players in T perform the actions that perform in the respective round of execution σ' (where T is a set of honest players).
- Execution σ' is on the graph G' with dealer D , with dealer's value $x_D = 1$, and corruption set H ; in each round, all players in H perform the actions that perform in the respective round of execution σ (where H is a set of honest players).

Note that T, H are admissible corruption sets in G, G' respectively due to their t -locality. It is easy to see that $H \cup T$ is a cut which separates D from B in both G and G' 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 t -locally resilient on G with dealer D , w must decide on the dealer's message 0 in execution σ on G with dealer D , and must do the same in execution σ' on G' with dealer D . However, in execution σ' the dealer's message is 1. Therefore \mathcal{A} makes w decide on an incorrect message in (G', D) . This contradicts the assumption that \mathcal{A} is locally safe. \square

We can show that if we drop the requirement for t -local safety, then the theorem does not hold. Intuitively, the reason is that an *ad hoc* protocol that assumes certain topological properties for the network may be t -locally resilient in a family of graphs that have the assumed topological properties. Indeed, Pelc and Peleg [14] introduced another algorithm for the uniform model, the *Relaxed Propagation Algorithm* (RPA) which uses knowledge of the topology of the network and they proved that there exists a graph G'' with dealer D for which RPA is 1-locally resilient and CPA is not. So if we use RPA in an *ad hoc* setting assuming that the network is G'' then this algorithm will be t -locally resilient for (G'', D) while CPA will not. Non- t -local safety of RPA can easily be shown. This shows that there exists non-safe algorithms of higher resilience than CPA.

Corollary 3 (CPA Uniqueness). *Given a graph G and dealer D , if there exists an ad hoc Broadcast algorithm which is t -locally resilient in (G, D) and t -locally safe, then CPA is t -locally resilient in (G, D) .*

Proof. Immediate from Theorems 1,2. □

3.3 Hardness of pLPC

Ichimura and Shigeno in [8] prove that the *set splitting* problem, known as NP-hard [6], can be reduced to the problem of computing the minimum integer t such that a t -local pair cut exists in a graph G . By generalizing the notion of the t -local pair cut to that of t -plp cut and defining the pLPC problem analogously one can use a nearly identical proof to that of [8] and show that the pLPC problem is NP-hard. For completeness the proof is given in the appendix.

Theorem 4. *pLPC is NP-hard.*

Therefore, computing the necessary and sufficient condition for CPA to work is NP-hard. Observe that this negative result also has a positive aspect, namely that a polynomially bounded adversary is unable to compute an optimal attack unless $P = NP$.

4 Known topology Networks

4.1 The Path Propagation Algorithm

Considering only safe Broadcast algorithms, the uniqueness of CPA in the *ad hoc* model implies that an algorithm that achieves Broadcast in cases where CPA does not, must operate under a weaker model e.g., assuming additional information on the topology of the network. It thus makes sense to consider the setting where players have full knowledge of the topology of the network. In this section we propose the *Path Propagation Algorithm* (PPA) and show that is of optimal resilience in the full-knowledge model. For convenience we will use the following notions: a set $S \subseteq V \setminus D$ is called a *cover* of a set of paths \mathcal{P} if and only if $\forall p \in \mathcal{P}, \exists s \in S$ s.t. $s \in p$ (s is a node of p). With $tail(p)$ we will denote the last node of path p . The description of PPA follows.

Protocol 1: *Path Propagation Algorithm (PPA)*

Input (for each node v): dealer's label D , graph G , $t(v) = \max \#$ corruptions in $N(v)$.

Message format: pair (x, p) , where $x \in X$ (message space), and p is a path of G (message's propagation trail).

Code for D : send message (x_D, D) to all neighbors, decide on x_D and terminate.

Code for $v \neq D$: upon reception of (x, p) from node u do:

if $(v \in p) \vee (tail(p) \neq u)$ then discard the message else send $(x, p||v)$ ² to all neighbours.
 if $decision(v) \neq \perp$ then send message $(decision(v), v)$ to all neighbors.

function $decision(v)$

(* dealer propagation rule *)

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

(* honest path propagation rule *)

if v receives messages $(x, p_1), \dots, (x, p_n)$ and $\nexists t$ -local cover of $\mathcal{P} = \{p_1, \dots, p_n\}$
 then return x else return \perp .

²By $p||v$ we denote the path consisting of path p and node v , with the last node of p connected to v .

The correctness of the honest path propagation rule is trivial: if a path is entirely corruption free, then value x , which is relayed through that path, is correct. Checking whether $\text{tail}(p) \neq u$ we ensure that at least one corrupted node will be included in a faulty path. Observe that each player can check the validity of the honest path propagation rule only if it has knowledge of the corruption function t and the network's topology.

4.2 A necessary and sufficient condition

We will now show that the non-existence of a t -local pair cut is a sufficient condition for PPA to achieve Broadcast in the t -locally bounded model in networks of known topology (proof in the Appendix).

Theorem 5 (Sufficiency). *Given a graph G with dealer D and corruption function t , if no t -local pair cut exists in (G, D) then all honest players will decide through PPA on x_D .*

Using the same arguments as in the proof of the necessity of condition $t < LPC(G, D)$ [14] it can be seen that the non-existence of a t -local pair cut is a necessary condition for any algorithm to achieve Broadcast under the non-uniform model.

Theorem 6 (Necessity). *Given a graph G with dealer D and corruption function t , if there exists a t -local pair cut in (G, D) then there is no t -locally resilient algorithm for (G, D) .*

Thus the non-existence of a t -local pair cut proves to be a necessary and sufficient condition for the existence of a t -locally resilient algorithm in both the uniform and the non-uniform model. Therefore PPA is of optimal resilience.

4.3 On the hardness of Broadcast in known networks

In order to run PPA we have to be able to deduce whether a corruption-free path exists among a set of paths broadcasting the same value. Formally, given a graph $G(V, E)$, a set of paths \mathcal{P} and a node u (the one that executes $\text{decision}(u)$) we need to determine whether there exists a t -local cover T of \mathcal{P} . We call this problem the Local Path Cover Problem, $LPCP(G, D, u, t, \mathcal{P})$ and show that is NP-hard (proof in the Appendix).

Theorem 7. *It is NP-hard to compute $LPCP(G, D, u, t, \mathcal{P})$.*

The above theorem implies that PPA may not be practical in some cases, since its decision rule cannot be always checked efficiently. It remains to show whether any other algorithm which has the same resiliency as PPA can be efficient. The following theorem provides an indication that the answer is negative, by showing that algorithms which behave exactly as PPA w.r.t. decision are unlikely to be efficient (proof in the Appendix).

Theorem 8. *Assuming $P \neq NP$, no safe fully polynomial protocol Π can satisfy the following: for any graph G , dealer D , corruption function t , and admissible corruption set C executing protocol Π_C , a node u decides through PPA on a value x iff u will decide on x by running Π on (G, D, t, C, Π_C) .*

5 Partial knowledge

Until now we have presented optimal resilience algorithms for Broadcast in two extreme cases, with respect to the knowledge over the network topology: the *ad hoc* model and the full-knowledge model. A natural question arises: is there any algorithm that works well in settings where nodes have partial knowledge of the topology?

To address this question we devise a new, generalized version of PPA that can run with partial knowledge of the topology of the network. More specifically we assume that each player v only has knowledge of the topology of a certain connected subgraph G_v of G which includes v . Namely if we consider the family \mathcal{G} of connected subgraphs of G we use the *topology view function* $\gamma : V \rightarrow \mathcal{G}$,

where $\gamma(v)$ represents the subgraph over which player v has knowledge of the topology. We also define the *joint view* of a set S as the subgraph $\gamma(S)$ of G with node-set $V(\gamma(S)) = \bigcup_{u \in S} V(\gamma(u))$ and edge-set $E(\gamma(S)) = \bigcup_{u \in S} E(\gamma(u))$. We will call an algorithm which achieves Broadcast for any t -local corruption set in graph G with dealer D and view function γ , (γ, t) -*locally resilient* for (G, D) .

Now given a corruption function t and a view function γ we define the Generalized Path Propagation Algorithm (GPPA) to work exactly as PPA apart from a natural modification of the path propagation rule.

Generalized path propagation rule: Player v receives the same value x from a set \mathcal{P} of paths that are completely inside $\gamma(v)$ and is able to deduce (from the topology) that no t -local cover of \mathcal{P} exists.

Remark. Note that GPPA generalizes both CPA and PPA. Indeed, if $\forall v \in V, \gamma(v) = \mathcal{N}(v)$, then $\text{GPPA}(G, D, t, \gamma)$ coincides with $\text{CPA}(G, D, t)$. If, on the other hand, $\forall v \in V, \gamma(v) = G$ then $\text{GPPA}(G, D, t, \gamma)$ coincides with $\text{PPA}(G, D, t)$.

We also notice that, quite naturally, as γ provides more information for the topology of the graph, resilience increases, with CPA being of minimal resilience in this family of algorithms, and PPA achieving maximal resilience.

To prove necessary and sufficient conditions for GPPA being t -locally resilient we need to generalize the notion of t -plp cut as follows:

Definition 7 (type 1 (γ, t) -partial local 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 will be called a type 1 (γ, t) -partial local pair cut (plp1 cut) if there exists a partition $C = C_1 \cup C_2$ s.t. C_1 is t -local and C_2 is t -local in the graph $\gamma(B)$.*

Definition 8 (type 2 (γ, t) -partial local 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 will be called a type 2 (γ, t) -partial local pair cut (plp2 cut) if there exists a partition $C = C_1 \cup C_2$ s.t. C_1 is t -local and $\forall u \in B, C_2 \cap N(u)$ is t -local in the graph $\gamma(u)$.*

We can now show the following two theorems. The proofs build on the techniques presented for CPA and PPA and are included in the Appendix.

Theorem 9 (sufficient condition). *Let t be corruption function and γ be a view function, if no (γ, t) -plp2 cut exists in G with dealer D then $\text{GPPA}(G, D, t, \gamma)$ is (γ, t) -locally resilient for G, D .*

Theorem 10 (necessary condition). *Let t be a corruption function, γ be a view function and \mathcal{A} be a t -locally safe ad hoc Broadcast algorithm. If a (γ, t) -plp1 cut exists in graph G with dealer D , then \mathcal{A} is not (γ, t) -locally resilient for G, D .*

One can argue that increased topology knowledge implies increased resilience for GPPA compared to CPA; for example, the sufficient condition of GPPA holds in settings where the sufficient condition of CPA does not hold. An overview of our results concerning the t -local model with respect to the level of topology knowledge appears in Figure 2.

Notice that the reason for which GPPA is not optimal is that nodes in $\gamma(v)$ do not share their knowledge of topology. An optimal resilience protocol would probably include exchange of topological knowledge among players.

6 General Adversary

Hirt and Maurer in [7] 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 Z are in \mathcal{Z} if $Z \in \mathcal{Z}$.

Let us now redefine some notions that we have introduced in this paper in order to extend our results to the case of a general adversary. We will call an algorithm that achieves Broadcast for any corruption set $T \in \mathcal{Z}$ in graph G with dealer D , \mathcal{Z} -*resilient*. We next generalize the notion of a t -local pair cut.

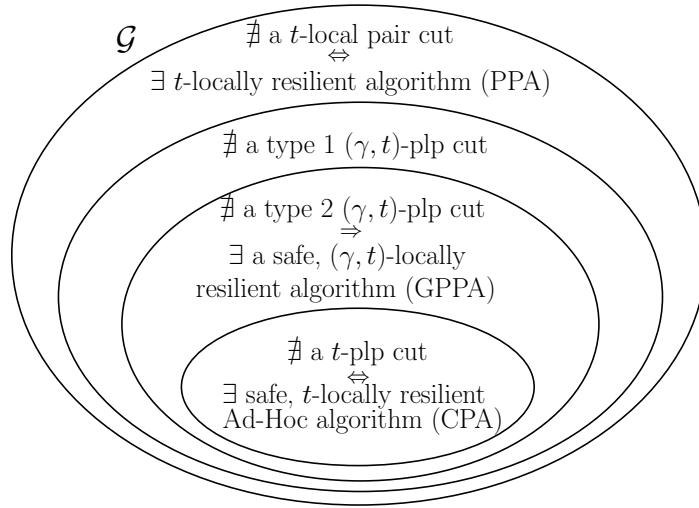


Figure 2: Overview of conditions concerning the existence of t -locally resilient algorithms with respect to the level of topology knowledge. Note that \mathcal{G} refers to the family of pairs (G, D) .

Definition 9 (\mathcal{Z} -pair cut). A cut C of G for which there exists a partition $C = C_1 \cup C_2$ and $C_1, C_2 \in \mathcal{Z}$ is called a \mathcal{Z} -pair cut of G .

Known Topology Networks

We adapt PPA in order to address the Broadcast problem under a general adversary. The Generalized \mathcal{Z} -PPA algorithm can be obtained by a modification of the path propagation rule of PPA (Protocol 1).

\mathcal{Z} -PPA Honest Path Propagation Rule: player v receives the same value x from a set \mathcal{P} of paths and is able to deduce that for any $T \in \mathcal{Z}$, T is not a cover of \mathcal{P} .

Moreover, the following theorems can be easily shown using essentially the same proofs as for Theorems 5, and 6 and replacing the notion of t -local pair cut with that of \mathcal{Z} -pair cut.

Theorem 11 (Sufficiency). Given a graph G , dealer D , and an adversary structure \mathcal{Z} , if no \mathcal{Z} -pair cut exists, then all honest players will decide on x_D through \mathcal{Z} -PPA.

Theorem 12 (Necessity). Given a graph G , dealer D , and an adversary structure \mathcal{Z} , if there exists a \mathcal{Z} -pair cut then there is no \mathcal{Z} -resilient Broadcast algorithm for (G, D) .

Ad Hoc Networks

Since in the *ad hoc* model the players know only their own labels, the labels of their neighbors and the label of the dealer it is reasonable to assume that a player has only local knowledge on the actual adversary structure \mathcal{Z} . Specifically, given the actual adversary structure \mathcal{Z} we assume that each player v knows only the *local adversary structure* $\mathcal{Z}_v = \{A \cap \mathcal{N}(v) : A \in \mathcal{Z}\}$.

As in known topology networks, we can describe a generalized version \mathcal{Z} -CPA of CPA, which is an *ad hoc* Broadcast algorithm for the general adversary model. In particular, we modify step (3) of CPA (Protocol 2) in the following way.

\mathcal{Z} -CPA Certified Propagation Rule: if a node v is not a neighbor of the dealer, then upon receiving the same value x from all its neighbors in a set $N \subseteq \mathcal{N}(v)$ s.t. $N \notin \mathcal{Z}_v$, it decides on value x .

In order to argue about the topological conditions which determine the effectiveness of \mathcal{Z} -CPA we generalize the notion of partial t -local pair 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 $\forall u \in B$, $\mathcal{N}(u) \cap C_2 \in \mathcal{Z}_u$.

Analogously to CPA Uniqueness, we can now prove \mathcal{Z} -CPA Uniqueness in the general adversary model (proofs in Appendix).

Theorem 13 (Sufficient Condition). *Given a graph G , dealer D , and an adversary structure \mathcal{Z} , if no \mathcal{Z} -pp cut exists, then \mathcal{Z} -CPA is \mathcal{Z} -resilient.*

Theorem 14 (Necessary Condition). *Let \mathcal{A} be a safe ad hoc Broadcast algorithm. Given a graph G , dealer D , and an adversary structure \mathcal{Z} , if a \mathcal{Z} -pp cut exists then \mathcal{A} is not \mathcal{Z} -resilient for G, D .*

Complexity of \mathcal{Z} -CPA. Regarding the computational complexity of \mathcal{Z} -CPA one can observe that it is polynomial if and only if for every player v there exists a polynomial (w.r.t. the size of G) algorithm \mathcal{B} which given a set $S \subseteq \mathcal{N}(v)$ decides whether $S \in \mathcal{Z}_v$. Since \mathcal{Z} -CPA is clearly polynomial in round complexity and communication complexity, if such an algorithm \mathcal{B} exists, \mathcal{Z} -CPA is fully polynomial.

7 Dealer Corruption

We have studied the problem of Broadcast in the case where the dealer is honest. In order to address the general case in which the dealer may also be corrupted one may observe that for a given adversary structure \mathcal{Z} and graph G , \mathcal{Z} -resilient Broadcast in *ad hoc* networks can be achieved if the following conditions both hold:

1. $\nexists Z_1, Z_2, Z_3 \in \mathcal{Z}$ s.t. $Z_1 \cup Z_2 \cup Z_3 = V$.
2. $\forall v \in V$ there does not exist a \mathcal{Z} -pp cut for G with dealer v .

Condition 1 was proved by Hirt and Maurer [7] sufficient and necessary for the existence of secure multiparty protocols in complete networks. \mathcal{Z} -resilient Broadcast in the general case where the network is incomplete can be achieved by simulating any protocol for complete graphs (e.g. the protocol presented in [4]) as follows: each one-to-many transmission is replaced by an execution of \mathcal{Z} -CPA. It is not hard to see that the conjunction of the above two conditions is necessary and sufficient for Broadcast in incomplete networks in the case of corrupted dealer. Similarly in networks of known topology, there exists a \mathcal{Z} -resilient Broadcast algorithm if condition 1 holds and for every $v \in V$ a \mathcal{Z} -pair cut does not exist for graph G with dealer v . Naturally, the above observations hold also in the special case of a locally bounded adversary.

8 Open questions

A number of questions arise from the results presented in this paper:

- Necessary and sufficient criteria for Broadcast on known topology and ad-hoc networks are NP-hard to compute. So what is the best attack a polynomially bounded adversary could deploy? Similar issues may be raised from the point of view of system designers. Defining a meaningful approximation objective is essential in answering such questions.
- We conjecture that in the known topology locally bounded setting no safe, fully polynomial algorithm can achieve optimal resiliency. We have provided an indication towards proving this.
- Regarding the partial knowledge model discussed in Section 5, GPPA is not of optimal resilience. Devising an algorithm with this property would be of great interest, especially in the general adversary model. One direction towards this, is to consider discovering network topology under a Byzantine adversary, as studied in [13, 3].
- In the *ad hoc* general adversary setting, we proved that \mathcal{Z} -CPA is unique, thus having optimal resilience. We conjecture that it is also unique w.r.t. polynomial time complexity, i.e., if a safe protocol achieves Broadcast in polynomial time then so does \mathcal{Z} -CPA.

References

- [1] Danny Dolev. The byzantine generals strike again. *J. Algorithms*, 3(1):14–30, 1982.
- [2] Danny Dolev, Cynthia Dwork, Orli Waarts, and Moti Yung. Perfectly secure message transmission. *J. ACM*, 40(1):17–47, January 1993.
- [3] 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.
- [4] Matthias Fitzi and Ueli M. Maurer. Efficient byzantine agreement secure against general adversaries. In Shay Kutten, editor, *DISC*, volume 1499 of *Lecture Notes in Computer Science*, pages 134–148. Springer, 1998.
- [5] Juan A. Garay and Yoram Moses. Fully polynomial byzantine agreement for $n > 3t$ processors in $t + 1$ rounds. *SIAM J. Comput.*, 27(1):247–290, 1998.
- [6] M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [7] 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.
- [8] 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.
- [9] Chiu-Yuen Koo. Broadcast in radio networks tolerating byzantine adversarial behavior. In Soma Chaudhuri and Shay Kutten, editors, *PODC*, pages 275–282. ACM, 2004.
- [10] 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.
- [11] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [12] 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.
- [13] Mikhail Nesterenko and Sébastien Tixeuil. Discovering network topology in the presence of byzantine faults. *IEEE Trans. Parallel Distrib. Syst.*, 20(12):1777–1789, 2009.
- [14] Andrzej Pelc and David Peleg. Broadcasting with locally bounded byzantine faults. *Inf. Process. Lett.*, 93(3):109–115, 2005.
- [15] Lewis Tseng and Nitin H. Vaidya. Iterative approximate byzantine consensus under a generalized fault model. In Davide Frey, Michel Raynal, Saswati Sarkar, Rudrapatna K. Shyamasundar, and Prasun Sinha, editors, *ICDCN*, volume 7730 of *Lecture Notes in Computer Science*, pages 72–86. Springer, 2013.
- [16] Lewis Tseng, Nitin H. Vaidya, and Vartika Bhandari. Broadcast using certified propagation algorithm in presence of byzantine faults. *CoRR*, abs/1209.4620, 2012.

Appendix

A The Certified Propagation Algorithm

Protocol 2: Certified Propagation Algorithm (CPA) for the Non-Uniform Model

Input (for each node v): Dealer's label D , labels of v 's neighbors, corruption bound $t(v)$.

Message format: A single value $x \in X$.

Code for D : send value $x_D \in X$ to all neighbors, decide on x_D and terminate.

Code for $v \in \mathcal{N}(D)$: upon reception of x_D from the dealer, decide on x_D , send it to all neighbors and terminate.

(* certified propagation rule *)

Code for $v \notin \mathcal{N}(D) \cup D$: upon reception of $t(v) + 1$ messages with the same value x from $t(v) + 1$ distinct neighbors, decide on x , send it to all neighbors and terminate.

B Proof of Theorem 4

Proof. We show that the *set splitting* problem known as NP-hard [6] can be reduced to the *pLPC* problem. Given a collection S of 3-element subsets of a finite set X , the set splitting problem asks whether there is a partition of X into two subsets X_1 and X_2 such that no subset in S is entirely contained in either X_1 or X_2 . Let $S+$ be a multiple collection adding dummy subsets $\{v\}$ to S such that the cardinality of $\{s \in S+ : v \in s\}$ is at least six for each $v \in X$. A complete graph with vertex set $S+$ and a copy of it are denoted by K_{S+} and K'_{S+} , respectively. We construct a graph G_{SSP} with vertex set $V(G_{SSP}) = V(K_{S+}) \cup V(K'_{S+}) \cup X$ and edge set $E(G_{SSP}) = E(K_{S+}) \cup E(K'_{S+}) \cup \{(v, s), (v, s') : v \in X, s \in S+, v \in s\}$, where s is a node in $V(K'_{S+})$ which is a copy of $s \in S+$. If a subgraph of G_{SSP} deleting $C (\subseteq V(G_{SSP}))$ has at least two connected components and $X \setminus C \neq \emptyset$, C contains $\mathcal{N}(v) \cap V(K_{S+})$ or $\mathcal{N}(v) \cap V(K'_{S+})$ for some $v \in X$. Since each $v \in X$ has at least six neighbor in both $V(K_{S+})$ and $V(K'_{S+})$, C is a t -local pair side cut with $t \geq 3$. We next consider the case of $C = X$. We can partition X into two 2-local sets in G_{SSP} , if and only if the set splitting problem has a desired partition X_1 and X_2 . Therefore, we have $pLPC(G_{SSP}, 2) = true$, if and only if the set splitting problem has a desired partition. Now we can easily show that NP-hardness for $pLPC(G, t)$ without a dealer implies NP-hardness for the case with a dealer. If $pLPC(G, t, D)$ could be solved with a polynomial-time algorithm then solving $pLPC(G, t, v)$ for every node in V would suffice to build a polynomial algorithm for $pLPC(G, t)$ which is a contradiction. Therefore to compute $pLPC(G, t, D)$ is NP-hard. \square

C Proof of Theorem 5

Proof. All players in $\mathcal{N}(D)$ decide due to the *dealer propagation rule*, since the dealer is honest. We next show the rest of the players will decide due to the *honest path propagation rule*.

Let v be any player in $V \setminus \mathcal{N}(D)$ and assume that no t -local pair cut exist in (G, D) exist. Let T be a t -local set and consider the execution σ_T of PPA where T is the corruption set. Let \mathcal{P} be the set of all paths connecting D with v and are composed entirely by nodes in $V \setminus T$ (honest nodes). Observe that $\mathcal{P} \neq \emptyset$, otherwise T is a cut separating D from v and T is trivially a t -local pair cut, a contradiction. Since paths in \mathcal{P} are entirely composed by honest nodes it is easy to see that v will receive the correct value through all paths in \mathcal{P} .

We next prove that under any t -local corruption set T' at least one path in \mathcal{P} is completely corruption free.




Figure 3: An instance of the reduction graph G for variables $\{x_1, x_2, x_3\}$ and clause $c_1 = \{x_1 \vee x_2 \vee \neg x_3\}$.

Assume that $\exists T' : t$ -local cover of \mathcal{P} . Then obviously $T \cup T'$ is a cut separating D from v , since every path that connects D with v contains at least a node in $T \cup T'$. Moreover the cut $T \cup T'$ can be partitioned in the sets $T \setminus T', T'$ which are trivially t -local and thus, $T \cup T'$ is a t -local pair cut, a contradiction. Hence, under any t -local corruption set T' at least one path in \mathcal{P} is entirely corruption free.

Consequently, in execution σ_T , node v will receive the correct value through every path in \mathcal{P} along with the corresponding propagation trail and will decide on the correct value due to the honest path propagation rule, because \mathcal{P} is not covered by any t -local set. Moreover honest nodes will not decide on the wrong value due to this rule, because the set of paths transmitting this value will always have a t -local cover (the corruption set T). \square

D Proof of Theorem 7

Proof. We will describe a reduction from $3SAT$ to $LPCP(G, D, u, t, \mathcal{P})$. For every variable x_i we construct a gadget G_{x_i} shown on the left of Figure 3. We will make use of a parameter μ that will serve as a constant corruption function (that is, our hardness result holds even for the uniform model). We will use several copies of the complete graphs $K_{\mu+1}$ and $K_{2\mu}$. Node D is connected to every vertex of a $K_{\mu+1}$ copy. Every vertex of that $K_{\mu+1}$ copy is connected with the ‘upper’ μ vertices of a $K_{2\mu}$ copy; let us call this ‘upper’ node set X_i . Symmetrically for the lower part, node u is connected to every vertex of another $K_{\mu+1}$ copy and every vertex of that $K_{\mu+1}$ copy is connected to the ‘lower’ μ vertices of $K_{2\mu}$, let us call this set X'_i . Now assuming that \mathcal{P} contains those paths in G_{x_i} that are of length 5 and connect D to u (and no other path in G_{x_i}) it is easy to show that :

Lemma 15. *If $LPCP(G, D, u, \mu, \mathcal{P}) = 1$ with μ -local cover T , then either $X_i \subseteq T$ or $X'_i \subseteq T$.*

$T \cap G_{x_i}$ is a cut of G_{x_i} . Since the only μ -local cuts in G_{x_i} are X_i and X'_i , the claim is immediate.

Now for every clause $c_i = c_{i_1} \vee c_{i_2} \vee c_{i_3}$ in C we construct the gadget shown on the right of Figure 3. Node D is connected to every vertex of $K_{\mu+1}$. Every vertex of $K_{\mu+1}$ is connected to the first literal of the clause, say l_{i_1} . Literal l_{i_1} is connected to l_{i_2} , and l_{i_2} to l_{i_3} . And symmetrically,

node u is connected to every vertex of another copy of $K_{\mu+1}$ and every vertex of $K_{\mu+1}$ is connected to l_{i_3} . Let us call this subgraph of G , G_{c_i} . Assuming that all paths from D to u of length 6 that go through G_{c_i} are contained in \mathcal{P} we show that:

Lemma 16. *if $LPCP(G, D, u, \mu, \mathcal{P}) = 1$ with μ -local cover T , then $l_{i_1} \in T$ or $l_{i_2} \in T$ or $l_{i_3} \in T$.*

The proof is by contradiction: if no l_{i_j} node belongs to T , then it must be $K_{\mu+1} \subseteq T$, contradicting the t -locality of T .

The last thing we need to establish is that if $X_i \subseteq T$ (respectively $X'_i \subseteq T$), no $\neg x_i$ (resp. x_i) literal of G_{c_j} is in T . We achieve this by adding a node v_{ij} connecting X_i (resp. X'_i) to $\neg x_i$ (resp. x_i) for each appearance of these literals in some G_{c_j} . The following holds because If both X_i and $\neg x_i$ are in T , then T is not μ -local since $|N(v_{ij}) \cap T| = \mu + 1$.

Lemma 17. *If $LPCP(G, D, u, \mu, \mathcal{P}) = 1$ with μ -local cover T , then $X_i \subseteq T$ (resp. $X'_i \subseteq T$) $\Rightarrow \neg x_i \notin T$ (resp. $x_i \notin T$).*

So for graph G that is constructed as described above and for path set \mathcal{P} consisting of the paths used for proving Lemmata 15 and 16 we have that $LPCP(G, D, u, \mu, \mathcal{P}) = 1$ iff there exists a truth assignment A which makes every clause in C true. The ' \Rightarrow ' direction follows from the lemmata proved above. The truth assignment A is constructed as follows: if $X_i \subseteq T$ (resp. $X'_i \subseteq T$) then $\neg x_i$ (resp. x_i) is true in A . The ' \Leftarrow ' direction comes naturally by setting T contain X_i if x_i is true by A , otherwise T contains X'_i ; T also contains all literals in G_{c_j} that are set true by A . Then T is a μ -local cover of \mathcal{P} and $LPCP(G, D, u, \mu, \mathcal{P}) = 1$. \square

E Proof of Theorem 8

Proof. We will show that if such Π existed then it would be a polynomial time solver for the 3-SAT problem. Let us consider what happens when Π is run on the graph G that we used in the proof of Theorem 7, with dealer D and the corrupted nodes being the ones that connect the ‘‘clause’’ gadgets with the ‘‘variable’’ gadgets (e.g. $C = \{v_1, v_2, v_3\}$ in Figure 3). The adversary protocol Π_C is: the corrupted nodes don't send or relay any messages.

The 3-SAT instance used to make G has a solution iff $LPCP(G, D, u, t, \mathcal{P}) = 1$, i.e. a μ -local cover C_1 on \mathcal{P} exists, where \mathcal{P} is the set of paths we used in the proof of Theorem 7. It can be seen from the decision rule of PPA that, while running PPA on G , u will not decide on any value iff a μ -local cover C_1 on \mathcal{P} exists. Moreover a node u does not decide through PPA on a value x iff u does not decide on x by running Π on (G, D, t, C, Π_C) .

So u decides on x_D while running Π on G , with dealer D and corruption set C which runs the Π_C protocol iff 3-SAT does not have a solution. Apparently if Π existed then 3-SAT would have a polynomial time solver. \square

F Proof of Theorem 9

Proof. Suppose no (γ, t) -plp2 cut exists. $T \cup N(D)$ is a cut on G non including node D . From the definition of (γ, t) -plp2 cut we have that there exists $u_1 \in V \setminus (T \cup N(D) \cup D)$ s.t. $N(D) \cap N(u)$ is not t -local on $\gamma(u_1)$. But since all the honest nodes in $N(D) \cap N(u)$ have decided, u will receive the value x_D from paths starting from these nodes of length 1. Finding a t -local corruption set covering these paths is impossible since it would have to include all these nodes, and from the above, it would not be t -local. So u_1 will decide on the dealer's value x_D . We can use the same argument inductively to show that every honest node will eventually decide on the correct value x_D through GPPA. Let $C_k = (N(D) \setminus T) \cup \{u_1, u_2, \dots, u_{k-1}\}$ be the set of the nodes that have decided until a certain round of the protocol. Then $C_k \cup T$ is a cut. Since T is t -local by the same argument as before there exists an undecided node u_k s.t. $C_k \cap N(u_k)$ is not t -local on $\gamma(u_k)$. Using the same argument as before u_k will decide on the correct value. Eventually all honest players will decide on x_D . Thus GPPA is t -locally resilient in G . \square

G Proof of Theorem 10

Proof. Assume that there exists a (γ, t) -plp1 cut $C = T \cup H$ in graph G with dealer D and with T being the t -local set of the partition (Figure 1). $\gamma(B)$ is the joint view of the nodes in B . G' is the graph that results from G if we remove edges from $A \setminus \gamma(B)$ s.t. the set H becomes t -local in G' . The existence of a set of edges that guarantees such a property is implied by the second property of the (γ, t) -plp1 cut. Suppose that there exists a t -locally safe Broadcast algorithm \mathcal{A} which is t -locally resilient in graph G with dealer D . We can argue the same way we did on Theorem 2 which leads to a contradiction. \square

H Proof of Theorem 13

Proof. Suppose that \mathcal{Z} -CPA is not \mathcal{Z} -resilient. Then we can split the graph in 3 parts: A being the honest decided nodes, B being the honest undecided nodes 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 a \mathcal{Z} -pp cut which is a contradiction. Hence, \mathcal{Z} -CPA is \mathcal{Z} -resilient. \square

I Proof of Theorem 14

Proof. Let $C = C_1 \cup C_2$ be the \mathcal{Z} -pp cut which partitions $V \setminus C$ in sets $A, B \neq \emptyset$ s.t. $D \in A$. 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 a node in B could decide with \mathcal{Z} being the adversary structure. Then using the standard argument employed in Theorem 2, an attack on the safeness of the algorithm would be possible in the same setting with \mathcal{Z}' being the adversary structure. The details of the proof are similar and are based on the difficulty of the honest players in B to distinguish which scenario they participate in, with respect to the adversary structure: the one with \mathcal{Z} or the one with \mathcal{Z}' . \square