

Pseudonymous Broadcast and Secure Computation from Cryptographic Puzzles

JONATHAN KATZ*[†]

ANDREW MILLER*

ELAINE SHI*[‡]

Abstract

In standard models of secure computation, point-to-point channels between parties are assumed to be authenticated by some pre-existing means. In other cases, even stronger pre-existing setup—e.g., a public-key infrastructure (PKI)—is assumed. These assumptions are too strong for open, peer-to-peer networks, where parties do not necessarily have any prior relationships and can come and go as they please. Nevertheless, these assumptions are made due to the prevailing belief that nothing “interesting” can be achieved without them.

Taking inspiration from Bitcoin, we show that precise bounds on computational power can be used in place of pre-existing setup to achieve weaker (but nontrivial) notions of security. Specifically, under the assumption that each party can solve cryptographic puzzles only at a bounded rate (and the existence of digital signatures), we show that *without prior setup* and *with no bound on the number of corruptions*, a group of parties can agree on a PKI with which they can then realize *pseudonymous* notions of authenticated communication, broadcast, and secure computation. Roughly, “pseudonymous” here means that inputs/outputs are bound to pseudonyms rather than parties’ true identities.

*Dept. of Computer Science, University of Maryland. **Email:** {jkatz, amiller, elaine}@cs.umd.edu.

[†]Work supported by NSF awards #0964541, #1111599, and #1223623.

[‡]Work supported by NSF award #1314857, a Sloan Research Fellowship, and a Google Faculty Research Award.

1 Introduction

The majority of the distributed-computing literature assumes a model in which parties have well-known identities and, moreover, there exists *authenticated*, point-to-point channels between each pair of parties. Such authentication is established either through some physical property of the underlying network, or using pre-existing setup assumptions such as a public key infrastructure (PKI) or pairwise shared keys. It is widely believed that without pre-existing identities and authenticated channels, the scope of meaningful tasks that can be achieved is limited (see, e.g., [5, 28]); furthermore, without a full PKI (e.g., with only pairwise shared keys), protocols can only tolerate a small constant fraction of faulty parties. [29]

The recent phenomenal success of Bitcoin [27] and other decentralized cryptocurrencies, however, has stirred vibrant interest in a new distributed-computing model where parties do not have well-known identities or pre-existing trust. Bitcoin offers empirical evidence that interesting security properties *can* be achieved in such networks, assuming computational resource bounds.

In this paper, we initiate a theoretical treatment of Bitcoin-like networks, to understand what is feasible in such distributed systems. We show that, indeed, a wide class of tasks can be realized without setup, but based instead on the assumption that an adversary has *bounded computational resources*. In more detail, we show feasibility of the following:

Consensus on a set of public keys. Our fundamental idea is to first bootstrap a (pseudonymous) PKI, and then realize (pseudonymous) secure computation using this PKI. To establish the PKI, we devise a protocol for Interactive Set Agreement (ISC) among the parties. Every honest party inputs their public key into this ISC protocol, and at the end of the protocol, all honest parties agree on a set (“*agreement*”). We prove that every honest party’s public key will be included in this set (“*validity*”); and moreover, an adversary cannot introduce more public keys disproportionate to the amount of computational resource it possesses (“*boundedness*”).

To design this consensus protocol, we face a few challenges. First, an adversary can try to introduce many fake identities (“Sybils”); intuitively, we defeat this by requiring the adversary to solve computational puzzles. Second, we assume a network model where honest parties’ messages are guaranteed to be delivered to all parties within a bounded delay; however, an adversary may specify a message to be delivered to only a subset of the honest parties. In this way, the adversary can try to cause honest parties to reach inconsistent decisions. Intuitively, we defeat the second type of attack through techniques reminiscent of the Dolev-Strong authenticated Byzantine agreement [12]. By combining these ideas in non-trivial manners, we craft novel protocols that utilize either a *sequential* puzzle (that cannot be parallelized) or a *parallelizable* puzzle, and ensure security against an *arbitrary* number of corrupted parties.

Pseudonymous secure multi-party computation (MPC). Once we instantiate a PKI through our interactive set agreement protocol, we hope to leverage it to establish authenticated channels, and then in turn, implement *broadcast*. Based on these, we hope to realize general secure multi-party computation using known techniques [18, 6].

However, a challenge arises in formalizing secure computation in this setting. Our ISC protocol ensures that each party can establish a public key (i.e., *pseudonym*) and must stick to this pseudonym throughout, but does not bind a party’s pseudonym to its true identity. However, standard notions of secure computation require *a priori* knowledge of parties’ true identities—specifically, to define the functionality to be computed, the ideal world typically refers to each party’s input by their true identities. To address this discrepancy, we formalize a new model of

secure computation called *pseudonymous* secure computation. We then show that our ISC protocol allows us to securely realize a pseudonymous PKI, from which we can implement broadcast primitives and secure computation for general functionalities with pseudonymous security.

Scope. We stress that our work is not intended to model the actual Bitcoin protocol, and we readily admit that in our initial attempt to provide a formal and clean treatment we have made some simplifying assumptions. For example, we assume that the network guarantees (bounded-delay) delivery of messages to every party and that the number of parties is fixed (and known) at the outset. Relaxing these assumptions is left for future work. Regardless, our results show that distributed-computation models based on assumptions about computational resources rather than authenticated identities are powerful, and can lead to a broad class of applications.

Related work. Cryptographic puzzles (also called *proofs of work*) have been analyzed and proposed for a variety of other purposes, including timed-release encryption [30], spam prevention [4, 13, 23], DoS resistance [20, 21], and defense against Sybil attacks [7]. Aspnes et al. [3] studied Byzantine agreement without a PKI in a model where computational puzzles can be solved at some bounded rate. That work also assumes pre-existing authenticated channels between honest parties. Moreover, their feasibility results do not extend to an unbounded number of corruptions. The success of Bitcoin [27] has renewed interest in the use of proof-of-work puzzles for coordinating peer-to-peer networks. For example, both [26] and [17] present formal models in which the core Bitcoin protocol can be shown to achieve relaxed notions of consensus, again under the assumption that a majority of the network is honest.

Okun [28] studied distributed-computing problems in several variant models of anonymous networks. The weakest model considered there, the “port-unaware” model, is most similar to ours. However, our model is weaker still: in the port-unaware model, each corrupt process can send at most one message to a correct process in each round, whereas in our model the adversary can deliver an arbitrary number messages to honest parties. Okun’s positive result crucially relies on this message bound, and is thus inapplicable in our model. Similarly our communication model is related to (but weaker than) the models considered in [10, 11].

Our positive result for pseudonymous secure computation are most closely related to the prior work of Barak et al. [5], who show that a weak form of secure computation can be realized without authenticated channels. Roughly speaking, secure computation in their setting means that an adversary can partition the honest parties into disjoint sets and then run the protocol with each of those subsets, in each case substituting inputs of its choice for the inputs of the other honest parties. The guarantees we provide are stronger, albeit under a stronger assumption. Specifically, our notion of pseudonymous secure computation ensures that all honest parties’ inputs are incorporated into a *single* computation; moreover, it ensures unanimous abort and, in the case of an honest majority, guaranteed output delivery. We also show how to achieve pseudonymous broadcast even with an unbounded number of corruptions.

Secure computation without a PKI, but when $1/3$ or more of the parties can be corrupted, is also studied in [15, 16, 19]. Some of these works also achieve secure computation with unanimous abort, but none realize broadcast (with guaranteed output delivery) in the sense we do here. More importantly, all these works assume pre-existing authenticated channels between parties.

Concurrent and independent work. Concurrently with our work, Andrychowicz and Dziembowski [2] show related results in a model that is similar to ours; however, we identify several differences. At a technical level, our models are incomparable: we assume the adversary can not

pre-compute puzzle solutions (in practice, this would be achieved using a random beacon); they do not make this assumption, but instead assume a bound on the number of messages an adversary can send in each round. They consider only parallelizable puzzles, whereas we consider both sequential and parallelizable puzzles. The round complexity of their ISC protocol is effectively $O(\kappa^2 f)$, where κ is a security parameter, which is incomparable to the round complexity of our protocol that uses parallelizable puzzles. Finally, they do not formalize pseudonymous secure computation.

2 Our Model

Our basic underlying model is derived from a standard setting for secure computation (namely the stand-alone model from Canetti [8]). However, our model fundamentally differs in that *we do not assume authenticated channels* between the parties. In addition, we work in a hybrid world where the parties can access a functionality that models their ability to solve a bounded number of computational puzzles per round—here, assumed for simplicity to be one. We provide further details in what follows.

Network model. Our underlying model consists of n parties in a fully connected, synchronous network. Since we do not assume authenticated channels, however, some aspects of the model bear explanation. The fact that there are no authenticated channels means that when a party receives a message, it cannot tell from which other party that message originated, or whether a message it receives in one round is from the same party as some other message it received in a different round. The only thing we assume is that any message sent by an honest party in some round is received by all other parties (including the sender himself) at the end of that round. Honest parties cannot send a message to some specific receiver, since we do not wish to assume that parties have any prior knowledge of each others’ identities. The fact that an honest party’s message is *diffused* throughout the entire network is reminiscent of flooding protocols as well as how Bitcoin works.

We consider an adversary who corrupts some parties and can cause them to behave arbitrarily. The adversary can also inject messages into the network, meaning (in particular) that an honest party may receive more than $n - 1$ messages in any given round. In contrast to honest parties, we allow the adversary to send a message to any desired subset of the honest parties. We assume a *rushing* adversary who receives the messages from the honest parties in the current round before deciding on its own messages for that round. The only limitation we place on the adversary is that it may not drop or modify honest parties’ messages. Thus, to recap, if an honest party sends a message in some round, then each honest party receives that message (along with whatever other messages the adversary chooses to send to that party) at the end of that round.

The sequential puzzle functionality, \mathcal{F}_{puz} . We wish to model the existence of computational puzzles that are cheap to verify but expensive (i.e., time-consuming) to solve and, in particular, can only be solved at some bounded rate. Each puzzle is also tied to a specific value which is presented along with a solution in order to enable verification. Puzzle solutions are assumed to be uniform random strings for simplicity, though we only need them to be unpredictable.

We model the above requirements by working in a hybrid world where there is a reactive functionality \mathcal{F}_{puz} which all parties can access twice per round. Let λ denote a (statistical) security parameter. The functionality maintains a set T (initially empty) of puzzle/solution pairs (x, h) , with $x \in \{0, 1\}^*$ and $h \in \{0, 1\}^\lambda$; if $(x, h) \in T$ at some moment in time then we say that h is a *solution for x* . The functionality in each round behaves as follows:

1. Receive from each party P_i an input (`solve`, $x^{(i)}$). For $i = 1, \dots, n$, check if a pair $(x^{(i)}, h^{(i)})$ has been stored in T , and if so return $h^{(i)}$ to P_i ; otherwise, choose uniform $h^{(i)} \in \{0, 1\}^\lambda$, return $h^{(i)}$ to P_i , and store $(x^{(i)}, h^{(i)})$ in T .
2. Receive from each party P_i an arbitrary-length vector (`check`, $(x_1^{(i)}, h_1^{(i)}), \dots$). Return to each party P_i the vector of values $(b_1^{(i)}, \dots)$ where $b_j^{(i)} = 1$ iff $(x_j^{(i)}, h_j^{(i)}) \in T$.

Parties may send messages to each other in between these two calls to the functionality. As in [8], we assume synchrony in the hybrid world. This means that if a corrupted party does not provide input to \mathcal{F}_{puz} by the end of the next “clock tick,” the functionality still returns output to those parties who did provide input.

We stress that this functionality requires all parties to submit their `solve` requests before receiving any of the puzzle solutions. Thus, while an adversary who corrupts f parties can solve f puzzles per round, the puzzle values cannot depend on each other.

Full details of this execution model are deferred to the Appendix.

2.1 Discussion

Instantiation of sequential puzzles. As described above, an essential characteristic of the puzzles modeled by \mathcal{F}_{puz} is that they cannot be solved any *faster* by the adversary than by a single honest process, even though the adversary can corrupt many parties and therefore solve *more* puzzle instances in total. This captures puzzles that require an inherently *sequential* computation to solve. Most cryptographic puzzles (e.g., [4] and [9]) do not satisfy this property. If the cost of verification did not matter, then we could obtain a trivial sequential puzzle through iterated hashing as $\text{sol} := \mathcal{H}^{[d]}(\text{puz})$, which takes exactly d sequential steps to solve. Otherwise, if a trusted party can be relied on to generate puzzle instances, then [30] suffices; but this is undesirable in our setting. Recently, Mahmoody et al. [24] construct puzzles in the bounded random-oracle model that come closest to realizing \mathcal{F}_{puz} ; however, there still remains an approximation gap: in fact, an adversary may gain some constant factor of parallel speed-up with better-than-negligible probability (something not possible with \mathcal{F}_{puz}), and verifying a solution requires nonzero work (in contrast to \mathcal{F}_{puz} , where verification is free). We leave for future work the problem of expanding our model to account for this gap; here we choose to introduce \mathcal{F}_{puz} rather than working directly in a random-oracle model in order to abstract the properties we need, and thus simplify the protocol description (and analysis).

Other notions of puzzles. The sequential puzzles modeled by \mathcal{F}_{puz} are one form of proof-of-work puzzle; however, several other notions have been discussed in the literature. Here we discuss a two alternatives and how they relate to our model.

In Appendix D, we present an ideal functionality intended to model parallelizable proof-of-work puzzles. Here, a puzzle requires some number N of computational steps to solve; however, these steps can be performed in parallel, implying that an adversary controlling f processors can solve puzzles at a f -times faster rate. (Alternatively, this can be viewed as an adversary whose processors run faster than honest processors by a factor of f .) A good example of puzzles that match this abstraction are those developed by Coelho [9].

The cryptographic puzzles used in Bitcoin are not only parallelizable, but can in fact be solved by any number of concurrent processes without communication. The probability with which a given process produces a puzzle solution is proportional to its computational power; hence, regardless of

the distribution of computing resources, puzzle solutions are found according to a Poisson process. Miller et al. [26] point out that this property is essential to the operation of Bitcoin, since it guarantees that independent participants do not duplicate much work; in [25], it is argued that this process is integral to Bitcoin’s incentive structure, since it ensures that even weak participants have a proportional chance of finding the next puzzle solution and thereby earning a reward.

Further remarks on our simplifying assumptions. We have aimed for the simplest (reasonable) model of cryptographic puzzles that captures the features of interest. Yet our model (and results) can be easily adapted or generalized. The assumption that the communication rate is equal to the puzzle-solving rate (i.e., that exactly one puzzle can be solved by each party in each communication round) is without much loss of generality, since the puzzle difficulty can be adjusted so this is the case. Alternately, as long as the number of puzzles that can be solved per round is bounded, our protocols can be easily modified so that they continue to provide the claimed guarantees.

Seemingly more worrisome is the assumption that each corrupted party has the same computational abilities as the honest parties, and so solves puzzles at the same rate. This, too, is not an essential feature, if we simply view n as the total available “units of computational power” in the network, rather than as the total number of parties. (In deed, in our model there is little distinction between a party and a processor.) In that case f would represent not the total number of corrupted entities, but the portion of computational power controlled by the attacker.

As discussed earlier, make puzzle verification “free” in our model to prevent denial-of-service attacks in which the adversary overwhelms an honest party with (incorrect) puzzle solutions. An alternate way to deal with such attacks would be to change our network model and assume some fixed bound on the number of messages the adversary can send to honest parties.

3 Interactive Set Consistency in the \mathcal{F}_{puz} -Hybrid Model

We first define Interactive Set Consistency (ISC).

Definition 1. *A protocol for n parties, in which each party P_i begins with an input value v_i , realizes Interactive Set Consistency in the presence of f corrupted parties if the following holds with all but negligible probability (in an associated security parameter) in the presence of any adversary controlling up to f parties:*

Boundedness *Each honest party P_i outputs a (multi)set V_i containing at most n values.*

Agreement *Each honest party P_i outputs the same (multi)set V .*

Validity *For each honest party P_i , it holds that $v_i \in V_i$.*

ISC is related to Interactive Consistency (IC) [29], with the difference being that the latter has a stronger validity requirement: all honest parties agree on a *vector* \vec{V} (rather than a multiset), with $\vec{V}[i] = v_i$ for each honest party P_i . ISC can be viewed as a pseudonymous version of IC.

We now describe a protocol for realizing ISC in the \mathcal{F}_{puz} -hybrid model. First we introduce some useful terminology.

Definition 2. *A signed message is a tuple (pk, σ, msg) where σ is a valid signature on msg with respect to public key pk .*

If $s = (pk, \sigma, msg)$ is a signed message, then we define $msg(s) = msg$ and $pk(s) = pk$.

Definition 3. A **puzzle graph** is a tuple $(\text{sol}, \text{pk}, \text{children})$, where $\text{sol} \in \{0, 1\}^\lambda$ is a puzzle solution, $\text{pk} \in \{0, 1\}^*$ is an identity string, and children is a (possibly empty) set of puzzle graphs. A puzzle graph $g = (\text{sol}, \text{pk}, \text{children})$ is recursively defined to be **valid** if sol is a solution for the puzzle $\text{pk} \parallel \{\text{sol}(c) \mid c \in \text{children}\}$ (where the latter are ordered lexicographically), and either $\text{children} = \emptyset$ or else every graph in children is valid.

If $h = (\text{sol}, \text{pk}, \text{children})$ is a puzzle graph, then we define $\text{pk}(h) = \text{pk}$.

Definition 4. We define that **puzzle graph h is at depth $\ell > 0$ in puzzle graph g** according to the following inductive rules:

- A puzzle graph g is at depth 1 in itself.
- A puzzle graph h is at depth $(\ell + 1)$ in g if for some $c \in \text{children}(g)$, h is at depth ℓ in c .

Definition 5. A puzzle graph g is **height- ℓ** iff ℓ is the greatest integer such that g contains a depth- ℓ subgraph.

The following lemma will be key to our protocol, as it guarantees that a puzzle graph with sufficient height must contain a subgraph generated during the first round:

Lemma 1. *If a party receives a puzzle graph g in round r , then the height of that graph is at most r . Furthermore, any depth- r subgraph in g must contain a solution output by the \mathcal{F}_{puz} functionality in round 1.*

Proof. A puzzle graph contains a solution sol that must have been output by the functionality \mathcal{F}_{puz} in some round; if the graph has children, then each of those children must contain a puzzle solution output by \mathcal{F}_{puz} in an earlier round. From this observation, the lemma follows by induction. A height-1 graph received in round 1 must contain a single node with no children, computed in that round. A height- $(r + 1)$ graph received in round $r + 1$ must contain a child obtained during round r at the latest. \square

Intuition behind our protocol. Our protocol is inspired by the Dolev-Strong protocol for broadcast (which works by assuming a pre-established PKI) but we integrate computational resources in a nontrivial manner. In each round r , a party *accepts* a value if it has received a collection of r signatures on that value; the party then adds its own signature to the collection and relays it to all other parties. However, since in our setting we do not have a PKI and parties do not know each others' public keys, we must add an additional constraint (using the \mathcal{F}_{puz} functionality) that prevents the adversary from utilizing more than one public key per corrupted party. Our constraint is based on the observation in Lemma 1, that a depth- r subgraph of a puzzle graph received in round r must have been solved in round 1. Thus in our protocol, a correct party only considers a public key “valid” if it comes along with a puzzle graph containing that public key at sufficient depth.

Our protocol is defined in Figure 1. Note that the protocol only uses $f + 1$ communication rounds, though we include a final “computation round” for convenience in the description. We now prove that the protocol realizes ISC.

Lemma 2. *For every honest party P_i , it holds that $\text{pk}_i \in \text{accepted}_j$ for every honest P_j .*

Proof. It is immediate that $\text{pk}_i \in \text{accepted}_i$. For $j \neq i$, it follows from the protocol description that P_j adds pk_i to accepted_j in round 2. \square

An ISC Protocol the \mathcal{F}_{puz} -Hybrid Model

Initially, each party P_i generates a keypair (sk_i, pk'_i) for a digital signature scheme. We set $pk_i = pk'_i \parallel v_i$, where v_i is the input of P_i , and refer to pk_i as the *identity* of party P_i .

The subroutine $solve(pk, children)$ returns the puzzle graph $g = (sol, pk, children)$, where sol is the solution returned from querying the \mathcal{F}_{puz} functionality with $(solve, pk \parallel \{sol(c) \mid c \in children\})$ (where the latter are in lexicographic order).

We define $Sign_i(msg)$ to return (pk_i, σ, msg) , where σ is a signature on msg computed using sk_i .

- **Round 1.** In the first round, each party P_i computes the puzzle graph $g_{i,1} = solve(pk_i, \emptyset)$. It then sends $g_{i,1}$ and $Sign_i(pk_i)$ to all parties, and sets $accepted_i = \{pk_i\}$.
- **Round 2 through $f + 2$.** For $r = 2, \dots, f + 2$, party P_i does:
 1. Set $G_{new} = \emptyset, S_{new} = \emptyset$.
 2. Let G be the set of valid puzzle graphs received in the previous round.
 3. Let S be the set of valid signed messages (pk, σ, msg) received in the previous round for which $pk \in accepted_i$. (S includes only one signed message per (pk, msg) pair; duplicates are discarded.)
 4. For each $g \in G$, and for each puzzle graph h that is depth $r - 1$ in g , let $S_h = \{s \in S \mid msg(s) = pk(h)\}$. If $pk(h) \notin accepted_i$ and $|S_h| \geq r - 1$, then:
 - (a) add $pk(h)$ to $accepted_i$,
 - (b) add g to G_{new} ,
 - (c) set $S_{new} = S_{new} \cup S_h \cup \{Sign_i(pk(h))\}$.
 5. If $r \leq f + 1$, send $g_{i,r} = solve(pk_i, G_{new})$ and S_{new} to all parties.

Output the set $V_i = \{v_j \mid pk'_j \parallel v_j \in accepted_i\}$.

Figure 1: Our ISC protocol for the \mathcal{F}_{puz} -hybrid model.

We say that an honest party P_i *accepts pk in round r* if P_i adds pk to $accepted_i$ at the end of round r . We say it *accepts pk by round r* if it accepts pk in round $r' \leq r$.

Lemma 3. *If an honest party accepts pk in round $r \leq f + 1$, then every honest party accepts pk by round $r + 1$.*

Proof. The proof is by induction on $r \geq 1$. For the base case, when $r = 1$, observe that each honest party P_i only accepts its own key pk_i and, as noted in the proof of Lemma 2, every other honest party accepts pk_i in round 2.

Suppose the lemma holds at round $r - 1 \geq 1$, and say honest party P_i accepts pk in round r . We know that P_i must have received (in round $r - 1$) at least $r - 1$ signed messages s' with $msg(s') = pk$ and $pk(s')$ accepted by P_i by round $r - 1$. These $r - 1$ signed messages must be signed using distinct public keys; furthermore, none of those public keys can be equal to pk_i because P_i only signs pk after it accepts it. Applying the inductive hypothesis, every other honest party must have accepted those $r - 1$ public keys by round r ; from the proof of Lemma 2, we know that every honest party has also accepted pk_i by round r . Since P_i sends all the signed messages it has received, plus its own signature on pk , in round r , it follows that in round r every other honest party P_j receives at least r signed messages s with $msg(s) = pk$ and $pk(s) \in accepted_j$.

Next, observe that if P_i accepts pk in round r , then it must have received a puzzle graph g containing a depth- $(r - 1)$ subgraph h such that $pk(h) = pk$. By solving an additional puzzle, P_i obtains a puzzle graph $g_{i,r}$ in which h is at depth r and sends $g_{i,r}$ to all other parties in round r .

It follows from both the above that every other honest party will accept \mathbf{pk} by round $r + 1$. \square

Lemma 4. *Each honest party accepts at most n distinct keys.*

Proof. If an honest party accepts \mathbf{pk} in some round $r \geq 2$, then it must have received a height- $(r - 1)$ puzzle graph g in the previous round. By Lemma 1, any height- $(r - 1)$ ground received in round $r - 1$ must contain a puzzle solved during round 1; since at most n such puzzles in total can be computed in a round, at most n distinct keys can be accepted. \square

Theorem 1. *For any $f < n$ there is a polynomial-time ISC protocol with $f + 1$ rounds of communication, secure against f corrupted parties.*

Proof. Validity follows from Lemma 2, and boundedness follows from Lemma 4. To prove agreement, we must show that if an honest party P_i accepts \mathbf{pk} in round $f + 2$, then every honest party accepts \mathbf{pk} by round $f + 2$. So, say P_i accepts \mathbf{pk} in round $f + 2$. Then P_i must have received $f + 1$ signatures on \mathbf{pk} from previously-accepted keys (not including its own). Observe that among these $f + 1$ previously-accepted keys, at least 1 must belong to some honest party P_j : by Lemma 2, the keys of all $n - f$ correct parties are accepted by P_i in round 1, and by Lemma 4 at most n keys in total are accepted among all the correct parties. Since a correct party only signs \mathbf{pk} after accepting it, P_j must have accepted \mathbf{pk} in round f or earlier. Therefore, by Lemma 3, if P_j accepts \mathbf{pk} in round f , then every correct party accepts \mathbf{pk} in round $f + 1$ or earlier. \square

Message Complexity. In the worst case, each party publishes n^2 signed messages, n for each accepted key. Since no more than $n(f + 1)$ puzzle solutions are solved in total, any puzzle graph can be represented using only $O(\lambda n f)$ bits (i.e., the graphs should be represented in a way that avoids duplicating shared children). Therefore since each party accepts at most n keys, it publishes at most $O(\lambda n^2 f)$ message bits. As a further optimization, each party may keep track of which subgraphs it has already published (e.g., after accepting a key in an earlier round) and avoid publishing duplicates. This reduces the worst case message cost to $O(\lambda n f)$ bits per party.

3.1 A Lower Bound on the Round Complexity

Our ISC protocol is round-optimal for *deterministic*¹ protocols:

Theorem 2. *Say $n - f \geq 2$. (Otherwise, ISC is trivial.) Any deterministic ISC protocol in the \mathcal{F}_{puz} -hybrid model that tolerates f faults must have at least $f + 1$ rounds of communication.*

A proof of the above is given in Appendix C. The proof is inspired by the lower bound on the round complexity of broadcast with a PKI [12, 1], but requires several modifications.

¹Technically, our protocol is not deterministic because of the initial key-generation step. Our proof can be extended to apply to deterministic protocols given access to an ideal functionality corresponding to a digital signature scheme.

4 Pseudonymous MPC from Sequential Puzzles

In the previous section, we showed a protocol for ISC in the \mathcal{F}_{puz} -hybrid model. This protocol allows the participating parties to agree on a set of public keys, with honest parties' public keys guaranteed to be included; thus, this effectively bootstraps a (pseudonymous) PKI.

In this section, we show that ISC can be used to achieve (pseudonymous) secure computation of general functionalities. We first show that the standard Dolev-Strong algorithm [12] can be run following ISC in order to implement a pseudonymous version of *secure broadcast*; this, in turn, can be used as part of standard protocols for secure multi-party computation (e.g., the GMW protocol [18]) to implement pseudonymous secure computation. In general, any other multi-party computation protocol (such as BGW [6], for example) could be used in place of BGW; the fairness and resilience properties would carry over from such protocol, since our ISC protocol provides the best possible resilience ($f < n$) and fairness condition (guaranteed termination).

4.1 Definition of Pseudonymous Security

We motivate the need to define a relaxed, *pseudonymous* notion of security by pointing out that the standard notion of security [8] cannot be achieved in a setting where the underlying network does not offer authenticated communication. Intuitively, the issue is that there is nothing distinguishing the parties from each other; thus, it is not even clear which of the n parties should play the role of P_1 , which the role of P_2 , etc. As a consequence, we must either restrict attention to functions that treat each of their n inputs *symmetrically*, or else we must allow the attacker to arbitrarily choose the order in which the parties' n inputs are entered into the computation. We choose the second approach, which is in fact more general than the first.

The relaxed notion of pseudonymous secure computation that we consider may still be useful in many circumstances. First, many functions of interest (such as majority, average, etc.) are invariant under permutations of the input values, and are therefore unaffected by the relaxation. Also, in many scenarios the ordering of the inputs may be irrelevant; for example, in a game of online poker against anonymous opponents, each party might only be concerned with having the game played honestly and correctly.

The formal definition. To define pseudonymous secure computation, we define an appropriate ideal world which is a relaxation of the one normally considered when defining secure computation. (We have already defined our real world in Section 2, and as usual [8] we will define a real-world protocol to be secure if the actions of any adversary attacking the protocol can be simulated by an adversary in the ideal world.) The main relaxation is that in pseudonymous MPC *we allow the adversary to choose an arbitrary permutation on the inputs*.

Let $f : (\{0, 1\}^*)^n \rightarrow (\{0, 1\}^*)^n$ be a function to evaluate. Each party P_i holds input $x_i \in \{0, 1\}^\lambda$, where λ is the security parameter. The ideal-world adversary \mathcal{S} is an interactive Turing machine that controls a set $C \subset \{1, \dots, n\}$ of corrupted parties, and we denote the set of uncorrupted parties by $I = [n] \setminus C$. The adversary \mathcal{S} receives the inputs of the corrupted parties, as well as an auxiliary input z . Execution in the ideal world proceeds as follows:

- **Input substitution:** Each honest party $P_i \in I$ sends its input x_i to the trusted party evaluating f . The adversary \mathcal{S} sends arbitrary inputs x_i on behalf of each corrupted party $P_i \in C$. Let $\vec{x} = (x_1, \dots, x_n)$ denote the input values sent by the n parties.²

²We assume parties have identifiers in the ideal world, but these identifiers are only used as a formalism to allow

- **Permutation:** \mathcal{S} chooses a permutation π on $\{1, \dots, n\}$.
- **Computation:** Let \vec{x}' be the vector obtained by applying the permutation π to the entries of \vec{x} ; i.e., $x'_{\pi(i)} = x_i$. The trusted party computes $(y_1, \dots, y_n) = f(x'_1, \dots, x'_n)$, and returns $y_{\pi(i)}$ to party P_i .
- **Output:** Each honest P_i outputs the value given to it by the trusted party. In addition, \mathcal{S} outputs an arbitrary function of its view.

4.2 Protocol for Pseudonymous Multi-party Computation

The GMW protocol for multi-party computation makes use of two communication primitives: private authenticated channels, and broadcast. We will need to emulate (the pseudonymous version of) these in the \mathcal{F}_{puz} -hybrid model. The basic idea is for each party to generate a public encryption/signing keypair, and then run the ISC protocol so that all the parties agree on a set of public keys. This effectively amounts to bootstrapping a pseudonymous PKI that includes a pseudonym for each honest party and at most f pseudonyms for the adversary. Thereafter, standard encryption and signing techniques can be used to realize (pseudonymous) authenticated and private channels between the parties; the Dolev-Strong algorithm can be used to give pseudonymous broadcast.

In more detail, the protocol for evaluating a function f proceeds as follows:

1. First, each party P_i generates a signing and encryption keypair $\mathbf{pk}_i, \mathbf{sk}_i$. Next, each party executes the ISC protocol 1, using \mathbf{pk}_i as its input value. At the end of the protocol, each party obtains a set of public keys which it can order lexicographically to give a vector $(\mathbf{pk}'_1, \dots, \mathbf{pk}'_n)$. We remark that this effectively defines a permutation of the parties, and that each party can determine its permuted index.
2. Next, each party P_i executes the GMW protocol [18]. Note that the GMW protocol uses both broadcasts and private messages. When the protocol calls for a broadcast, the Dolev-Strong algorithm is executed using $\vec{\mathbf{pk}}'$ as the PKI. When the protocol calls for a private message, the signing and encryption keys are used as described above.

In Appendix B we prove the following theorem:

Theorem 3 (Pseudonymously Secure SFE). *In the \mathcal{F}_{puz} -hybrid model, we achieve SFE for general functions while ensuring pseudonymous security against a non-adaptive adversary. Specifically, input completeness and guaranteed termination can be achieved with honest majority, i.e., $f < n/2$. In the presence of an arbitrary number of corruptions, we ensure security with unanimous abort.*

5 Conclusion

This work is inspired by peer-to-peer networks like Bitcoin, and makes an initial attempt to understand what meaningful security properties can be achieved in such distributed networks. Contrary to the widely held belief that nothing interesting can be achieved in networks without authenticated channels or a PKI, we show that by placing a strict bound on each party's computational resources it is possible set up a (pseudonymous) PKI and achieve (pseudonymous) notions of broadcast and secure computation. Although our work does not directly apply to the actual Bitcoin protocol, our results show that distributed-computation models resembling Bitcoin can lead to rich applications with non-trivial security guarantees.

\mathcal{S} to address the corrupted parties; the identifiers have no inherent meaning.

References

- [1] Marcos Kawazoe Aguilera and Sam Toueg. A simple bivalency proof that t -resilient consensus requires $t + 1$ rounds. *Information Processing Letters*, 71(3):155–158, 1999.
- [2] Marcin Andrychowicz and Stefan Dziembowski. Distributed cryptography based on proofs of work. Cryptology ePrint Archive, Report 2014/796, 2014. <http://eprint.iacr.org/>.
- [3] James Aspnes, Collin Jackson, and Arvind Krishnamurthy. Exposing computationally-challenged Byzantine impostors. Technical report, Computer Science Dept., Yale University, 2005.
- [4] Adam Back. Hashcash—a denial of service counter-measure. <http://www.hashcash.org/papers/hashcash.pdf>, 2002.
- [5] Boaz Barak, Ran Canetti, Yehuda Lindell, Rafael Pass, and Tal Rabin. Secure computation without authentication. In *Advances in Cryptology—CRYPTO 2005*, pages 361–377. Springer, 2005.
- [6] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *20th Annual ACM Symposium on Theory of computing*, pages 1–10. ACM, 1988.
- [7] Nikita Borisov. Computational puzzles as sybil defenses. In *Sixth IEEE International Conference on Peer-to-Peer Computing*, pages 171–176. IEEE, 2006.
- [8] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, 2000.
- [9] Fabien Coelho. An (almost) constant-effort solution-verification proof-of-work protocol based on merkle trees. In *Progress in Cryptology—AFRICACRYPT 2008*, pages 80–93. Springer, 2008.
- [10] Jeffrey Considine, Matthias Fitzi, Matthew Franklin, Leonid A Levin, Ueli Maurer, and David Metcalf. Byzantine agreement given partial broadcast. *Journal of Cryptology*, 18(3):191–217, 2005.
- [11] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, Anne-Marie Kermarrec, Eric Ruppert, et al. Byzantine agreement with homonyms. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 21–30. ACM, 2011.
- [12] Danny Dolev and H. Raymond Strong. Authenticated algorithms for Byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983.
- [13] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Advances in Cryptology – CRYPTO’92*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer Verlag, 1993.
- [14] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*, 2014.

- [15] Matthias Fitzi, Nicolas Gisin, Ueli M. Maurer, and Oliver von Rotz. Unconditional Byzantine agreement and multi-party computation secure against dishonest minorities from scratch. In *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 482–501. Springer Verlag, 2002.
- [16] Matthias Fitzi, Daniel Gottesman, Martin Hirt, Thomas Holenstein, and Adam Smith. Detectable Byzantine agreement secure against faulty majorities. In *21st ACM Symposium Annual on Principles of Distributed Computing*, pages 118–126. ACM Press, 2002.
- [17] Juan A Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. Technical report, 2014.
- [18] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game, or a completeness theorem for protocols with honest majority. In *19th Annual ACM Symposium on Theory of Computing*, pages 218–229. ACM Press, 1987.
- [19] Shafi Goldwasser and Yehuda Lindell. Secure multi-party computation without agreement. *Journal of Cryptology*, 18(3):247–287, 2005.
- [20] Bogdan Groza and Bogdan Warinschi. Cryptographic puzzles and DoS resilience, revisited. *Designs, Codes and Cryptography*, 2013.
- [21] Ari Juels and John G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*, volume 99, pages 151–165, 1999.
- [22] Joshua A Kroll, Ian C Davey, and Edward W Felten. The economics of bitcoin mining or, bitcoin in the presence of adversaries. In *Workshop on Economics in Information Security (WEIS)*, 2013.
- [23] Ben Laurie and Richard Clayton. Proof-of-work proves not to work. In *Workshop on Economics and Information Security*, 2004.
- [24] Mohammad Mahmoody, Tal Moran, and Salil Vadhan. Publicly verifiable proofs of sequential work. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 373–388. ACM, 2013.
- [25] A Miller, A Juels, E Shi, B Parno, and J Katz. Permacoin: Repurposing bitcoin work for long-term data preservation. *IEEE Security and Privacy*, 2014.
- [26] Andrew Miller and Joseph J LaViola Jr. Anonymous byzantine consensus from moderately-hard puzzles: A model for bitcoin, 2014.
- [27] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://bitcoin.org/bitcoin.pdf>, 2008.
- [28] Michael Okun. *Distributed Computing Among Unacquainted Processors in the Presence of Byzantine Failures*. PhD thesis, Hebrew University of Jerusalem, 2005.
- [29] Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.

- [30] Ronald L Rivest, Adi Shamir, and David A Wagner. Time-lock puzzles and timed-release crypto. Technical report, Massachusetts Institute of Technology, 1996.

A Formal Specification of our Model

A formal specification of our \mathcal{F}_{puz} -hybrid execution model is given by Figures 2 and 3.

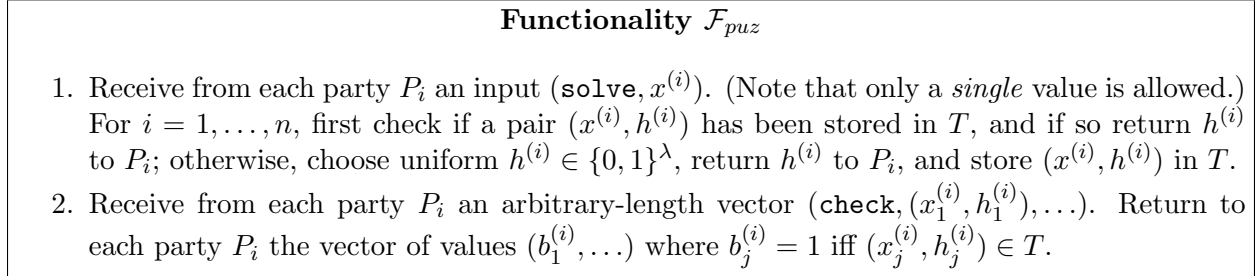


Figure 2: The \mathcal{F}_{puz} functionality

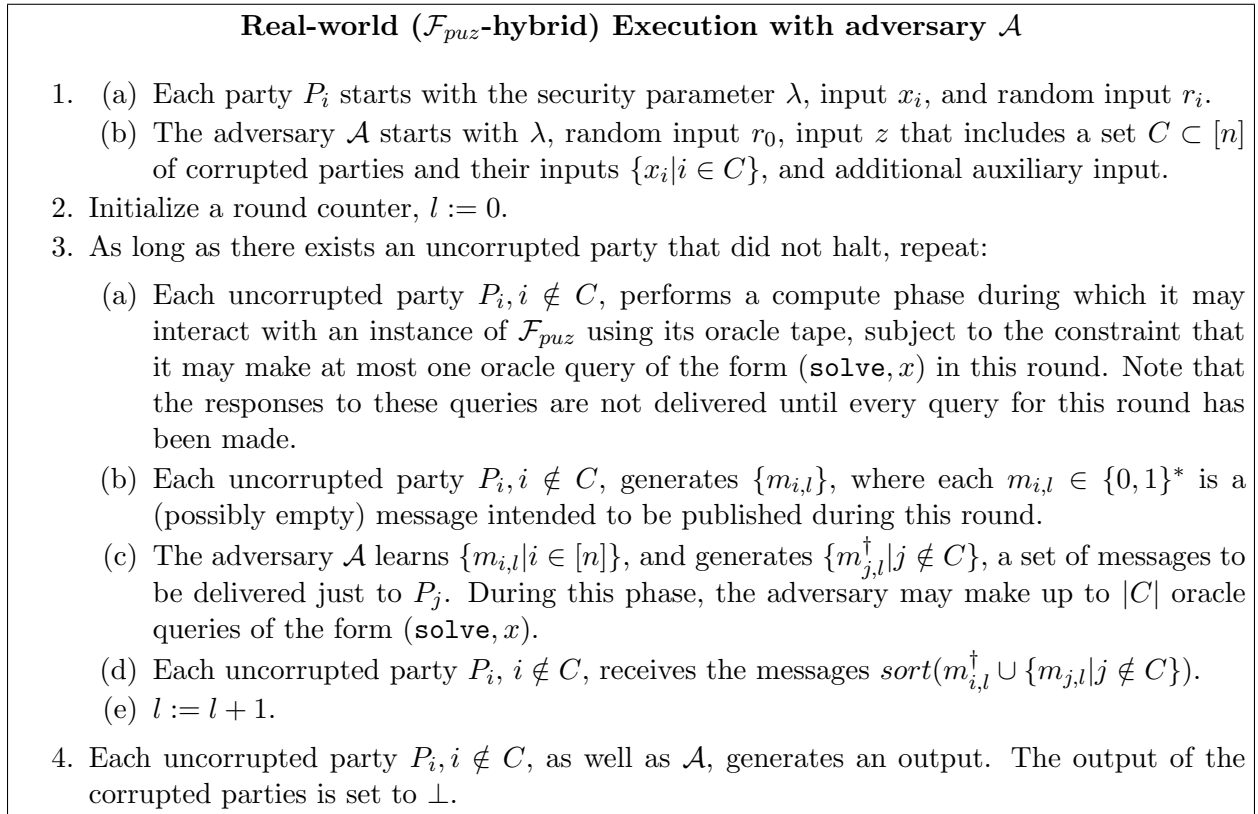


Figure 3: A summary of the nonadaptive \mathcal{F}_{puz} -hybrid computation.

B Proof of our Pseudonymously Secure Multi-party Computation Protocol

We restate and prove the following theorem from Section 4.2:

Theorem 3 (Pseudonymously Secure SFE). *In the \mathcal{F}_{puz} -hybrid model, we achieve SFE for general functions while ensuring pseudonymous security against a non-adaptive adversary. Specifically, input completeness and guaranteed termination can be achieved with honest majority, i.e., $f < n/2$. In the presence of an arbitrary number of corruptions, we ensure security with unanimous abort.*

Proof. First we prove that the protocol described above satisfies the case for honest majority. Since the underlying GMW protocol is executed intact, the existing simulator for this protocol, \mathcal{S}' , can be reused, with the permuted player indices substituted for the correct indices. The main thing for us to show is that the real-world execution of the ISC protocol in the first phase can indeed be simulated.

Our simulator \mathcal{S} first generates keypairs pk_i and sk_i for each simulated party. The view of the real-world adversary consists of the messages sent by the honest parties in each round; this communication pattern can be simulated perfectly. According to the *agreement* property of the ISC definition, if a corrupted party does not complete the protocol, its public key may be replaced with \perp ; in any case, all uncorrupted parties arrive at a consistent vector pk with high probability. Since the ISC protocol sorts this vector in lexicographical order, the simulator \mathcal{S} instructs the trusted party T to use this permutation.

After completing the ISC phase, our simulator \mathcal{S} executes the GMW simulator \mathcal{S}' in a sandbox, placing each party P_i with the corresponding permuted index. Note that since \mathcal{S}' is intended for a model with broadcast and private channels, we must be able to simulate the protocols used to implement these channels. It is trivial to simulate the broadcast protocol since the simulator already knows the signing keys; the private channel can be trivially simulated by encrypting random strings.

Finally in the case without honest majority (i.e., $n > f > \lfloor n/2 \rfloor$) the GMW [18] protocol guarantees privacy but not fairness. \square

C Lower Bound on Round Complexity in the \mathcal{F}_{puz} Model

Dolev and Strong first showed that assuming authentication, any deterministic Byzantine agreement protocol tolerating f faults (where $n > f + 1$) must have at least $f + 1$ messaging rounds [12].

Our lower bound proof is inspired a simple bivalency-based proof of this theorem [1]. We basically must show that having the additional puzzle functionality does not change this lower bound, even though this additional functionality can be used to further restrict behavior of the adversary, since the adversary can only solve a limited number of computational puzzles per round.

For the purpose of our lower bound, we limit our concern only to *compliant* protocols (as defined below) that are effectively deterministic. We must carefully explain what this means in our setting, since the \mathcal{F}_{puz} functionality itself is randomized and our positive results make use of digital signature algorithms.

Definition 6. *A compliant protocol in the \mathcal{F}_{puz} -hybrid model must ignore the security parameter and treat the public/private keys, digital signatures, and puzzle solutions as opaque strings. Protocols*

may transmit opaque strings along communication channels and include them as input to oracle queries and digital signature algorithms, but cannot otherwise inspect or modify their bits.

This definition in particular rules out protocols that use the signature algorithm or interaction with \mathcal{F}_{puz} in order to perform coin flips.

Theorem 2. *Say $n - f \geq 2$. (Otherwise, ISC is trivial.) Any deterministic ISC protocol in the \mathcal{F}_{puz} -hybrid model that tolerates f faults must have at least $f + 1$ rounds of communication.*

Proof. For the purpose of proving our lower bound, we limit ourselves to constructing adversaries that are also compliant in the same sense, and in fact only need to *crash* processes (rather than induce Byzantine behavior). Furthermore, our adversaries crash at most one process per round.

We will state a lemma that given these constraints, the same reasoning used in deterministic models applies equally to compliant protocols in our model with high probability.³ First we define some notation useful for describing executions in our model.

A *partial run* is a finite sequence of system configurations (including the input values to each process and the internal state of each process and the adversary) that transition according to our model. The last element in an r -round *partial run* is the system configuration at the end of r rounds. Note that since our execution model is randomized (due to coin flips taken by the puzzle functionality and digital signature scheme), given an initial configuration our model defines a probability distribution over r -round partial runs.

We define an equivalence relation among partial runs in our model. Two partial runs are equivalent if they are equal up to a bijective transformation of keys, signatures, and puzzle solutions. A partial run s can be represented as a pair $(state, \{o_i\})$, where $\{o_i\}$ is the sequence of every occurrence of an opaque string in a configuration and $state$ is the sequence of system configurations with each occurrence of an opaque string taken out (and replaced with a symbol `opaque`). Another partial run s' is equivalent to s if $s' = (state, \{F(o_i)\})$ for some bijective function $F : \{0, 1\}^* \rightarrow \{0, 1\}^*$. This implies that two partial runs are equivalent if their initial conditions are the same, any decisions made by correct processes are the same, any processes that have crashed did so at the same time, and the content of the internal states of each process and messages delivered differ only in the representation of the opaque strings associated with signatures and puzzles.

Lemma 5. *Consider a compliant protocol and compliant adversary (as per Definition 6). Then for each round r , there exists a single equivalence class of r -round partial runs (which we call the canonical partial run) such that with high probability, a partial run sampled from the distribution of r -round partial runs belongs to this equivalence class.*

Proof. As the protocols make only a (polynomially) bounded number of queries (to \mathcal{F}_{puz} or to digital signature routines), with high probability the processes do not generate colliding keys or puzzle solution strings, nor does any process “guess” a puzzle solution it has not received from \mathcal{F}_{puz} or the communication channel. \square

³Note that an alternative to invoking the above lemma would be to redefine digital signatures and our puzzle functionality to inherently use opaque “tags” rather than concrete bit-strings - indeed this is the approach taken in [12]. We have chosen our approach because it allows us to use a simple and concrete representation of the model for presenting our positive results.

Hereafter, we are only concerned with the canonical partial runs (rather than partial runs that involve some unlikely string collision), and therefore the term *partial run* should be taken to mean the class of *canonical partial runs*.

Let s_r be an r -round partial run. We say that s_{r+1} is a *one-round extension* of s_r if there exists some adversary \mathcal{A} such that s_r and s_{r+1} are respectively the canonical r -round and $r + 1$ -round partial runs given \mathcal{A} . A partial run s is an *extension* of s_r if it is a partial run obtained by extending s_r one round at a time and in which all processes have decided.

Having established a basis for applying deterministic reasoning to our model, the lower bound proof can be carried out with simple modifications to the bivalency proof in [1]. We describe these necessary modifications and then refer the reader to the text of [1] for the remainder of the proof.

The lower bound of [12, 1] applies to binary consensus, a somewhat different problem than ISC in which processes must decide on a single bit). We can imagine that each process decides on 2^λ bits, one for each possible input value (at most n of these may be set to 1). The proof in [1] relies on a notion of *valences*, which we now define in our setting. A partial run s_r is *v -valent* (for some value v) if in all extensions of s_r , the correct processes decide on a set S where $v \in S$, and *\bar{v} -valent* if in all extensions $v \notin S$. When v is clear from context, we abbreviate v -valent with 1-valent and \bar{v} -valent in order to more closely match the notation in [1]. A partial run is *univalent* (with respect to v) if it is either v -valent or \bar{v} -valent, and *bivalent* if it is neither v -valent nor \bar{v} -valent.

The proof in [1] involves constructing pairs of executions that are indistinguishable to some of the processes in the network. We next prove a lemma that guarantees that these constructions remain indistinguishable, even in our setting where the processes may additionally interact with \mathcal{F}_{puz} .

Lemma 6. *For $r < f$, let s_r be an r -round partial run in which at most r processes have crashed. Consider the following two one-round extensions of s_r : s_{r+1} , in which one process crashes before delivering its message to some process p , and s'_{r+1} in which the same process crashes after delivering its message to p . Let s' and s denote the runs extending s'_{r+1} and s_{r+1} respectively in which p crashes at the beginning of the next round (before sending any messages). (Note that in the special case $r + 1 = f$, the processes terminate so the crash is unnecessary.) Let p' (distinct from p) be a correct process. Process p' cannot distinguish between runs s and s' (i.e., its behavior in both runs is the same).*

Proof. In the case of an ordinary message passing network, this is immediate, since p is the only process that observes any difference in round $r + 1$ and p crashes before it can communicate to any other process. However in our setting we must prove this holds in spite of the additional \mathcal{F}_{puz} functionality, since in round $r + 1$, process p may affect \mathcal{F}_{puz} by solving a puzzle before crashing. Without loss of generality, assume that in run s , p solves a puzzle puz and in s' solves a puzzle puz' . However, note that a) \mathcal{F}_{puz} records a puzzle solution sol randomly sampled from a space of size 2^λ , b) the recorded solution is returned only to process p , which crashes before it can transmit any information about it to another process, and c) the other processes may only observe the difference through interaction with \mathcal{F}_{puz} by calling sending (`check, puz, sol`) or (`check, puz', sol`) (i.e., by correctly *guessing* the puzzle solution). As the processes only make a polynomial number of queries, this occurs with negligible probability. \square

Now that we have defined the notions of (canonical) partial runs and valences in our setting and established conditions under which indistinguishability holds despite the additional \mathcal{F}_{puz} , our theorem follows from the proof in [1] verbatim.

D ISC in the Parallelizable Proof-of-Work Model

While in our main result we have considered inherently sequentially puzzles (i.e., puzzles that take the adversary an entire round to solve, even with all f corrupted processes working together in parallel), we are also interested in modeling puzzles that the adversary can solve faster using its parallel resources. Equivalently, we may be interested in modeling an adversary that can compute sequentially f times faster than a corrupt process.

We can model this via a modified functionality \mathcal{F}_{parpuz} , which differs from \mathcal{F}_{puz} in that it allows the adversary to make multiple rounds of interaction with the functionality within a single communication round.

- Receive from each *uncorrupted* party P_i an input (`solve`, $x^{(i)}$). (Note that only a *single* value is allowed.) For $i = 1, \dots, n$, first check if a pair $(x^{(i)}, h^{(i)})$ has been stored in T , and if so return $h^{(i)}$ to P_i ; otherwise, choose uniform $h^{(i)} \in \{0, 1\}^\lambda$, return $h^{(i)}$ to P_i , and store $(x^{(i)}, h^{(i)})$ in T .
- For up to f iterations, a single corrupted party P_i may request (`solve`, x_i), and the input is processed immediately as above.
- Receive from each party P_i an arbitrary-length vector (`check`, $(x_1^{(i)}, h_1^{(i)}), \dots$). Return to each party P_i the vector of values $(b_1^{(i)}, \dots)$ where $b_j^{(i)} = 1$ iff $(x_j^{(i)}, h_j^{(i)}) \in T$.

As before, we allow parties to call \mathcal{F}_{parpuz} with a `check` instruction any (polynomial) number of times. Each of the honest parties is allowed to call \mathcal{F}_{parpuz} with a `solve` instruction only once per communication round; moreover, all of the `solve` instructions for a round must be sent before any honest party receives its puzzle solution. However, the corrupted parties can call \mathcal{F}_{parpuz} , one after another in sequence, up to a total of f times within an overall communication round.

Protocol intuition for the \mathcal{F}_{parpuz} -hybrid model. Our protocol proceeds in two overall phases. In the first phase, called the “mining” phase, each correct process constructs a chain of $O(f^2)$ puzzle solutions associated with that process’s public key. In the second phase, the “communication” phase, the processes publish their puzzle chains and propagate the puzzle chains they receive from others. To ensure agreement, each process also signs and relays signatures according to the Dolev-Strong algorithm. Signatures corresponding to public keys without associated puzzle chains are ignored. The communication phase ends after $f + 1$ rounds.

Intuitively, the protocol works because every correct process is able to create a valid puzzle chain for its own key, yet the corrupt processes are only able to create at most f puzzle chains before the protocol terminates.

Definition 7. A length- ℓ puzzle chain is defined inductively as follows:

- A puzzle graph g is a length-1 puzzle chain if $\text{children}(g) = \emptyset$.
- A puzzle graph g is a length- $(\ell + 1)$ puzzle chain if $\text{children}(g) = \{s\}$, $\text{pk}(s) = \text{pk}(g)$, and s is a length- ℓ puzzle chain.

The protocol is defined in Figure 4.

Lemma 7. For every correct process P_i , its identity pk_i is accepted by every other correct process in round 2.

An $O(f^2)$ -Round ISC Protocol using Parallelizable Puzzles

Initially, each process P_i generates a signing keypair $(\text{sk}_i, \text{pk}_i)$ according to some digital signature scheme, where pk_i is the concatenation of the underlying public key and the i^{th} input value. Hereafter we refer to pk_i as the identity of process P_i . The subroutines *solve* and *sign_i* are the same as in 1.

- **Mining Phase.** In the first round, each process P_i constructs a puzzle graph

$$g_{i,1} := \text{solve}(\text{pk}_i, \emptyset)$$

where pk_i denotes the sender's identity, and the message includes an empty history. In each of the following rounds, up to round $r_{\text{mine}} = f(f+1) + 1$, each process extends its chain by one puzzle solution

$$g_{i,r} := \text{solve}(\text{pk}_i, \{g_{i,r-1}\}).$$

The mining phase ends after r_{mine} rounds. We abbreviate $g_{i,r_{\text{mine}}}$ by g_i .

- **Communication Phase.** Initially, in round 1 of the communication phase, each process P_i
 1. sets $\text{accepted} := \{\text{pk}_i\}$,
 2. publishes g_i ,
 3. and publishes $\text{sign}_i(\text{pk}_i)$.

Thereafter, in each round $2 \leq r \leq f+2$, process P_i

1. Receives a set of signed messages S , such that $\forall s \in S, \text{pk}(s) \in \text{accepted}$ and s has a distinct pk and m (i.e., $\forall s' \in S \setminus \{s\}, (\text{pk}(s), \text{msg}(s)) \neq (\text{pk}(s'), \text{msg}(s'))$). Invalid or redundant signatures are discarded.
2. Receives a set of length- r_{mine} graphs G such that $\forall g \in G, \text{pk}(g) \notin \text{accepted}$, and each g has a distinct id (i.e., $\forall g' \in G \setminus \{g\}, \text{pk}(g) \neq \text{pk}(g')$). Redundant or invalid puzzle chains are discarded.
3. For each $g \in G$, let $S_g := \{s \in S \mid \text{id}(s) = \text{id}(g)\}$. If $|S_g| \geq r$, then P_i :
 - (a) sets $\text{accepted} := \text{accepted} \cup \{\text{pk}(g)\}$,
 - (b) publishes g ,
 - (c) and publishes $S_g \cup \{\text{sign}_i(\text{pk}(g))\}$.

At the end of round $f+2$ (of the communication phase), process P_i outputs the set accepted .

Figure 4: Our ISC protocol for the $\mathcal{F}_{\text{parpuz}}$ -hybrid model.

Proof. This follows immediately from the protocol definition. In round 1 of the communication phase, each correct process P_i broadcasts its own puzzle chain and signature on pk_i , and every other process receives these messages and accepts pk_i in the next round. \square

Lemma 8. *If a correct process P_i accepts a key pk in round $r < f+2$ (of the communication phase), then every correct process accepts pk in round $r+1$ or earlier.*

Proof. The proof is by induction on the round number $r \geq 1$. For the base case, when $r = 1$ observe that each correct process P_i only accepts its own key pk_i , and by Lemma 7 every other correct process accepts pk_i in round 2.

Suppose the lemma holds at round r , and suppose a correct process P_i accepts pk in round r .

First, note that in round $r+1$, every other correct process receives at least r signatures over pk from distinct previously-accepted keys (see the proof of Lemma 3).

Second, observe that P_i publishes a length- r_{mine} chain g for pk in round r , and therefore every other correct party receives g in round $r+1$. Thus, every correct process accepts pk in round $r+1$,

completing the proof by induction. \square

Lemma 9. *Among the correct processes, at most n distinct keys are ever accepted.*

Proof. Each accepted key requires a length- r_{mine} chain, and each node in the chain must be associated with a consistent key. Each uncorrupted process only solves puzzles associated with its own key. Since the corrupted processes can solve (in total) at most f puzzles per round, they can solve at most $f(r_{mine} + f + 1)$ puzzles before the protocol terminates. Therefore, at most $\lfloor f(r_{mine} + f + 1)/r_{mine} \rfloor = f + \lfloor f \cdot (f + 1)/r_{mine} \rfloor = f$ length- r_{mine} chains can be computed by corrupt processes. Therefore, length- r_{mine} chains are found corresponding to at most n distinct keys. \square

Theorem 4. *There exists a polynomial-time ISC protocol with $O(f^2)$ rounds of communication, secure against $f < n$ number of corrupted parties.*

Proof. The proof of this theorem is identical to that of Theorem 1, substituting Lemmas 7,8, and 9 for Lemmas 2,3, and 4. \square

E Future Work

There are many directions for future research. In our work we have made several simplifying assumptions, and relaxing any of them is an important next step. For example, we have assumed that there is a fixed number n of parties that is public knowledge at the outset of the protocol. Is anything achievable if n is unknown? Alternately, what can be said in a *partially synchronous* network where the maximum communication delay is unknown? In a different direction, it would be interesting to explore a model in which parties were not *malicious* or *honest*, but are instead only *rational*. Here, we could assign some “cost” to solving puzzles rather than assuming a strict upper bound on how many puzzles a party could solve per round. This is the direction taken in other work analyzing Bitcoin [14, 22]. It would be useful to reconcile this with approaches to modeling rationality in secure multi-party computation.