Efficient Agreement Over Byzantine Gossip

Ran Cohen [*]   Julian Loss [†]   Tal Moran [‡]

January 12, 2024

Abstract

Byzantine agreement (BA) asks for a set of parties to reach agreement in an adversarial setting. A central question is how to construct *efficient* BA protocols that scale well with the number of parties. In particular, the *communication complexity* is a critical barrier for large-scale implementations.

State-of-the-art, scalable BA protocols typically work by sampling a small, unpredictable committee of parties that will send messages in each round. These messages must reach all honest parties, to allow the next round's committee to function. In practice, this is usually accomplished by propagating messages over a *gossip network*, implemented over a partial communication graph. Most formulations of gossip networks have an ideal guarantee that *every* message delivered to any honest party will be delivered to *every* other honest party. Unfortunately, realizing this guarantee necessarily makes the protocol vulnerable to denial-of-service attacks, since an adversary can flood the network with many messages that the protocol must deliver to all parties.

In this paper, we make several contributions towards realizing the goal of efficient, scalable byzantine agreement over a gossip network:

1. We define "gossip with abort," a relaxed gossip model that can be efficiently realized with minor modifications to existing gossip protocols, yet allows for significant savings in communication compared to the full point-to-point model.

2. Our protocols work in a *graded* PKI model, in which honest parties only have partial agreement about the set of participants in the protocol. This model arises naturally in settings without trusted setup, such as the "permissionless" setting underlying many blockchain protocols.

3. We construct a new, player-replaceable BA protocol in the graded PKI model. The concrete communication complexity of our protocol, for typical parameter values, is more than 25 times better than the current state-of-the-art BA protocols in the honest-majority setting.

Contents

1 Introduction

Byzantine agreement (BA) [32, 26] is a cornerstone for interactive protocols that require reaching consensus in hostile environments. It considers n parties who wish to jointly agree on one of their inputs in the presence of malicious agents. The problem is traditionally captured by requiring the protocol to satisfy *consistency*, *validity*, and *termination*, even when a subset of up to f colluding and cheating parties try to prevent it.

The motivation for this work is massively-scalable BA protocols, which are intended to run with thousands (or even hundreds of thousands) of participants, over the public internet. We focus on *synchronous* BA, where the protocol is assumed to proceed in rounds. Such protocols are already being deployed in the context of blockchain protocols, e.g., [11, 13].

The goal: reducing complexity. A central question in this setting is bounding the resources needed for BA as a function of the number of parties n (e.g., the round and communication complexity). The classical communication model, used in numerous seminal results, considers a complete communication graph where parties are connected via *pairwise channels*. In this setting, if every party naïvely sends a message of size ℓ to every other party, the communication overhead is $\Omega(n^2\ell)$. While this may be reasonable when the number of parties is small, communication that is quadratic in the number of parties becomes infeasible for massively-scalable protocols.

Luckily, a long series of results have shown that BA protocols with sub-quadratic communication complexity and expected constant rounds are possible [11, 1]. At a high level, the main technique for achieving this is to randomly sample a small committee, then run an "internal" BA protocol in which only committee members speak, rather than the entire set of parties. The committee size can be polylogarithmic in n; therefore, even if the internal BA protocol has quadratic complexity in the size of the committee, as a function of the total number of parties the communication complexity can still be $O(n \cdot \mathrm{poly}(\log n))$.

Player-replaceable protocols. Moreover, to guarantee security against adaptive adversaries (who could potentially corrupt an entire committee once they learn the identities of its members), state-of-the-art consensus protocols usually require an additional property: *player-replaceability* [29]. In a player-replaceable protocol, parties do not keep secret state across rounds (other than their setup information); hence, a different committee can be sampled in each round of the protocol. Intuitively, an adversary cannot tell who will be on the committee until the members reveal themselves by sending messages, but once the messages are sent, corrupting the committee is no longer helpful, since the next committee will be sampled independently.

Since any party can potentially be selected to a committee at any round, all honest parties need to receive the messages sent by committee members. Technically, it is possible for the committee members to send direct messages to every other party over point-to-point links, and indeed, this is how the formal communication model is defined in many protocols. In practice, however, propagating messages to all parties is usually accomplished using a *gossip protocol*.

Gossip. The function of a gossip protocol (also referred to as a *diffusion channel* [22]) is to propagate messages to all parties, but without requiring full connectivity between them. A simplified model to achieve this involves connecting parties using a low-degree communication graph, such that each party has a small number of peers, but the distance between any two parties is still small. Each party forwards every new message it receives to all of its neighbors.

Currently deployed protocols (such as libP2P's gossipsub [34], or Bitcoin's p2p network [16]) use more complex heuristics (such as choosing the communication graph randomly, and updating it dynamically), but for the purposes of this paper they behave similarly, and our constructions based on low-degree communication graphs can be adapted to these gossip protocols with very minor modifications.

The real power of gossip. Gossip protocols are used in practice to propagate messages in player-replaceable consensus protocols (here we include both BA and state machine replication, which can be thought of as a "continuous" BA). However, in most existing protocols, the formal communication model is slightly different:

- On one end of the scale, the protocols assume only a *multicast* functionality, that guarantees delivery of honestly generated messages to all honest parties, but not necessarily of adversarially generated messages. The actual gossip protocols guarantee something stronger (looking ahead, we will use these stronger guarantees to reduce the concrete communication complexity).

- At the other end of the scale, some consensus protocols assume that every message sent *or received* by an honest party will eventually be received by every honest party.

  This property can be achieved in the fully connected point-to-point communication setting by having parties always "echo" every message they receive to all of the other parties. Indeed, many classical BA protocols make use of this idea (in particular, the classical protocol for implementing the "gradecast" primitive, by Feldman and Micali [20] has an explicit echo step; this primitive subsequently formed the basis for a long line of Byzantine agreement protocols).

Unfortunately, forcing the gossip protocol to satisfy the latter guarantee opens the door to simple denial-of-service (DoS) attacks by a malicious adversary: since *every* message must be propagated, the adversary can send a huge number of different messages, and force the honest parties to retransmit them all. Locally checking the validity of a message might not prevent this attack: for example, an adversary may be able to generate many different messages that are each valid individually, and send each to a different honest party. Moreover, dealing with this attack by simply dropping messages that appear to be from a malicious user might cause consensus to fail, since the protocol security in the "strong" gossip model can rely on all honest parties receiving the same set of messages from adversary (albeit with different timings).

Gossip rounds. To summarize briefly, the setting we are targeting is one in which the number of participants is huge, making full connectivity infeasible and necessitating the use of protocols with subquadratic communication. In this type of setting, gossip protocols are already widely deployed as a communication substrate, and Algorand-style self-selection of small committees is the norm in order to reduce communication complexity.

Apart from the very rare exceptions of protocols that explicitly consider gossip, BA protocols are "translated" into this setting from the fully connected model. In this case, every round of communication in the original protocol corresponds to multiple rounds of gossip. We define a *gossip round* as the equivalent of the original (fully connected) protocol round: the time it takes for a message to reach all honest nodes in the communication graph. This allows us to make an apples-to-apples comparison between the round-complexity of our protocol and existing constructions.

Note that the length of each gossip round depends on two parameters: the *diameter* of the honest graph and the *latency* of its underlying point-to-point links. The diameter of the honest graph is an assumption about the gossip network's structure (or a bound if we construct the graph dynamically), and does not depend on the protocol. The link latency, on the other hand, usually depends on the size of the messages sent over the links (to prevent simple DoS attacks, existing gossip network generally receive and verify an entire message before retransmitting it).

1.1 Our Contributions

In this paper, we make several contributions towards realizing practical, massively-scalable BA:

- **Graded gossip with abort**. Our first contribution is identifying a gossip primitive that is powerful enough to allow significant communication-complexity gains compared to plain multicast over a fully connected network, but at the same time can be implemented with essentially the same cost as multicast over typical gossip communication graphs. We call the new primitive a *d-graded gossip with abort*.

- **Working directly in a graded PKI model**. Our construction of graded gossip with abort can be realized directly in the *graded* PKI model. In this model, honest parties may not know the number or the identities of other parties in advance, and may have only partial consensus on who is allowed to participate in the protocol. This type of PKI is much easier to construct in a setting where trusted setup is not available. In fact, our requirements from graded PKI are general enough that we can model partial consensus at a *message* level using the same definition (e.g., when parties need to self-select using a VRF, but do not fully agree on the number of participants). Since all of our protocols in this paper use *d*-graded gossip as a substrate, they also work directly in the graded PKI model.

  Looking ahead, since our BA protocol supports agreement on sets, it can be used to bootstrap a graded PKI into a full PKI with a single BA invocation (by using the BA to agree on the set of valid public keys).

- **Constructing gradecast and graded crusader agreement**. We show how to realize two classic primitives: *gradecast* and *graded crusader agreement*, over gossip with abort. This allows "transparent" compilation of any BA protocol based entirely on these primitives to the gossip-with-abort model, thereby improving communication complexity by an almost linear factor. Examples include the phase-king protocol of Berman, Garay and Perry [7], which can be formulated using crusader agreement, and the simple gradecast protocols of Ben-Or, Dolev, and Hoch [5, 6].

- **Concretely efficient BA**. We construct a new, player-replaceable BA protocol in the graded PKI model. The *concrete* communication complexity of our protocol, for typical parameter values, is more than 25 times better than the current state-of-the-art BA protocols in the honest-majority setting (e.g., for 800 participants, it requires less than 2MiB of communication per gossip peer, as opposed to more than 50MiB for previous protocols). We note that the efficiency gains kick in even when all parties are honest (in this case our protocol is still more than 18 times better than the state-of-the-art. (See Section 4.5.1 for a detailed complexity comparison)

  Our protocol achieves these gains by making slightly stronger (albeit still reasonable) assumptions about the adversarial corruption model and network connectivity of honest parties (honest parties remain connected throughout the protocol, and the adversary can corrupt parties adaptively, but corruption is delayed for several communication rounds after a party is selected by the adversary).

  Our protocol is competitive in terms of latency, and terminates in expected 21 gossip rounds. Moreover, because our protocol does not have any "large" communication rounds or "large" messages (as opposed to, e.g., Micali and Vaikuntanathan [30]), in practice the length of each gossip round can be made much shorter, so the real-world latency of our protocol will usually be better than state-of-the-art protocols that have fewer rounds but larger messages.

1.2 Technical Overview

We proceed to explain our contributions in more detail.

Graded gossip with abort. As mentioned above, our starting point is the observation that the standard gossip primitive, as defined, e.g., in [4, 9, 8, 28], which guarantees that *every* message received by an honest party by round r will be received by *every* other honest party by round $r + 1$, is too strong for an efficient realization against malicious adversaries.

In the terminology of secure multiparty computation (MPC), one can view the standard definition of gossip as *gossip with guaranteed output delivery*, which forces honest parties to output bad values entered by the adversary. Our relaxation can be viewed as a form of *gossip with restricted abort* (hereafter, referred to as *gossip with abort*), by enabling honest parties to output $\perp$ when the message sender is malicious. Honest messages, on the other hand, are still guaranteed to be delivered.[1]

We refer the reader to Section 3.2 for the formal definition. Loosely speaking, we require the following properties from d-graded gossip with abort:

- **Validity:** Once an honest party P_i gossips a value v, it is guaranteed that all honest parties output (P_i, v) together with the highest grade d.

- **Consistency:** In case an honest party outputs (P_j, v) with grade $g > 1$ in round r, then every honest party will output (P_j, v^*) with grade g^* by round $r + 1$ such that $v^* \in \{v, \perp\}$ and $|g - g^*| \le 1$.

- **Uniqueness:** For each gossip session, if an honest party outputs (P_j, v) and (P_j, v'), then $v' \in \{v, \perp\}$.

- **Unforgeability:** In case an honest party outputs (P_j, v) and P_j is honest, then P_j indeed gossiped the value v in a previous round.

Intuitively, *validity* ensures that honest messages are always delivered with the highest grade, and *consistency* ensures that if a corrupted party tries to equivocate and send different values to different honest parties, e.g., send v to P_i and $v' \ne v$ to P_j, then it could be that P_i outputs v and P_j outputs v' in round r, but later, in round $r + 1$ both parties will output $\perp$. The underlying idea is that it is sufficient for the honest parties to detect a cheating sender, but they do not have to agree on the contents of the messages that incriminate the sender (it could be that different honest parties will obtain a different pair of messages).

We note that the definition does not guarantee that all honest parties identify a cheating sender. For example, consider $d = 2$; then, it could be that an honest P_i outputs (P_k, v) with grade $g = 2$ in round r, and later an honest P_j outputs $(P_k, \perp)$ with grade $g = 1$ in round $r + 1$. In this case, *consistency* does not ensure that P_i will also output $(P_k, \perp)$ in round $r + 2$, since P_j received grade $g = 1$ (*consistency* ensures this for grade $g > 1$). Increasing the number of different grades (by increasing d) can prevent this special case (since inconsistencies can only happen if a party received a message with grade 1). Intuitively, this is the reason we need $d > 2$ for the d-graded gossip on which we base our efficient BA protocol.

The graded-PKI model. More formally, we define d-graded gossip with abort in a model that generalizes the standard public-key infrastructure (PKI) model for signatures, called the *d-graded PKI model* [3, 23, 15]. Similarly to standard PKI, in this model every party holds a private signing key and obtains a vector of verification keys; unlike standard PKI, the parties do not necessarily agree on the verification keys. Instead, each party assigns a grade in $\{0, \ldots, d\}$ to each verification key, and it is guaranteed that: (1) for every honest party[2] P_i, all other honest parties agree on the verification key of P_i and give it the highest grade d, and (2) for every verification key vk, if two honest parties received vk then the corresponding grades differ by at most 1.

---

[1] Security with *restricted* abort [12] enables a specific set of parties to abort the protocol if corrupted; in our setting, only a corrupted sender can cause an abort.

[2] A party is *honest* for the purposes of this definition if it is not corrupted at any point during the protocol.

At a high level, grades represent "bounded disagreement" between the parties. This notion of graded PKI is particularly suitable for permissionless protocols in which parties do not necessarily agree on the set of participates, and is often the first step in constructing a full PKI "from scratch."

Realizing graded gossip with abort in the graded PKI model. Our first technical result is a simple d-graded gossip with abort protocol in the d-graded PKI model. We consider an incomplete communication network in which each party knows the set of its neighbors (but need not have a global view of the network); in addition, we consider a public bound on the *honest diameter* D, i.e., that every pair of parties is connected by a path of length at most D consisting of honest nodes (this is called a *strongly connected network* in [4]).

At a high level, gossip with abort involves a small modification to the naïve flooding protocol (on the partial communication graph). In order to gossip a value v, the sender P_i signs v to obtain the signature σ, and sends $(\mathsf{vk}_i, v, \sigma)$ to all its neighbors. Upon receiving a new message (vk, v, σ), a party P_j gossips this message to all neighbors (from which it did not previously receive it).

The difference between naïve flooding and gossip with abort is that if two different messages were received from the same public key vk, the information that will be flooded is a proof that vk behaved maliciously (consisting of a pair of different messages signed by the same key). The key insight is that we do not care about the consistency of the messages comprising the malfeasance proofs. Thus, if party P_i sent a malfeasance proof for a key vk to a neighbor, it will not send any additional messages for vk to that neighbor (even if it receives different malfeasance proofs for vk).

Even when a cheating sender floods the honest parties with many messages, each honest party sends at most two messages for that sender; therefore, the communication cost that honest parties send is at most $2|E|$ times the message size, where E is the set of edges in the communication graph.

Basing gossip with abort on existing flooding protocols. While we prove our construction directly, given a predetermined communication graph, the same idea can be applied much more widely. For example, both practical flooding implementations (such as Ethereum and Bitcoin's underlying gossip networks) as well as theoretical constructions, such the flooding protocol of Liu-Zhang et al. [27] (which also takes into account resource weights of participating nodes), can easily be adapted to achieve d-Graded Gossip in the d-Graded PKI model, with almost no overhead. For example, [27] provide a "Skeleton for Flooding Protocols" in which each node maintains a list of already-relayed messages. The change required is to consider a message as already relayed if two different messages from the same key have already been relayed (even if the message value is different from either). The underlying proofs of propagation can be used essentially unchanged to show that if two or more different messages were signed by the same key, all honest parties will receive this information (albeit possibly not the same two messages).

Graded threshold gossip. Next, we identify a new primitive, f-*threshold-gossip*, which abstracts the operation of "propagate signatures and check if $f + 1$ distinct parties signed a value," that can also be efficiently (and very simply) implemented over gossip-with-abort. The high-level idea is that each party can count malfeasant public-keys as "supporting" every possible value (intuitively, if we know a public key is malicious, then the adversary *could have* signed every value). This allows us to maintain consistency of the output even though honest parties may not all receive the same signatures. For example, if an honest party P_i received $f + 1$ signatures for a value v over gossip with abort, it is possible that P_j received some aborts instead of signatures for v. However, gossip with abort guarantees that the *sum* of signatures and aborts will have to be at least $f + 1$ for P_j. We refer the reader to Section 3.3 for the formal definition and realization.

Efficient Byzantine agreement. Using the building blocks described above, we construct a new and efficient BA protocol. Our construction is inspired by the expected constant-round BA protocol of Abraham et al. [2]. However, their protocol requires parties to send "certificates" proving that $f + 1$ parties signed some value, necessitating quadratic communication.

Instead, the parties in our new protocol communicate only using gradecast (which we realize using graded gossip with abort) and threshold-gossip. The consistency properties of these primitives let us remove the requirement to explicitly form certificates. However, because the arrival times of messages may be inconsistent, the protocol construction is non-trivial, and requires careful analysis. Section 4 describes the protocol in detail.

We note that disagreement about message timing is one of the main challenges for consensus even in the full PKI model. By abstracting it as a "graded agreement" on the messages, we can include additional sources of bounded disagreement essentially "for free" (e.g., by combining the bounded disagreement about in-protocol message timing with the bounded disagreement about key validity.) Thus, the relaxation to a graded PKI arises naturally from our BA protocol construction.

Bootstrapping graded to full PKI using set agreement. Our BA protocol generalizes standard byzantine agreement to agreement on sets of strings. Briefly, in a set agreement,[3] the inputs of the parties are sets of strings, and the output is a set. The validity guarantees for set agreement ensure that every string in the intersection of honest inputs is contained in the output, while every string in the output set is contained in the union of honest inputs. (See section 4.1 for the formal definition of set agreement.)

This generalization allows us to easily bootstrap a graded PKI into a full PKI. Each party uses as its input set the keys that (locally) received grade d. Since honest keys always receive grade d, they are in the intersection of all honest inputs and thus must appear in the output set. On the other hand, if a key appeared in the output set, it must have been included in the input set of at least one honest party, and hence must have grade at least $d - 1$ for all honest parties.

1.3 Related Work

While gossip protocols [17, 24] are widely studied in the literature, their use to improve the communication of consensus protocols has only recently been explored. In the context of blockchain protocols, the famous protocol of Nakamoto utilizes gossip as a means of communication to keep the communication among a large potential number of miners as low as possible. While early works have treated the gossip layer in the Nakamoto protocol as a black box, a recent work by Coretti et al. [14] gave a new Byzantine-resilient gossip protocol and studies the effect of using it as a communication layer in the Nakamoto protocol. In closely related work, Matt et al. [28] studied the effect of adaptive corruptions on gossip protocols. In the context of classical consensus protocols (i.e., not based on the longest chain rule), recent works by Momose and Ren [31] and Tsimos et al. [33] showed how to reduce the communication complexity of Byzantine agreement by implementing particularly expensive protocol steps over a gossip network. These works are similar in spirit to our own; however, our protocol is optimized for concrete efficiency in a permissionless setting. Moreover, our protocol is the first to have an expected constant gossip-round complexity, and can be directly executed over a graded PKI.

Organization of the paper. Preliminaries can be found in Section 2. Section 3 presents the graded-PKI model and the constructions of graded gossip, gradecast, threshold gossip, and graded crusader agreement. In Section 4 we present our byzantine agreement on sets protocol and prove its security.

---

[3]Not to be confused with k-set agreement [10].

2 Preliminaries

We consider synchronous protocols that proceed in a round-by-round manner. The parties communicate over a partial communication graph, where every party is connected to each of its neighbors by an ideally authenticated channel; that is, the adversary can view communication sent over those channels, but cannot drop or inject messages. We denote the number of parties in our protocols (i.e., the number of nodes in the graph) by n; however, the parties are not aware of n and each party is only aware of its local neighbor-set.

All our constructions assume the existence of digital signatures which are existentially unforgeable under chosen message attack. We denote the security parameter by λ, and all our security guarantees hold with all but negligible probability in λ. Note that since a party in a permissionless protocol may not have a globally-agreed identifier, we identify parties by their corresponding signature verification keys. For example, in the graded gossip protocol, the output is of the form (vk, v) where vk is the verification key of the sender.

We consider PPT adversaries that have weak adaptive-corruption capabilities; namely, that are delayed adaptive [28]. That is, the adversary can announce that it wishes to corrupt a party dynamically during the course of the protocol, but it only gets hold of the corrupted party after some time elapses. In all our constructions, we consider this delay to be the duration of a single gossip execution. Once a party is corrupted, the adversary can see its internal state and instruct this party to behave in an arbitrary manner; that is, the adversary is malicious (aka Byzantine).

3 Graded PKI and Graded Gossip

In this section, we present the building blocks with which we construct the BA protocol in Section 4. In Section 3.1 we define the graded PKI model; in Section 3.2 we define and construct graded gossip with abort and gradecast; in Section 3.3 we use those primitives to construct threshold gossip; finally, in Section 3.4 we show how to realize graded crusader agreement.

In the classical BA model, the set of participating parties is public and known ahead of time. When considering an honest majority, BA protocols cannot exist in the plain model [26, 21], and usually rely on a *Public-Key Infrastructure* (PKI): each party has an associated public verification key for a signature scheme, which is known to all participants.

Modern, massively-scalable BA protocols, on the other hand, often target a *permissionless* setting, in which there is no initial consensus about the set of participants or their public keys. Previous works [3, 23, 15] have shown how to construct a PKI "from scratch" in this setting based on resource-bound assumptions, such as Proofs of Work (PoW).

These protocols for achieving PKI typically achieve a weaker primitive first: a *graded PKI* (also known as a *ranked PKI* [3]), in which parties have only partial agreement on the participating identities. The graded PKI abstraction can also capture partial consistency in validating eligibility proofs for committee sampling. For example, the Algorand-style method of determining eligibility by comparing a party's VRF output to a difficulty threshold may yield only partial consistency if honest parties can disagree slightly about the threshold value.

Our BA protocol can utilize a graded PKI directly, obviating the (expensive) conversion to a full PKI. Below, we formally define the graded PKI model we use, and show how it can be used to efficiently implement a graded version of "gossip with abort."

3.1 d-Graded PKI

In the *d-graded* PKI model, every *honest* party P_i holds a signing key sk_i for a digital-signature scheme, and a corresponding verification key vk_i. In addition, there exists a protocol with a method $\mathsf{GradeKey}$ that accepts a verification key vk and session id sid and outputs a grade $g \in \{0, 1, \ldots, d\}$.

We denote $\mathsf{GradeKey}_i(\mathsf{vk}, \mathsf{sid})$ the result of party P_i executing $\mathsf{GradeKey}$ on input $(\mathsf{vk}, \mathsf{sid})$. While $\mathsf{GradeKey}$ is a protocol, rather than a deterministic function of the public key and session (e.g., the grade can depend on interaction with other parties, and additional local information), we assume without loss of generality that for a given party P_i and key/session $(\mathsf{vk}, \mathsf{sid})$, $\mathsf{GradeKey}_i(\mathsf{vk}, \mathsf{sid})$ always returns the same value (this is without loss of generality since P_i can always store a table of inputs that have already been queried).

A d-graded PKI must satisfy the following properties (except with negligible probability in the security parameter):

1. **Validity:** For every two honest parties P_i and P_j and every session id sid, if $\mathsf{GradeKey}_i(\mathsf{vk}_i, \mathsf{sid}) = d$ then it holds that $\mathsf{GradeKey}_j(\mathsf{vk}_i, \mathsf{sid}) = d$ (i.e., if a party sees its own key/session pair with the highest grade, then every other honest party will also give it the highest grade).

2. **Graded Consistency:** For every key/sessionid pair $(\mathsf{vk}, \mathsf{sid})$, and every two honest parties P_i and P_j,
$$|\mathsf{GradeKey}_i(\mathsf{vk}, \mathsf{sid}) - \mathsf{GradeKey}_j(\mathsf{vk}, \mathsf{sid})| \leq 1.$$

The significance of d. Note that one can always trivially construct a d'-graded PKI from a d-graded PKI when $d' < d$, by defining $\mathsf{GradeKey}(\mathsf{vk}, \mathsf{sid})' = \max\{0, \mathsf{GradeKey}(\mathsf{vk}, \mathsf{sid}) - (d - d')\}$. However, the converse may not be true (e.g., given a 2-graded PKI, it might be hard to construct a 3-graded PKI). Thus, in our constructions, we specify the *minimal d* they require.

3.1.1 Player Replaceable Protocols and Active Parties

In each round, only a subset of the parties may be allowed to speak. These parties are considered **active**. The typical use case of a player-replaceable protocol involves electing a small committee of parties in each round, such that only parties on the committee are active.

We abstract the mechanism for determining the active parties into the graded PKI. This allows us to model protocols in which there is only partial agreement about which parties are active in each round. To do this, we consider the round as part of the session id, and denote a party P_i *active* in session sid if $\mathsf{GradeKey}_i(\mathsf{vk}_i, \mathsf{sid}) = d$. (This use of the graded PKI abstraction is why the validity definition is only required to hold when $\mathsf{GradeKey}_i(\mathsf{vk}_i, \mathsf{sid}) = d$.)

See Section 3.1.3 for an example of how this can be implemented.

3.1.2 Bounding Corruptions in the Graded-PKI Model

In the standard model, it is customary to bound the number of parties corrupted by the adversary. This works in the full PKI model as well, since every public key is associated with a single party. In the graded PKI model, however, the adversary can create many different keys without corrupting additional parties.

Moreover, for player-replaceable protocols, the honest-majority requirement must hold separately for each committee. Thus, we need to modify the notion of an honest majority to take this into account.

f-faulty and f-bounded executions. Let H be the set of honest verification keys (keys generated by parties that are still honest at the end of the protocol). Define $V_{j,\mathsf{sid}} = \{\mathsf{vk} \mid \mathsf{GradeKey}_j(\mathsf{vk}, \mathsf{sid}) > 0\}$ to be the set of verification keys accepted by party P_j for session id sid during a protocol execution and $F_{j,\mathsf{sid}} = V_{j,\mathsf{sid}} \setminus H$ to be the subset of those keys that are not honest. We say a protocol execution is f-*faulty* if for every honest party P_j and every sid, it holds that $|F_{j,\mathsf{sid}}| \leq f$, and that an execution is f-*bounded* if for every sid it holds that $|\bigcup_i F_{i,\mathsf{sid}}| < f$.

We emphasize that honest parties might disagree on the set of accepted keys, so an f-faulty execution might not be f-bounded.

Definition 3.1 (*f*-faulty adversaries)**.** An adversary $\mathcal{A}$ is *f-faulty* (resp., *f*-bounded) with respect to a protocol Π, if the probability that $\mathcal{A}$ interacting with an execution of Π is not *f*-faulty (resp., *f*-bounded) is negligible in the security parameter (where the probability is over the coins of the adversary, of the honest parties, and of the setup).

Note that in the full PKI setting, an adversary that can corrupt at most f parties is f-bounded against any protocol in which $\mathsf{GradeKey}(vk, sid)$ checks whether vk is a valid key. Similarly, consider in a Algorand-style trusted PKI setting, with VRF keys which are used to sample committees of expected size λ out of a population of N parties, such that each committee is denoted with a separate session id. In this case, an adversary that can corrupt at most a c-fraction of the population $(c + \varepsilon) \cdot \lambda$-bounded for every constant $\varepsilon > 0$ (λ is the security parameter).

3.1.3 Examples of Graded PKI

In prior work [3, 23, 15], a graded PKI is a setup protocol, executed by the parties ahead of time, such that each party outputs a vector of verification keys with grades (with the same consistency requirements we have).

In this work, we use the graded PKI to capture graded eligibility proofs for committee selection as well. In this case, the public keys may not be known in advance to honest parties. Below, we give three examples of graded PKI implementations to give some motivation for our slightly different interface.

Full PKI. A full PKI (where parties agree on the entire set of valid public keys ahead of time) can trivially serve as a d-graded PKI for any d: $\mathsf{GradeKey}(vk, sid) = d$ if vk is a valid key in the full PKI, and 0 otherwise.

PoW-based PKI Setup. Andrychowicz and Dziembowski [3] show how to construct a graded PKI (which they call a *ranked key set*) without setup assumptions, using Proofs of Work instead. This protocol runs as a setup phase, requiring interaction between all participating parties and outputs the entire set of public keys with non-zero grades to each party. In this case, $\mathsf{GradeKey}(vk, sid)$ is the grade output by the protocol if vk is in the set, and 0 otherwise.

VRF-based Eligibility. A common method of randomly electing a committee with expected size n, out of a population of N parties, is Algorand-style self-sortition [11]: each party computes a VRF of the session id (which typically includes the round number), and self-selects if the output (interpreted as a value in the range $[0, 1]$) is less than n/N. The VRF output also serves as a publicly-verifiable proof of eligibility.

However, in some cases the number of parties N is not in full consensus. For example, parties can self-select by running a proof-of-work solver locally, and N is estimated locally by each party based on observed published proofs. In this case, parties can still generate a graded eligibility proof, assuming that for every two honest parties P_i and P_j, the disagreement about N is bounded by some fixed constant k.

For a public key vk and session id sid, let $\mathsf{GradeKey}(vk, sid)$ return g if a VRF proof was previously received that shows $\mathsf{VRF}_{vk}(sid) < n/(N - (k-1) - (d-g) \cdot k)$. In this case, an honest party P_i considers itself active in session sid if $\mathsf{VRF}_{vk}(sid) < n/N$. Since every other honest party has $N' \geq N - k$, if another honest party considers itself active, it will receive grade d.

3.2 d-Graded Gossip

In this section, we define a weaker gossip primitive, d-Graded Gossip (with abort), that can be realized in the d-graded PKI model, for $d \geq 3$.

A protocol π is a *d-graded gossip* protocol if it has the following API (i.e., input/output interface) and satisfies the following properties: The API is defined by the method $\texttt{GradeGoss}(\mathsf{sid}, v)$, where v is a value and sid is an arbitrary unique session id (string). The output of a party is of the form $(\mathsf{vk}, \mathsf{sid}, v, g)$ for a public key vk, a session id sid, a value v, and a grade g; note that a party may produce more than one output per $(\mathsf{vk}, \mathsf{sid})$ pair.

The session id will be used to bound communication complexity by allowing each party to gossip at most one message for each sid. (Typically, the session id would contain the protocol round number or a specific role in the protocol.)

The *d*-graded gossip protocol π is defined in the *d*-graded PKI model, and guarantees the following, for every round r and every session id sid:

1. **Validity:** If P_i is honest, invokes $\texttt{GradeGoss}(\mathsf{sid}, v)$ at round r, and this is the first time that P_i made any call to $\texttt{GradeGoss}(\cdot)$ with sid, then every honest party will receive $(\mathsf{vk}_i, \mathsf{sid}, v, d)$ from π at round $r+1$, and will never receive $(\mathsf{vk}_i, \mathsf{sid}, v', g')$ from π for $(v', g') \neq (v, d)$.

2. **Consistency:** If an honest party received $(\mathsf{vk}, \mathsf{sid}, v, g)$ from π at round r with $g > 1$, and it remains honest in round $r+1$, then every honest party will receive $(\mathsf{vk}, \mathsf{sid}, v^*, g^*)$ from π by round $r+1$, for $v^* \in \{v, \bot\}$ and $|g^* - g| \leq 1$.

3. **Uniqueness:** If an honest party received $(\mathsf{vk}, \mathsf{sid}, v, g)$ and $(\mathsf{vk}, \mathsf{sid}, v', g')$ from π, then $v' \in \{v, \bot\}$.

4. **Unforgeability:** For all v and g, if P_i is honest and an honest party received $(\mathsf{vk}_i, \mathsf{sid}, v, g)$ from π, then P_i invoked $\texttt{GradeGoss}(\mathsf{sid}, v)$.

When a party receives $(\mathsf{vk}, \mathsf{sid}, \bot, g)$, we call this an **equivocation proof** for $(\mathsf{vk}, \mathsf{sid})$.

3.2.1 Partial Communication Graphs

Our graded gossip implementation assumes an underlying directed communication graph $G = (V, E)$, where the vertices V represent parties, every two parties P_i and P_j such that $(P_i, P_j) \in E$ have a direct communication link (that lets P_i send messages to P_j), and every party knows their immediate neighbors in G.

We note that no party is required to know the entire graph G; thus, it is possible to construct G "locally."

Definition 3.2 (Honest Distance). The *honest distance* between two parties $P_i, P_j \in V$ is the length of the shortest path between P_i and P_j in G that passes only through honest nodes.

If there is no honest path between P_i and P_j the honest distance is defined to be ∞.

Note that in an adaptive corruption setting, the honest distance between two nodes can grow when additional nodes are corrupted.

Definition 3.3 (Honest Diameter). The *honest diameter* of a graph G is the maximum honest distance between two honest nodes.

3.2.2 Implementing *d*-Graded Gossip from *d*-Graded PKI

Protocol 1 is a minor modification of the naïve flooding protocol, in which each party may send an equivocation proof instead of a message if the sender is malicious. Each party maintains a table of "equivocating keys." Denote by $B_{\mathsf{vk}, \mathsf{sid}}$ the table entry for key vk in session sid; that is, $B_{\mathsf{vk}, \mathsf{sid}} = 1$ if key vk was "caught" equivocating in session sid (we initialize all values to 0, so for every pair $(\mathsf{vk}, \mathsf{sid})$ for which the value $B_{\mathsf{vk}, \mathsf{sid}}$ has not explicitly been set, $B_{\mathsf{vk}, \mathsf{sid}} = 0$). Note that this table is local to each party (and honest parties may not necessarily agree on the value of $B_{\mathsf{vk}, \mathsf{sid}}$).

Gossip vs. Communication Rounds. The formal definition of graded gossip with abort specifies that all parties receive honest messages within one "round". This makes the definition compatible with the common definitions of gossip and multicast (and makes the higher-level protocols using it much easier to read). This clearly cannot work if there are honest parties that are not directly connected (as is the case in our setting). Thus, we treat the actual communication rounds as "subrounds" of the gossip protocol. The actual number of communication rounds for each "gossip round" (which is what we mean by "round" in the remainder of the paper) depends on the gossip protocol implementation (in Protocol 1, it is an upper bound on the honest communication distance between two honest parties).

Protocol 1: $\pi^{(d)}_{\text{graded-gossip}}$

The protocol is parameterized by a maximum message size L and a maximum grade d. Each "gossip round" of the protocol consists of D subrounds; subrounds are the underlying synchronous communication rounds, and are counted consecutively from the beginning of the protocol (subround t is in protocol round $\lfloor t/D \rfloor$).

In each subround:

1. When a party P_i calls $\texttt{GradeGoss}(\text{sid}, v)$, it signs (sid, v) as $\sigma \leftarrow \texttt{Sign}_{\text{sk}_i}(\text{sid}, v)$. Next, P_i sends $(\text{sid}, v, \text{vk}_i, \sigma)$ to itself (and handles it as a received message).

2. When a party receives $(\text{sid}, v, \text{vk}, \sigma)$ from itself or a neighbor perform one of the following actions:

Case 1: $|v| > L$, $\texttt{GradeKey}(\text{vk}, \text{sid}) = 0$, or σ is not a valid signature under vk of (sid, v). In this case, drop the message.

Case 2: $B_{\text{vk,sid}} = 1$ (we already have an equivocation proof for (vk, sid)). In this case, drop the message.

Case 3: $B_{\text{vk,sid}} = 0$ and $\texttt{GradeKey}(\text{vk}, \text{sid}) > 0$, but a previous message $(\text{sid}, v', \text{vk}, \sigma')$ was received, such that $v' \neq v$:

    i. set $B_{\text{vk,sid}} = 1$

    ii. send $(\text{sid}, v, \text{vk}, \sigma)$ to all neighbors.

    iii. output $(\text{vk}, \text{sid}, \bot, \texttt{GradeKey}(\text{vk}, \text{sid}))$

Case 4: $B_{\text{vk,sid}} = 0$, $\texttt{GradeKey}(\text{vk}, \text{sid}) > 0$ and no previous message $(\text{sid}, v', \text{vk}, \sigma')$ was received:

    i. send $(\text{sid}, v, \text{vk}, \sigma)$ to all neighbors.

    ii. output $(\text{vk}, \text{sid}, v, \texttt{GradeKey}(\text{vk}, \text{sid}))$

Claim 3.4. *Let P_i be an honest party that outputs $(\text{vk}, \text{sid}, v, g)$ at subround t of $\pi^{(d)}_{\text{graded-gossip}}$ (Protocol 1), such that $g > 1$. For all $j \geq 0$, every node whose honest distance from P_i at subround $t + j$ is at most j will have output $(\text{vk}, \text{sid}, v', g')$ by subround $t + j$, where $v' \in \{v, \bot\}$ and $g' \geq g - 1$.*

Proof. Consider the event that the adversary succeeds in forging an honest sender's signature; by the security of the signature scheme this happens with negligible probability. We proceed by conditioning on the complementary event and prove the claim by induction on j. The base case of $j = 0$ is trivial. For $j + 1$, consider a node $P_{i'}$ at distance $j + 1$ from P_i. By definition, since $P_{i'}$ is at honest distance $j + 1$ from P_i, there is a path of honest nodes of length $j + 1$ from P_i to $P_{i'}$ at subround $t + j + 1$. Let P_{pred} be the predecessor of $P_{i'}$ on this path; then, P_{pred} has

honest distance j from P_i at subround $t + j + 1$, and since honest distance can only grow with time, it has honest distance j from P_i at subround $t + j$. Thus, by the induction hypothesis, P_{pred} must have output $(\mathsf{vk}, \mathsf{sid}, v', g')$ at some subround $t + j'$, for $j' \leq j$ and $v' \in \{v, \perp\}$ and $g' \geq g - 1$.

First, note that since P_i output $(\mathsf{vk}, \mathsf{sid}, v, g)$ and $g > 1$, it must be that $\mathsf{GradeKey}_i(\mathsf{vk}, \mathsf{sid}) = g > 1$. Thus, by graded consistency of the PKI, it holds that

$$\mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}) \geq \mathsf{GradeKey}_i(\mathsf{vk}, \mathsf{sid}) - 1 = g - 1 > 0. \tag{3.1}$$

Next, consider the two cases which could cause P_{pred} to produce output in the byzantine gossip protocol:

Case 1: In Protocol 1, case 3 was satisfied for P_{pred} at subround $t + j'$. This means P_{pred} received $(\mathsf{sid}, v_1, \mathsf{vk}, \sigma_1)$ for some v_1 and a valid signature σ_1, and also previously received $(\mathsf{sid}, v_2, \mathsf{vk}, \sigma_2)$ for $v_2 \neq v_1$ and a valid signature v_2. Let $(\mathsf{sid}, v_2^*, \mathsf{vk}, \sigma_2^*)$ be the first such message it received (at subround $t + j'' \leq t + j' \leq t + j$).

At that point, case 4 must have been satisfied, so P_{pred} must have sent $(\mathsf{sid}, v_2^*, \mathsf{vk}, \sigma_2^*)$ to all of its neighbors, and in particular to $P_{i'}$. At subround $t + j'$, it must have sent $(\mathsf{sid}, v_1, \mathsf{vk}, \sigma_1)$ to $P_{i'}$. If at subround $t + j'$ party $P_{i'}$ has not yet output $(\mathsf{vk}, \mathsf{sid}, \perp, \mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}))$, then $B_{\mathsf{vk}, \mathsf{sid}} = 0$. By Inequality 3.1, $\mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}) > 0$, and thus this second message will satisfy case 3, causing $P_{i'}$ to output $(\mathsf{vk}, \mathsf{sid}, \perp, \mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}))$. Since $\mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}) \geq g - 1$ this satisfies the induction hypothesis.

Case 2: In Protocol 1, case 4 was satisfied for P_{pred} at subround $t + j'$. This means P_{pred} received $(\mathsf{sid}, v, \mathsf{vk}, \sigma)$ and a valid signature σ, and sent it to all of its neighbors, and in particular $P_{i'}$ would receive this message by subround $t + j' + 1$. Assume $P_{i'}$ has not already output $(\mathsf{vk}, \mathsf{sid}, v', \mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}))$ for $v' \in \{v, \perp\}$ (otherwise we are already done). In that case, for $P_{i'}$ it must hold that $B_{\mathsf{vk}, \mathsf{sid}} = 0$. By Inequality 3.1, $\mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}) > 0$; hence, one of the following must be the case:

Case 2.1: $P_{i'}$ previously received a message $(\mathsf{vk}, \mathsf{sid}, v', \dots)$ for $v' \neq v$. Then, case 3 would be satisfied for $P_{i'}$ at subround $t + j' + 1$; hence, $P_{i'}$ would output $(\mathsf{vk}, \mathsf{sid}, \perp, \mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}))$ at subround $t + j' + 1$.

Case 2.2: $P_{i'}$ did not receive any message with prefix $(\mathsf{vk}, \mathsf{sid})$. Then, case 4 would be satisfied for $P_{i'}$ at subround $t + j' + 1$; hence, $P_{i'}$ would output $(\mathsf{vk}, \mathsf{sid}, v, \mathsf{GradeKey}_{i'}(\mathsf{vk}, \mathsf{sid}))$ at subround $t + j' + 1$. $\qquad\square$

We proceed to show that the protocol $\pi_{\mathsf{graded\text{-}gossip}}^{(d)}$ satisfies the graded-gossip requirements, as long as the adversary does not disconnect honest parties on the graph (or makes the honest diameter too big). Given a communication graph G, an adversary $\mathcal{A}$, and an integer D, representing the assumed honest diameter in the protocol $\pi_{\mathsf{graded\text{-}gossip}}^{(d)}$ (see Definition 3.3), we define the event $\mathcal{E}_{\mathcal{A}, D}$ in which the honest diameter of G is D at the end of an execution of $\pi_{\mathsf{graded\text{-}gossip}}^{(d)}$ with the adversary $\mathcal{A}$; the event is defined over the random coins used by the honest parties and the adversary.

Theorem 3.5 (Graded Gossip Security). *For every delayed-adaptive PPT adversary $\mathcal{A}$ and every communication graph G, either $\Pr[\mathcal{E}_{\mathcal{A}, D}] = 0$ or, conditioned on $\mathcal{E}_{\mathcal{A}, D}$ it holds that $\pi_{\mathsf{graded\text{-}gossip}}^{(d)}$ (Protocol 1) is a d-graded gossip protocol for messages of length L in the d-graded PKI model.*

Proof. Unforgeability follows from the security of the signature scheme. If the protocol outputs $(\mathsf{vk}_i, \mathsf{sid}, v, g)$, then it must have received a message $(\mathsf{sid}, v, \mathsf{vk}_i, \sigma)$ where vk_i is the verification

key of P_i and σ is a valid signature on (sid, v). If P_i did not call $\mathtt{GradeGoss}(\mathsf{sid}, v)$, this implies σ is a forged signature.

Uniqueness can be seen by inspecting the code: if the protocol outputs $(\mathsf{vk}, \mathsf{sid}, v, g)$, then case 4 cannot be satisfied again for $(\mathsf{sid}, v', \mathsf{vk}, \sigma')$ for any v'; therefore, the only remaining output option is $(\mathsf{vk}, \mathsf{sid}, \bot, g)$.

Consistency is directly implied by Claim 3.4 and the bound on the honest diameter: round r starts at subround $D \cdot r$, while round $r + 1$ starts at subround $D \cdot r + D$, and for every two honest nodes P_i and $P_{i'}$, honest node $P_{i'}$ has honest distance at most D from P_i. Thus, by Claim 3.4 if an honest node vk outputs $(\mathsf{vk}, \mathsf{sid}, v, g)$ in round r then every honest node will output $(\mathsf{vk}, \mathsf{sid}, v, g')$ or $(\mathsf{vk}, \mathsf{sid}, \bot, g')$ by the end of round $r + 1$, with $g' \geq g - 1$.

For *Validity*, it suffices to note that an honest party P_i will never sign two different messages with the same sid; hence, case 3 can never be satisfied, and for every two honest parties P_i and $P_{i'}$ it holds that $\mathsf{GradeKey}_{i'}(\mathsf{vk}_i, \mathsf{sid}) = d$. By consistency, that implies every honest node will output $(\mathsf{vk}_i, \mathsf{sid}, v, d)$. $\qquad\square$

Theorem 3.6 (Graded Gossip Communication Complexity). *Let $|\mathsf{sid}| \leq \lambda$ and let the public keys and signatures be of size λ_{pk} and λ_{sig}, respectively (we assume both are polynomial in λ). When Protocol $\pi_{\mathsf{graded\text{-}gossip}}^{(d)}$ (Protocol 1) is implemented over a communication graph $G = (V, E)$ for messages of size L, the communication complexity for every $(\mathsf{vk}, \mathsf{sid})$ pair output by an honest party (regardless of the number of messages output) is at most*

$$2|E|(\lambda + \lambda_{pk} + \lambda_{sig} + L) = 2|E|(poly(\lambda) + L).$$

Proof. By inspection, every honest party sends a message with matching $(\mathsf{sid}, *, \mathsf{vk}, *)$ at most twice to every neighbor: once when case 4 is satisfied (after which it can never be satisfied again for the same $(\mathsf{sid}, \mathsf{vk})$ pair) and once when case 3 is satisfied (after which it can never be satisfied again for the same $(\mathsf{sid}, \mathsf{vk})$ pair). Thus, for every given $(\mathsf{sid}, \mathsf{vk})$ pair, at most two messages are sent by honest parties over each edge in the graph.

Since each message sent has the form $(\mathsf{sid}, v, \mathsf{vk}, \sigma)$, the length of each message is bounded by $\lambda + \lambda_{pk} + \lambda_{sig} + L$. Therefore, the total complexity is $2|E|(\lambda + \lambda_{pk} + \lambda_{sig} + L)$. $\qquad\square$

Corollary 3.7. *In an f-bounded execution with $|H|$ honest parties, The total communication complexity for every sid is at most $2(f + |H|)|E|(\lambda + \lambda_{pk} + \lambda_{sig} + L)$.*

Proof. The execution is f-bounded, hence the total number of keys that have non-zero grade is at most $f + |H|$. Since honest parties do not output messages from keys with grade less than 1, the total communication complexity for every sid is at most

$$2(f + |H|)|E|(\lambda + \lambda_{pk} + \lambda_{sig} + L). \qquad\square$$

Corollary 3.8. *If G has degree polynomial in the security parameter λ, then the communication complexity is bounded by $poly(\lambda) \cdot |V| \cdot L$.*

Proof. In this case, $|E| = |V| \cdot poly(\lambda)$, so

$$2|E|(poly(\lambda) + L) = 2|V| \cdot poly(\lambda) + 2|V| \cdot poly(\lambda) \cdot L = 2|V| \cdot poly(\lambda) \cdot L. \qquad\square$$

3.2.3 Implementing Gradecast from 3-Graded Gossip

Following Feldman and Micali [20], we define r-round gradecast as a protocol that has the following API and satisfies the following properties: The API is defined by the method $\mathtt{Gradecast}(\mathsf{sid}, v)$, where v is a value and sid is an arbitrary unique session id (string). The output of a party is of the form $(\mathsf{vk}, \mathsf{sid}, v, g)$ for a public key vk, a session id sid, a value v, and a grade $g \in \{0, 1, 2\}$.

An R-round gradecast protocol guarantees the following properties with overwhelming probability, for every round r and every session id sid:

1. **Validity:** If Gradecast(sid, v) is called by an honest party P_i in round r, then every honest party will output $(\mathsf{vk}_i, \mathsf{sid}, v, 2)$ to at round $r + R$.

2. **Weak consistency:** If an honest party outputs $(\mathsf{vk}, \mathsf{sid}, v, 2)$, then every honest party outputs $(\mathsf{vk}, \mathsf{sid}, v, g)$ with $g \in \{1, 2\}$.

We implement 3-round gradecast over 3-graded gossip in Protocol 2.

Protocol 2: $\pi_{\mathsf{gradecast}}$

The protocol is parameterized by a maximum message size L and uses the protocol $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ (Protocol 1) in the 3-graded PKI model.

1. When a party P_i calls Gradecast(sid, v) at round r, invoke the protocol $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ with GradeGoss(sid, (r, v)).

2. The protocol outputs $(\mathsf{vk}, \mathsf{sid}, v, 2)$ at round $r + 3$ if and only if all the following are satisfied:

   (a) It received a message $(\mathsf{vk}, \mathsf{sid}, (r, v), 3)$ from $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ at round $r + 1$.

   (b) It did not receive $(\mathsf{vk}, \mathsf{sid}, \bot, g)$ from $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ up to round $r + 3$ (for any grade g).

3. If it did not output yet, the protocol outputs $(\mathsf{vk}, \mathsf{sid}, v, 1)$ at round $r + 3$ if and only if all the following are satisfied:

   (a) It received a message $(\mathsf{vk}, \mathsf{sid}, (r, v), g)$, for $g \in \{2, 3\}$, from $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ up to round $r + 2$.

   (b) It did not receive $(\mathsf{vk}, \mathsf{sid}, \bot, g)$ from $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ up to round $r + 2$ (for any grade g).

4. If it did not output yet, the protocol outputs $(\mathsf{vk}, \mathsf{sid}, \bot, 0)$ at round $r + 3$ if *any* of the following are satisfied:

   (a) It received $(\mathsf{vk}, \mathsf{sid}, \bot, g)$ from $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ up to round $r + 2$ (for any grade g).

   (b) The first message from $\pi^{(3)}_{\mathsf{graded\text{-}gossip}}$ with prefix $(\mathsf{vk}, \mathsf{sid}, \dots,)$ was received at round $r + 3$.

Theorem 3.9 (Gradecast Security). *For every delayed-adaptive PPT adversary $\mathcal{A}$, if $\pi^{(d)}_{\mathsf{graded\text{-}gossip}}$ (Protocol 1) is a d-graded gossip protocol against $\mathcal{A}$, then $\pi_{\mathsf{gradecast}}$ (Protocol 2) is a 3-round gradecast protocol in the 3-graded PKI model against $\mathcal{A}$.*

Proof. We prove each property separately:

Validity. Since P_i is honest, it will only call GradeGoss($\cdot$) once for sid at round r; hence, by *3-graded gossip unforgeability* no equivocation proof $(\mathsf{vk}, \mathsf{sid}, \bot)$ will ever be received.

Further, P_i calls GradeGoss(sid, (r, v)) at round r; hence, by *3-graded gossip validity* all honest parties receive $(\mathsf{vk}, \mathsf{sid}, (r, v), 3)$ by round $r + 1$ and no equivocation. Thus, the protocol will satisfy the conditions to output $(\mathsf{vk}, \mathsf{sid}, v, 2)$ at round $r + 3$.

Weak Consistency. Suppose an honest party outputs a message $(\mathsf{vk}, \mathsf{sid}, v, 2)$ at round $r+3$. This means the party received a message $(\mathsf{vk}, \mathsf{sid}, (r, v), 3)$ at round $r+1$, and it did not receive an equivocation proof $(\mathsf{vk}, \mathsf{sid}, \bot, \cdot)$ up to round $r+3$. By *3-graded gossip consistency*, every other honest party must have received $(\mathsf{vk}, \mathsf{sid}, (r, v), g)$ or $(\mathsf{vk}, \mathsf{sid}, \bot, g)$ by round $r+2$, for $g \in \{2, 3\}$; however, it cannot be $(\mathsf{vk}, \mathsf{sid}, \bot, g)$ since then (by *3-graded gossip consistency*) every honest party must have received $(\mathsf{vk}, \mathsf{sid}, \bot, g')$ by round $r+3$ (for $g' > 0$). Thus, every honest party must have received $(\mathsf{vk}, \mathsf{sid}, (r, v), g)$ by round $r+2$, and must have output either $(\mathsf{vk}, \mathsf{sid}, v, 1)$ or $(\mathsf{vk}, \mathsf{sid}, v, 2)$ at round $r+3$. Thus, weak consistency is satisfied. $\qquad\square$

Observing that the gradecast protocol makes only a single call to graded gossip, and no other communication is required, gives us a simple bound on the communication complexity of Protocol 2:

Lemma 3.10 (Communication Complexity). *Let $\mathsf{BCost}(L)$ be the communication complexity for graded gossip with maximum message size of L bits, and let $|r|$ a bound on the encoding of the communication round. Then, the communication complexity of $\pi_{\mathsf{gradecast}}$ (Protocol 2) is at most $\mathsf{BCost}(|r| + L)$.*

3.3 d-Graded f-Threshold-Gossip from d-Graded Gossip

Threshold gossip abstracts the common pattern of requiring $f + 1$ distinct signatures on a gossiped value (where f is a bound on the number of corrupt public keys) in order to ensure that at least one honest party signed the value.

A protocol π is a d-graded f-threshold-gossip protocol if it has the following API and satisfies the following properties: The API is defined by the method $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}, r, S)$, where sid is the session id, r is the round number, and S is a set of values. The output of a party is of the form (sid, r, v, g) for a session id sid, a round number r, a value v, and a grade g.

A d-graded f-threshold-gossip protocol is defined in the d-graded PKI model, and guarantees the following properties with overwhelming probability, for every round r and every session id sid:

- **Graded f-Threshold-Completeness:** For every value v, if $f + 1$ honest parties $\{P_i\}$ called $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}, r, S_i)$ at round r and $v \in \bigcap_{\{i | P_i \text{ is honest}\}} S_i$, then every honest party will output (sid, r, v, d) by the beginning of round $r + 1$.

- **Graded f-Threshold-Soundness:** In every f-faulty execution, for all $r' > r$ and $g > 0$, if by the beginning of round r' an honest party outputs (sid, r, v, g), then there exist a set S and an honest party P_i such that $v \in S$ and P_i called $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}, r, S)$ at round r.

- **Graded Gossip:** For all $r' > r$ and $g > 1$: if, by the beginning of round r', an honest party P_i output (sid, r, v, g), then every honest party P_j will output (sid, r, v, g') by the beginning of round $r' + 1$, and $|g - g'| \leq 1$.

We emphasize that if an honest party receives a value with grade $g = 1$, it is not guaranteed that other honest parties will receive the value at all.

Protocol 3: $\pi_{\text{thresh-gossip}}^{(d,f)}$

The protocol is parameterized by a maximum grade d, a bound f on the corrupt public keys seen by an honest party, and a maximum message size L.

1. When a node calls T-GradeGoss(sid, r, S), invoke the protocol $\pi_{\text{graded-gossip}}^{(d)}$ with GradeGoss($\mathsf{sid}, (r, S)$).

2. For all $\mathsf{vk}, \mathsf{sid}, r, v$, if more than one message of the form $(\mathsf{vk}, \mathsf{sid}, r, v, g)$ were received from $\pi_{\text{graded-gossip}}^{(d)}$, discard all but the one with the maximal grade.

3. For all $s \in \{1, \ldots, d\}$, at round $r + s$:

   (a) Let $V_{\mathsf{sid},r,v,g}$ be the number of different **Valid** messages $(\mathsf{vk}, \mathsf{sid}^*, r^*, S^*, g^*)$ received from $\pi_{\text{graded-gossip}}^{(d)}$ up to round $r + s$ such that
      - $(\mathsf{sid}^*, r^*) = (\mathsf{sid}, r)$,
      - $v \in S^*$,
      - $g^* \geq g$, and
      - no message of the form $(\mathsf{vk}, \mathsf{sid}, r, \bot, g')$ was received from $\pi_{\text{graded-gossip}}^{(d)}$ with $g' \geq g$.

   (b) Let $M_{\mathsf{sid},r,g}$ be the number of different **Malicious** messages $(\mathsf{vk}, \mathsf{sid}^*, r^*, \bot, g^*)$ received from $\pi_{\text{graded-gossip}}^{(d)}$ up to round $r + s$ such that $(\mathsf{sid}^*, r^*) = (\mathsf{sid}, r)$ and $g^* \geq g$. (This counts how many distinct parties were identified as equivocating for (sid, r) with grade at least g.)

   (c) For all v, if $V_{\mathsf{sid},r,v,d+1-s} > 0$, $V_{\mathsf{sid},r,v,d+1-s} + M_{\mathsf{sid},r,d+1-s} > f$, and (sid, r, v, g) has not yet been output for any g, output $(\mathsf{sid}, r, v, d + 1 - s)$.

Theorem 3.11. If $\pi_{\text{graded-gossip}}^{(d)}$ (Protocol 1) is a d-graded gossip protocol, then $\pi_{\text{thresh-gossip}}^{(d,f)}$ (Protocol 3) is a d-graded f-threshold gossip protocol against f-faulty PPT adversaries in the d-graded PKI model.

Proof. For session id sid, round r, value v, and grade g, let $V_{\mathsf{sid},r,v,g}^{(i)}$ and $M_{\mathsf{sid},r,g}^{(i)}$ be the values of $V_{\mathsf{sid},r,v,g}$ and $M_{\mathsf{sid},r,g}$ for party P_i.

Graded f-threshold completeness. Let v be a value, let $\mathcal{P} = \{P_i\}$ be the set of honest parties that called GradeGoss(sid, r, S_i) at round r such that $v \in S_i$, and assume that $|\mathcal{P}| \geq f+1$. By d-graded gossip validity, for every $P_i \in \mathcal{P}$, every honest party will receive $(\mathsf{vk}_i, \mathsf{sid}, S_i, d)$ from $\pi_{\text{graded-gossip}}^{(d)}$ by round $r + 1$. Since $v \in S_i$, and no honest party ever equivocates, it holds that $V_{\mathsf{sid},r,v,d} \geq |\mathcal{P}| \geq f + 1$. Thus, every honest party will output (sid, r, v, d) in round $r + 1$.

Graded f-threshold soundness. Let P_i be an honest party, let v be a value, let $g > 0$ be a grade, and let r be a round number. Then, every key vk for which a message $(\mathsf{vk}, \mathsf{sid}, r, S, g')$ was received by P_i from $\pi_{\text{graded-gossip}}^{(d)}$ at round $r' > r$, for some grade g' and a set S such that $v \in S$, counts at most once in the sum $V_{\mathsf{sid},r,v,g}^{(i)} + M_{\mathsf{sid},r,g}^{(i)}$. Since the execution is f-faulty, if the sum is greater than f then at least one honestly generated key $\mathsf{vk}_{i'}$ belonging to an honest party $P_{i'}$ must have contributed to the sum. By d-graded gossip unforgeability, this can only occur if $P_{i'}$ called GradeGoss(sid, r, S) at round r; that is, $P_{i'}$ must have called T-GradeGoss(sid, r, S) at round r.

Graded gossip. Let P_i be an honest party and suppose that P_i output (sid, r, v, g) at round $r' = r + s$. This can only happen if $g = d + 1 - s$ and, for P_i, at round r' it holds that $V^{(i)}_{\mathsf{sid},r,v,d+1-s} + M^{(i)}_{\mathsf{sid},r,d+1-s} > f$.

By definition, every message received from $\pi^{(d)}_{\mathsf{graded\text{-}gossip}}$ that was counted for either $V^{(i)}_{\mathsf{sid},r,v,d+1-s}$ or $M^{(i)}_{\mathsf{sid},r,d+1-s}$ has the form $(\mathsf{vk}, \mathsf{sid}, r, S, g')$, where $g' \geq g$ and S is either a set of values or $\bot$. By d-*graded gossip consistency*, if P_i received such a message at round r', then every honest party must have received a message $(\mathsf{vk}, \mathsf{sid}, r, S', g'')$ by round $r' + 1$, such that $g'' \geq g' - 1 \geq g - 1$, and $S' \in \{S, \bot\}$. Let P_j be an arbitrary honest party.

By f-*threshold soundness*, for every honest party, and all $r' > r$, it must hold that $V_{\mathsf{sid},r,v,d} \geq 1$ (since at least one honest party called $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}, r, S'')$ for some S'' with $v \in S''$).

If P_j already output (sid, r, v, g') up to round r', then $g' \geq g$. If it did not output (sid, r, v, g') up to round r', then it must hold that $V^{(j)}_{\mathsf{sid},r,v,d-s} + M^{(j)}_{\mathsf{sid},r,d-s} > f$, since every message counted by P_i at round r' would be counted by P_j at round $r' + 1$, for $g' = g - 1 = d - s$.

Thus, every honest party that did not output $(\mathsf{sid}, r, v, \cdot)$ at round r' must output $(\mathsf{sid}, r, v, g - 1)$ at round $r' + 1$. $\qquad\square$

Observing that each threshold gossip invokes $\pi^{(d)}_{\mathsf{graded\text{-}gossip}}$ exactly once for each party gives us the following bound on its communication complexity:

Lemma 3.12 (Graded Threshold-Gossip Communication Complexity). *Let $\mathsf{BCost}(L)$ be the communication complexity for d-graded gossip with maximum message size of L-bits, and $|r|$ a bound on the encoding of the communication round. Then, the communication complexity for* $\pi^{(d,f)}_{\mathsf{thresh\text{-}gossip}}$ *(Protocol 3) is at most* $\mathsf{BCost}(|r| + L)$.

3.4 Graded Crusader Agreement from 4-graded f-Threshold Gossip

In the same way that gradecast is a relaxation of broadcast, *graded crusader agreement* [18] is a relaxation of byzantine agreement. Formally, a protocol π is an R-round graded crusader agreement protocol if it has the following API and satisfies the following properties: The API is defined by the method $\mathtt{GA}(\mathsf{sid}, v)$, where v is a value and sid is an arbitrary unique session id (string). The output of a party is of the form (sid, v, g) for a session id sid, a value v, and a grade $g \in \{0, 1, 2\}$.

An R-round graded crusader agreement guarantees the following properties with overwhelming probability, for every round r and every session id sid:

1. **Validity:** If all honest parties called $\mathtt{GA}(\mathsf{sid}, v)$ at round r, then all honest parties output $(\mathsf{sid}, v, 2)$ at round $r + R$.

2. **Graded agreement:** For every two honest parties P_i and $P_{i'}$: if P_i outputs (sid, v, g) and $P_{i'}$ outputs (sid, v', g'), then

   - $|g - g'| \leq 1$
   - $v' \in \{v, \bot\}$ or $v \in \{v', \bot\}$
   - If $g = 2$ then $v' = v$

For any f-faulty adversary, and assuming the number honest parties is at least $f + 1$, we can construct 4-round graded crusader agreement from 4-graded f-threshold gossip:

> **Protocol 4: π_{GA}**
>
> The protocol is parameterized by a bound f on the corrupt public keys seen by an honest party, and a maximum message size L. It makes use of the protocol $\pi_{\mathsf{thresh\text{-}gossip}}^{(4,f)}$ (Protocol 3) in the 4-graded PKI model.
>
> 1. When a node calls $\mathsf{GA}(\mathsf{sid}, v)$ at round r, it invokes $\pi_{\mathsf{thresh\text{-}gossip}}^{(4,f)}$ with $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}, r, \{v\})$.
>
> 2. At round $r + 4$:
>
> Case 1: If the node received $(\mathsf{sid}, r, v', 4)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ at round $r + 1$, and no other message $(\mathsf{sid}, r, v'', \cdot)$ for $v'' \neq v'$ up to round $r + 4$, output $(\mathsf{sid}, v', 2)$.
>
> Case 2: If the node has not yet produced output, received (sid, r, v', g) from $\pi_{\mathsf{thresh\text{-}gossip}}$ for $g \geq 3$ by round $r + 2$, and no other message $(\mathsf{sid}, r, v'', g'')$ for $v'' \neq v'$ and $g'' \geq 2$ up to round $r + 3$, output $(\mathsf{sid}, v', 1)$.
>
> Case 3: Otherwise output $(\mathsf{sid}, \perp, 0)$

Lemma 3.13. *If the number of honest parties is at least $f + 1$, and $\pi_{\mathsf{thresh\text{-}gossip}}^{(4,f)}$ is a 4-graded f-threshold gossip protocol, then π_{GA} (Protocol 4) is a 4-round graded crusader agreement protocol against f-faulty adversaries in the 4-graded PKI model.*

Proof. We prove each property separately.

Validity. If all honest parties called $\mathsf{GA}(\mathsf{sid}, v)$ at round r, then $f + 1$ honest parties called $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}, r, \{v\})$ at round r. By *4-graded f-threshold-gossip completeness*, every honest party will receive $(\mathsf{sid}, r, v, 4)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ by round $r + 1$. Since no honest party called $\pi_{\mathsf{thresh\text{-}gossip}}$ with any other value, by *4-graded f-threshold-gossip soundness*, no honest party will receive any other message $(\mathsf{sid}, r, v'', g'')$ for $v'' \neq v$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ up to round $r + 4$. Thus, every honest party will output $(\mathsf{sid}, v, 2)$ at round $r + 1$.

Graded agreement. Suppose node P_i output (sid, v, g) and node $P_{i'}$ output (sid, v', g').

Case 1: $g = 2$ or $g' = 2$. Assume without loss of generality that $g = 2$ (the other case is symmetric). In this case, P_i received $(\mathsf{sid}, r, v, 4)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ at round $r + 1$; hence, by the *graded gossip* property of *4-graded f-threshold-gossip*, $P_{i'}$ received $(\mathsf{sid}, r, v, g'')$ by round $r + 2$, for $g'' \geq 3$. Further, $P_{i'}$ cannot have received any $(\mathsf{sid}, r, v''', g''')$ for $v''' \neq v$ and $g''' \geq 2$ by round $r + 3$, since then by the *graded gossip* property P_i would have received $(\mathsf{sid}, r, v''', g'''')$ by round $r + 4$ for $g'''' \geq 1$; hence, would not have output with $g = 2$.

Thus, $P_{i'}$ must output either $(\mathsf{sid}, v, 1)$ or $(\mathsf{sid}, v, 2)$.

Case 2: $g = 1$ and $g' = 1$. In this case P_i received $(\mathsf{sid}, r, v, g'')$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ by round $r + 2$, for $g'' \geq 3$. By the *graded gossip* property of *4-graded f-threshold-gossip*, $P_{i'}$ received $(\mathsf{sid}, r, v, g''')$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ up to round $r + 3$, for $g''' \geq 2$. But this implies $v' = v$, otherwise $P_{i'}$ would not have output $(\mathsf{sid}, v, 1)$.

Case 3: $g = 0$ and $g' \leq 1$, or $g \leq 1$ and $g' = 0$. In this case, graded agreement is trivially satisfied. $\qquad\square$

4 Byzantine Agreement Over Graded PKI

In this section we construct an efficient BA protocol using gradecast and graded threshold gossip. As we showed in Section 3, these can be constructed directly in the graded-PKI model.

18

Our protocol requires 5-graded PKI, which we use to implement 5-graded f-threshold gossip.

4.1 Byzantine Agreement on Sets

In fact, our protocol achieves a stronger notion than traditional BA. We achieve byzantine agreement on *sets*, in the sense that each party begins with a set of values and outputs a set of values. It is guaranteed that: (1) all honest parties output the same set S^{out}, (2) this common set S^{out} is a superset of the intersection of all honest input sets, and (3) S^{out} is a subset of the union of all honest input sets. As discussed in section 1.2, this notion of agreement can be immediately used to boost a graded-PKI into a standard PKI.

Definition 4.1 (Byzantine Agreement on Sets)**.** We say that a protocol Π run by parties $P_1, \ldots, P_n$, where P_i initially holds a set S_i^{in}, achieves *Byzantine agreement on sets* if at the end of the protocol, every honest party outputs a set S_i^{out} such that the following four conditions are satisfied with overwhelming probability (in the security parameter λ):

- **Consistency:** for every two honest parties P_i and P_j, if P_i terminates with set S_i^{out} and P_j terminates with set S_j^{out} then $S_i^{\mathsf{out}} = S_j^{\mathsf{out}}$.

- **Inclusion Validity:** For every value v, if for every honest party P_i it holds that $v \in S_i^{\mathsf{in}}$, then for every honest party P_j it holds that $v \in S_j^{\mathsf{out}}$.

- **Exclusion Validity:** For every value v, if for every honest party P_i it holds that $v \notin S_i^{\mathsf{in}}$, then for every party P_j it holds that $v \notin S_j^{\mathsf{out}}$.

- **Termination:** All honest parties terminate the protocol with overwhelming probability.

Note that we can easily construct standard byzantine agreement in a black-box fashion from byzantine agreement on sets by using a singleton set as the input.

4.2 Protocol Definitions and Overview

The protocol is parameterized by the corruption threshold f and the messages' length L. The maximal grades is set to $d = 5$. For clarity, we denote by $\pi_{\mathsf{thresh\text{-}gossip}}$ the protocol $\pi_{\mathsf{thresh\text{-}gossip}}^{(5,f)}$ and by $\pi_{\mathsf{gradecast}}$ a gradecast protocol for messages on size L in the 5-graded PKI model. Denote by S_i^{in} the input set of P_i. Each execution of the protocol also has an associated session id sid.

p-**Weak Leader Election.** The protocol makes use of a leader-election subprotocol in each iteration. We abstract this as a special case of graded PKI, with an additional protocol $\pi_{\mathsf{quality}}(\mathsf{vk}, \mathsf{sid}, r)$. This protocol receives a public key vk, session id sid, and round index r as input, and returns a *quality score* in the range $[0, 1]$.

The leader in round r of session sid, in the view of party P_i, is the key $(\mathsf{vk}, \mathsf{sid}, r)$ such that $\mathsf{GradeKey}_i(\mathsf{vk}, (\mathsf{sid}, r)) > 0$[4] and for every other key $\mathsf{vk}' \neq \mathsf{vk}$ in the view of P_i, it holds that $\pi_{\mathsf{quality}}(\mathsf{vk}, \mathsf{sid}, r) > \pi_{\mathsf{quality}}(\mathsf{vk}', \mathsf{sid}, r)$. (Note that if a round has a leader, in P_i's view, it is unique, but a round may not have a leader at all.)

We define $\mathtt{IsLeader}^{(i)}(\mathsf{vk}, \mathsf{sid}, r)$ to return 1 if $(\mathsf{vk}, \mathsf{sid}, r)$ is the leader in the view of party P_i, and 0 otherwise.

We assume that the protocol π_{quality} to be a p-Leader Election protocol. A p-Leader Election protocol is defined in the d-graded PKI model, and guarantees the following properties with overwhelming probability, for every round r and every session id sid:

---

[4]Note that the $\mathsf{GradeKey}$ treats the tuple (sid, r) as a session id.

- **Quality Consistency**: For every pair of honest parties P_i and P_j, and every tuple $(\mathsf{vk}, \mathsf{sid}, r)$ such that $\mathsf{GradeKey}_i(\mathsf{vk}, (\mathsf{sid}, r)) > 0$ and $\mathsf{GradeKey}_j(\mathsf{vk}, (\mathsf{sid}, r)) > 0$, if both parties invoke $\pi_{\mathsf{quality}}(\mathsf{vk}, \mathsf{sid}, r)$ (not necessarily at the same time), their outputs are identical.

- **p-Honest Leader Election**: For every sid and round r, with probability at least p, there exists an honestly-generated vk such that for every honest party P_i it holds that $\mathtt{IsLeader}^{(i)}(\mathsf{vk}, \mathsf{sid}, r) = 1$.

These properties ensure that an honest leader is chosen with probability p at each round, and that when this occurs all honest parties will agree on the leader's identity.

Claim 4.2. *For every two honest parties P_i and P_j, if vk and vk' are honest keys, and $\mathtt{IsLeader}^{(i)}(\mathsf{vk}, \mathsf{sid}, r) = \mathtt{IsLeader}^{(j)}(\mathsf{vk}', \mathsf{sid}, r) = 1$, then $\mathsf{vk} = \mathsf{vk}'$ (i.e., if the leader is honest then P_i and P_j agree on the identity of the leader).*

Proof. We assume without loss of generality that every active honest party speaks in each round. Assume, in contradiction, that $\mathsf{vk} \neq \mathsf{vk}'$. By d-graded PKI completeness, both $\mathsf{GradeKey}_j(\mathsf{vk}, (\mathsf{sid}, r)) = d > 0$ and $\mathsf{GradeKey}_i(\mathsf{vk}', (\mathsf{sid}, r)) = d > 0$. Since $(\mathsf{vk}, \mathsf{sid}, r)$ is the leader in P_i's view it holds that $\pi_{\mathsf{quality}}(\mathsf{vk}, \mathsf{sid}, r) > \pi_{\mathsf{quality}}(\mathsf{vk}', \mathsf{sid}, r)$. But for the same reason, using P_j's view, it must be that $\pi_{\mathsf{quality}}(\mathsf{vk}', \mathsf{sid}, r) > \pi_{\mathsf{quality}}(\mathsf{vk}, \mathsf{sid}, r)$. By quality consistency, the outputs of π_{quality} are consistent. Thus, we have reached a contradiction. $\square$

Note that p-Weak Leader Election (also referred to as oblivious leader election [19, 25]) is trivially satisfied by the simple leader election functionality used in many BA protocols, where the leader is always in full consensus, but may be a corrupt party with probability $1 - p$. However, it is strictly weaker: with probability $1 - p$, honest parties are not guaranteed to agree about the leader's identity. This lets it model the Algorand-style leader election protocol based on VRFs, which is a common subprotocol in permissionless protocols. (The protocol is typically be implemented by self-selecting several potential leaders using a VRF, and a posteriori picking the one with the lowest VRF value as the leader; the VRF here translates to the quality score in our definition.)

Iterations and rounds. The protocol proceeds in sequential iterations, beginning with iteration 0; we denote the iteration number by j. Each iteration consists of 7 rounds, except for iteration 0 that consists of 8 rounds. Let $R(j, r) = 7j + r$ be the absolute round number for iteration j, round r (iteration 0 is the first iteration, and the preround is considered round -1; that is, iteration 0 consists of rounds -1 till 6).

Local variables. For each iteration j, party P_i maintains several variables:

- $L_i^{(j)}$ contains the set of values to which P_i is "locked," or $\bot$ if it is not locked. P_i will not commit to any other set in iteration j if it is locked.

- $\mathrm{HardLocked}_i^{(j)} \in \{0, 1\}$ is a boolean variable indicating whether P_i is "hard locked" to the set $L_i^{(j)}$. When P_i is hard locked to a set, it will commit to the set regardless of any proposal in iteration j.

- $T_i^{(j)}$ records the sets from "valid proposals" received in iteration j. As part of the security proof, we show that the output must always be such a set. Since, apart from the proposal and preround, the protocol only requires parties to exactly compare sets, it is ok to replace the encodings of sets in these rounds with the hash of the set.

4.3 Protocol Execution

The byzantine agreement protocol π_{BA} is parameterized by a bound f on the corrupt public keys seen by an honest party, and a maximum message size L. The protocol is defined in the 5-graded PKI model, and makes use of the 5-graded f-faulty threshold gossip protocol $\pi_{\mathsf{thresh\text{-}gossip}}^{(5,f)}$ (Protocol 3) and the 3-round gradecast protocol $\pi_{\mathsf{gradecast}}$ (Protocol 2).

Each protocol iteration proceeds as follows:

Round -1 (preround): (Executed only a single time)

- **If P_i is active:** P_i invokes the 5-graded $\pi_{\mathsf{thresh\text{-}gossip}}$ with $\mathsf{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{preround}, -1, S_i^{\mathsf{in}})$, informing everybody of its set S_i^{in}.

Valid Values. At the beginning of rounds $r \in \{0, 1, 2, 3\}$ of iteration 0, P_i forms a set of valid values $V_i^{(g)}$ by taking the union of all values v such that $(\mathsf{sid}||\textbf{preround}, -1, v, g)$ was received from $\pi_{\mathsf{thresh\text{-}gossip}}$ with $g \geq 5 - r$. (E.g., at round 0 it sets $V_i^{(5)}$ to be the set of values received with grade 5, at round 1 it sets $V_i^{(4)}$ be the set of values received with grades 5 or 4, etc.)

Round 0 (hard-lock):

- If $j > 0$ and by the beginning of round 0 of iteration j, party P_i received from $\pi_{\mathsf{thresh\text{-}gossip}}$ the message $(\mathsf{sid}||\textbf{commit-}(j-1), R(j-1, 5), S^*, g)$ on a set $S^* \in \bigcup_{\ell=0}^{j-1} T_i^{(\ell)}$, for $g \geq 4$, it sets $L_i^{(j)} \leftarrow S^*$ and $\mathrm{HardLocked}_i^{(j)} \leftarrow 1$.
- Otherwise, P_i sets $\mathrm{HardLocked}_i^{(j)} \leftarrow 0$.

Round 1 (soft-lock): (This round is not required in the first iteration)

- If $j > 0$ and up to the beginning of round 1, party P_i received from $\pi_{\mathsf{thresh\text{-}gossip}}$ the message $(\mathsf{sid}||\textbf{commit-}(j-1), R(j-1, 5), S^*, g)$ on a set $S^* \in \bigcup_{\ell=0}^{j-1} T_i^{(\ell)}$, for $g \geq 3$, P_i sets $L_i^{(j)} \leftarrow S^*$.
- Otherwise, P_i sets $L_i^{(j)} \leftarrow \perp$.

Round 2 (propose):

- **If P_i is active:**
  - If $j > 0$ and, up to the beginning of round 2, party P_i received from $\pi_{\mathsf{thresh\text{-}gossip}}$ the message $(\mathsf{sid}||\textbf{commit-}(j-1), R(j-1, 5), S^*, g)$ on a set $S^* \in \bigcup_{\ell=0}^{j-1} T_i^{(\ell)}$, for $g \geq 2$, P_i sets $S \leftarrow S^*$.
  - Otherwise, P_i sets $S \leftarrow V_i^{(4)}$.

  P_i invokes $\pi_{\mathsf{gradecast}}$ with $\mathsf{Gradecast}(\mathsf{sid}||\textbf{proposal-}j, S)$.

Round 3 (wait1): Wait for messages.

Round 4 (wait2): Wait for messages.

Round 5 (commit):

- For every proposal $(\mathsf{vk}, \mathsf{sid}||\textbf{proposal-}j, S, g)$ that P_i received from $\pi_{\mathsf{gradecast}}$ up to the beginning of round 5, such that:

  (a) $g \geq 1$

  (b) $S \subseteq V_i^{(2)}$,

P_i adds S to $T_i^{(j)}$. If no such proposals were received, P_i sets $T_i^{(j)} \leftarrow \perp$.

- **If P_i is active:** If $\mathrm{HardLocked}_i^{(j)} = 1$, then P_i invokes $\pi_{\mathsf{thresh\text{-}gossip}}$ with the message
  $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{commit-}j, R(j,5), \{L_i^{(j)}\})$.

  Otherwise, if P_i received a proposal $(\mathsf{vk}, \mathsf{sid}||\textbf{proposal-}j, S, g)$ such that all of the following conditions are met:

  (c) $\{S\} = T_i^{(j)}$ (i.e., S is the only set in $T_i^{(j)}$).

  (d) $\mathtt{IsLeader}^{(i)}(\mathsf{vk}, \mathsf{sid}||\textbf{proposal-}j, R(j,2)) = 1$

  (e) $g = 2$

  (f) $S \subseteq V_i^{(3)}$,

  (g) Either $V_i^{(5)} \subseteq S$ or $(\mathsf{sid}||\textbf{commit-}(j-1), R(j-1,5), S, g)$ was received from $\pi_{\mathsf{thresh\text{-}gossip}}$ for $g \geq 1$ by the beginning of round 5 of iteration j.

  (h) $L_i^{(j)} = \perp$ or $L_i^{(j)} = S$,

  then P_i invokes $\pi_{\mathsf{thresh\text{-}gossip}}$ with $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{commit-}j, R(j,5), \{S\})$.

Round 6 (notify):

Case 1: If P_i received $(\mathsf{sid}||\textbf{notify-}(j-1), R(j-1,6), S, 5)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ at round 0 of iteration j, such that $S \in \bigcup_{\ell=0}^{j-1} T_i^{(\ell)}$:

  - it outputs S
  - if P_i is active, it invokes $\pi_{\mathsf{thresh\text{-}gossip}}$ with $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{notify-}j, R(j,6), \{S\})$
  - P_i terminates. (P_i continues to participate in the $\pi_{\mathsf{thresh\text{-}gossip}}$ protocol for one more iteration.)

Case 2: Otherwise, **if P_i is active:** if it has received $(\mathsf{sid}||\textbf{commit-}j, R(j,5), S, 5)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ and $S \in T_i^{(j)}$, it calls $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{notify-}j, R(j,6), \{S\})$.

Iteration-end: At the end of round 6, P_i sets $j \leftarrow j+1$ and continues from round 0.

4.4 Security Analysis

We proceed to state the security guaranteed of the protocol.

Theorem 4.3. *Let $0 < p < 1$ be a constant, and assume that $\pi_{\mathsf{gradecast}}$ (Protocol 2) is a 3-round gradecast protocol, that π_{quality} is a p-leader-election protocol, and that $\pi_{\mathsf{thresh\text{-}gossip}}^{(5,f)}$ (Protocol 3) is a 5-graded f-threshold gossip protocol against f-faulty PPT adversaries in the 5-graded PKI model.*

Then, the protocol π_{BA} (Section 4.3) achieves byzantine agreement on sets against f-faulty PPT adversaries in the 5-graded PKI model.

Recall that a 5-graded PKI also defines a 3-graded PKI as required by our 3-round gradecast protocol, see Section 3.1.

The rest of this subsection is devoted to proving Theorem 4.3. In Section 4.4.1 we prove consistency, in Section 4.4.2 we prove termination, and in Section 4.4.3 we prove validity.

4.4.1 Consistency

Lemma 4.4 (Commit Consistency). *For all $j \geq 0$, if an honest party P_i called $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{commit-}j, R(j,5), \{S\})$ at iteration j, then for every honest party $P_{i'}$:*

  1. $S \in \bigcup_{\ell=0}^{j} T_{i'}^{(\ell)}$ and

2. if $P_{i'}$ called `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), \{S'\}$) in iteration j, then $S' = S$.

Proof. Consider the event that the adversary succeeds in forging an honest sender's signature; by the security of the signature scheme this happens with negligible probability. We proceed by conditioning on the complementary event and prove the lemma by induction on j.

Base case: For $j = 0$, the only way P_i can call `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), \{S\}$) is if P_i received ($\mathsf{vk}, \mathsf{sid}||\mathbf{proposal\text{-}}j, S, 2$) at round 3 of iteration j that satisfies conditions 5a to 5h.

By *gradecast weak consistency*, $P_{i'}$ must have received ($\mathsf{vk}, \mathsf{sid}||\mathbf{proposal\text{-}}j, S, g$) by round 4 of iteration j, for $g \geq 1$. By condition 5f, $S \subseteq V_i^{(3)}$; hence, for every honest $P_{i'}$ it holds that $S \subseteq V_{i'}^{(2)}$, and thus condition 5b is satisfied for $P_{i'}$. Since conditions 5a to 5b are all satisfied, $P_{i'}$ must have added S to $T_{i'}^{(j)}$ in round 5. This also implies that if $P_{i'}$ sent a commit message for S' in iteration j, then $S' = S$ (otherwise we would have $S', S \in T_{i'}^{(j)}$, hence $|T_{i'}^{(j)}| > 1$).

Induction Step: Assume the induction hypothesis is true up to some $j \geq 0$. Suppose P_i called `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), \{S\}$) at iteration $j+1$. There are exactly two options for how this could happen:

Case 1: Hard-locked: $\mathrm{HardLocked}_i^{(j+1)} = 1$.

In this case, P_i must have received ($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S, 4$) from $\pi_{\mathsf{thresh\text{-}gossip}}$ up to round 0 of iteration $j+1$. Thus, by the *graded-gossip* property of $\pi_{\mathsf{thresh\text{-}gossip}}$, $P_{i'}$ must have received ($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S, g$) from $\pi_{\mathsf{thresh\text{-}gossip}}$ up to round 1 of iteration $j+1$, for $g \geq 3$. Moreover, by *graded f-threshold-soundness*, at least one honest party must have called `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), \{S\}$) at iteration j. By the induction hypothesis, this means $S \in \bigcup_{\ell=0}^{j} T_{i'}^{(\ell)}$ and that no honest party called `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}j, \{S'\}$) for $S' \neq S$ in iteration j. Therefore, by *graded f-threshold-soundness*, $P_{i'}$ cannot have received ($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S', *$) for $S' \neq S$, thus, $P_{i'}$ will set $L_{i'}^{(j+1)} \leftarrow S$ in round 1 of iteration $j+1$.

To see that if $P_{i'}$ called `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), \{S'\}$), then $S' = S$, suppose it did. Then either:

- $\mathrm{HardLocked}_{i'}^{(j+1)} = 1$.
  In this case $S' = S$ because $P_{i'}$ calls `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), \{L_{i'}^{(j+1)}\}$), and we have just shown that $L_{i'}^{(j+1)} = S$, or

- $\mathrm{HardLocked}_{i'}^{(j+1)} = 0$.
  In this case $S' = S$ because otherwise condition 5h would not be satisfied.

Case 2: Good Proposal: $\mathrm{HardLocked}_i^{(j+1)} = 0$, and P_i received ($\mathsf{vk}, \mathsf{sid}||\mathbf{proposal\text{-}}j, S, 2$) from $\pi_{\mathsf{gradecast}}$ by round 5 of iteration $j+1$ that satisfies conditions 5a to 5h.

By the same arguments as for the base case, it must be that $P_{i'}$ received the proposal, it satisfies conditions 5a to 5b, and thus $P_{i'}$ will add S to $T_{i'}^{(j+1)}$.

To see that if $P_{i'}$ called `T-GradeGoss`($\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), \{S'\}$), then $S' = S$, suppose it did. Then either

- $\mathrm{HardLocked}_{i'}^{(j+1)} = 0$.
  In this case $S' = S$ because otherwise condition 5c would not be satisfied, or

- $\mathrm{HardLocked}_{i'}^{(j+1)} = 1$ and $L_{i'}^{(j+1)} = S'$.
  In this case, $P_{i'}$ must have received ($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S', 4$) from $\pi_{\mathsf{thresh\text{-}gossip}}$ up to round 0 of iteration $j+1$. By the consistent-gossip property of $\pi_{\mathsf{thresh\text{-}gossip}}$, P_i must have received ($\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S', 3$) from $\pi_{\mathsf{thresh\text{-}gossip}}$ up to round 1 of iteration $j+1$. By the graded f-threshold-soundness property, at least one honest

party must have called $(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S')$, and therefore, by the induction hypothesis, $S' \in \bigcup_{\ell=0}^{j} T_i^{(\ell)}$ and no honest party called $(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S'')$ for $S'' \neq S'$. By graded f-threshold-soundness, P_i could not have received $(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S'', *)$ for $S'' \neq S'$. Thus, in round 1 of iteration $j + 1$, P_i must have set $L_i^{(j+1)} \leftarrow S'$. By condition 5h, this implies $S' = S$. $\qquad\square$

Lemma 4.5 (Notify Implies Termination). *Let $j \geq 0$ be the first iteration at which an honest party P_i calls* `T-GradeGoss`$(\mathsf{sid}||\mathbf{notify\text{-}}j, R(j, 6), \{S\})$ *for some set S. Then, every honest party will terminate by the end of iteration $j + 2$ with output S.*

Proof. Since no honest party called `T-GradeGoss`$(\mathsf{sid}||\mathbf{notify\text{-}}j'...)$ for $j' < j$, by f-threshold-soundness P_i cannot have received $(\mathsf{sid}||\mathbf{notify\text{-}}(j-1), R(j-1, 6), S, 5)$, hence must be in case 2 of round 6 at iteration j. Since P_i called `T-GradeGoss`$(\mathsf{sid}||\mathbf{notify\text{-}}j, R(j, 6), \{S\})$ in case 2 of round 6 of iteration j, P_i must have received $(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S, 5)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ by the beginning of round 6 of iteration j.

By Lemma 4.4, this implies that for every honest party $P_{i'}$, $S \in \bigcup_{\ell=0}^{j} T_{i'}^{(\ell)}$ and for all $S' \neq S$, no honest party called `T-GradeGoss`$(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), \{S'\})$. By f-threshold-soundness, this implies that for all $S' \neq S$, no honest party received $(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S', *)$ (in any round).

Thus, no honest party called `T-GradeGoss`$(\mathsf{sid}||\mathbf{notify\text{-}}j, R(j, 6), \{S'\})$, and—again by f-threshold soundness—no honest party can receive $(\mathsf{sid}||\mathbf{notify\text{-}}j, R(j, 6), S', *)$, hence no honest party can terminate with output other than S in iteration $j + 1$.

Since P_i received $(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S, 5)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ by the beginning of round 6 of iteration j, then by the consistent-gossip property, $P_{i'}$ must have received $(\mathsf{sid}||\mathbf{commit\text{-}}j, R(j, 5), S, g)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ by the beginning of the next round, for $g \geq 4$. That is, in round 0 of iteration $j + 1$. Moreover, since $S \in \bigcup_{\ell=0}^{j} T_{i'}^{(\ell)}$, it follows that $P_{i'}$, and every honest party, must set $L_{i'}^{(j+1)} \leftarrow S$ and $\mathrm{HardLocked}_{i'}^{(j+1)} \leftarrow 1$ at round 0 of iteration $j + 1$.

Since $\mathrm{HardLocked}_{i'}^{(j+1)} \leftarrow 1$, $P_{i'}$ must have called `T-GradeGoss`$(\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), \{S\})$ in iteration $j + 1$. This holds for every honest party, and no honest party could have terminated yet (honest parties only terminate one iteration after first receiving a **notify**, and the first is in iteration $j + 1$). Thus for every $S' \neq S$,

- no honest party called `T-GradeGoss`$(\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), S')$,

- hence by graded f-threshold-soundness, no honest party receives $(\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), S', *)$,

- thus no honest party calls `T-GradeGoss`$(\mathsf{sid}||\mathbf{notify\text{-}}(j+1), R(j+1, 6), S')$,

- hence no honest party receives $(\mathsf{sid}||\mathbf{notify\text{-}}(j+1), R(j+1, 6), S', *)$

Moreover, all honest parties called `T-GradeGoss`$(\mathsf{sid}||\mathbf{commit\text{-}}(j+1), \{S\})$ in round 5 of iteration $j + 1$. Since there are at least $f + 1$ active honest parties, by the f-threshold-completeness property of $\pi_{\mathsf{thresh\text{-}gossip}}$, all honest parties will receive $(\mathsf{sid}||\mathbf{commit\text{-}}(j+1), R(j+1, 5), S, 5)$ by the beginning of round 6 of iteration $j + 1$. Therefore, all honest parties will call `T-GradeGoss`$(\mathsf{sid}||\mathbf{notify\text{-}}(j+1), R(j+1, 6), \{S\})$ at round 6 (whether due to case 1 or case 2). By graded f-threshold-completeness all honest parties will receive $(\mathsf{sid}||\mathbf{notify\text{-}}(j+1), R(j+1, 6), S, 5)$ by the end of iteration $j + 1$ and any party that was not in case 1 of round 6 in iteration $j + 1$. Honest parties that previously received $(\mathsf{sid}||\mathbf{notify\text{-}}(j), R(j, 6), S, 5)$, will terminate with output S at iteration $j + 1$; the remaining honest parties will terminate with output S in iteration $j + 2$ (due to case 1 of round 6). $\qquad\square$

Lemma 4.6 (Consistency). *Under the assumptions in Theorem 4.3, Protocol π_{BA} satisfies consistency.*

Proof. Honest parties only terminate if there exists j and a set S such that they received $(\mathsf{sid}\|\textbf{notify-}j, R(j, 6), S, 5)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$. Thus, by f-threshold-soundness, if any honest party terminated with output S, there must exist a minimal iteration j^* such that an honest party called $\texttt{T-GradeGoss}(\mathsf{sid}\|\textbf{notify-}j^*, R(j^*, 6), \{S\})$ at iteration j^*. By Lemma 4.5, all honest parties will terminate with output S in this case. $\qquad\square$

4.4.2 Termination

Claim 4.7 (Honest Proposal Lock Consistency). *For all $j \geq 0$, if the leader P_i in iteration j is honest and proposes a set S, then for every honest party $P_{i'}$ the proposal satisfies condition 5h in round 5.*

Proof. If $L_{i'}^{(j)} = \perp$ the condition is trivially satisfied. Otherwise, suppose $P_{i'}$ set $L_{i'}^{(j)} \leftarrow S'$ in round 0 or round 1 of iteration j. In this case, $P_{i'}$ must have received $(\mathsf{sid}\|\textbf{commit-}(j-1), R(j-1, 5), S', g)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ by the beginning of round 1 of iteration j, where $g \geq 3$.

By consistent-gossip, every honest party, including P_i must have received $(\mathsf{sid}\|\textbf{commit-}(j-1), R(j-1, 5), S', g')$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ by the beginning of round 2 of iteration j, with $g' \geq 2$. By f-threshold-soundness, at least one honest party called $\texttt{T-GradeGoss}(\mathsf{sid}\|\textbf{commit-}(j-1), R(j-1, 5), S')$ in iteration $j-1$, and by Lemma 4.4, $S' \in \bigcup_{\ell=0}^{j-1} T_i^{(\ell)}$ and every other honest commit in iteration $j-1$ was for S'. Therefore, by f-threshold-soundness, P_i could not have received $(\mathsf{sid}\|\textbf{commit-}(j-1), R(j-1, 5), S'', *)$ for $S'' \neq S'$.

Thus, it must be that P_i called $\texttt{Gradecast}(\mathsf{sid}\|\textbf{proposal-}j, S')$; that is, $S' = S$, and condition 5h is satisfied for $P_{i'}$. $\qquad\square$

Claim 4.8. *For every iteration j, if any honest party receives $(\mathsf{sid}\|\textbf{commit-}j, R(j, 5), S, g)$ for $g > 0$ then*

- *there exists an honest party P_i such that $V_i^{(5)} \subseteq S$ and*

- *for every $j' > j$ and every honest party $P_{i'}$, it holds that S satisfies conditions 5b and 5f for $P_{i'}$ at iteration j'*

Proof. The proof is by induction on j.
Base case ($j = 0$): If any honest party receives $(\mathsf{sid}\|\textbf{commit-}j, R(j, 5), S, g)$, then by f-threshold-soundness there must be an honest party P_i that called $\texttt{T-GradeGoss}(\mathsf{sid}\|\textbf{commit-}j, R(j, 5), S)$ in iteration j. Since this is the first iteration, $\texttt{HardLocked}_i^{(j)} = 0$, S must satisfy conditions 5b and 5f for P_i, hence $S \subseteq V_i^{(3)}$. Moreover, since no previous commit messages could have been sent, in order to satisfy *condition 5g* it must be that $V_i^{(5)} \subseteq S$. Finally, for every honest $P_{i'}$ and $j' > j$, it holds that $V_i^{(3)} \subseteq V_{i'}^{(2)}$, and therefore $S \subseteq V_{i'}^{(3)}$ and satisfies conditions 5b and 5f for $P_{i'}$ at iteration j'
Induction Step: Assume the hypothesis holds up to some $j \geq 0$. If an honest party receives $(\mathsf{sid}\|\textbf{commit-}(j+1), R(j+1, 5), S, g)$, then by f-threshold-soundness there must be an honest party P_x that called $\texttt{T-Gossip}(\mathsf{sid}\|\textbf{commit-}(j+1), R(j+1, 5), S)$ in iteration $j+1$.

Case 1: Some honest party received $(\mathsf{sid}\|\textbf{commit-}j^*, R(j^*, 5), S, g)$ for some $j^* \leq j$ up to iteration $j+1$. In this case, by the induction hypothesis there exists P_i that we need and conditions are satisfied for $j' = j+1 > j^*$. Otherwise,

Case 2: No honest party received $(\mathsf{sid}\|\textbf{commit-}j^*, R(j^*, 5), S, g)$ for $j^* \leq j$ and $g > 0$ up to iteration $j+1$. In this case, the proof is essentially identical to the base case, since it must be that $\texttt{HardLocked}_x^{(j+1)} = 0$ and satisfying *condition 5g* requires $V_x^{(5)} \subseteq S$. $\qquad\square$

Lemma 4.9 (Honest-Leader Termination). *For all $j \geq 0$, if the leader P_i in iteration j is honest and called $\mathit{Gradecast}(\mathsf{sid}||\textbf{proposal-}j, S)$, then all honest parties will terminate with output S by the end of iteration $j + 1$.*

Proof. Let $P_{i'}$ be an honest party. By 3-round gradecast validity, if P_i is honest and called $\mathtt{Gradecast}(\mathsf{sid}||\textbf{proposal-}j, S)$ at round 2, then $P_{i'}$ will receive $(\mathsf{vk}_i, \mathsf{sid}||\textbf{proposal-}j, S, 2)$ by round 5. Thus, it will satisfy conditions 5a and 5e for $P_{i'}$.

Case 1: $S = V_i^{(4)}$. In this case,

(a) Since $V_i^{(4)} \subseteq V_{i'}^{(3)}$ for every two honest parties $P_i, P_{i'}$, S must satisfy conditions 5b and 5f.

(b) Since $V_{i'}^{(5)} \subseteq V_i^{(4)}$ it follows that $V_{i'}^{(5)} \subseteq S$ hence it will satisfy condition 5g.

Case 2: $S \neq V_i^{(4)}$. In this case, it must be that P_i received $(\mathsf{sid}||\textbf{commit-}(j-1), R(j-1, 5), S, g)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ for $g \geq 2$ by round 2. Thus.

(a) By Claim 4.8 S will satisfy conditions 5b and 5f.

(b) By graded threshold soundness, P_i' will receive $(\mathsf{sid}||\textbf{commit-}(j-1), R(j-1, 5), S, g)$ from $\pi_{\mathsf{thresh\text{-}gossip}}$ for $g \geq 1$ by round 3, thus S will satisfy condition 5g due to the second part of the condition.

So far, we have shown that conditions 5a, 5b and 5e to 5g must be satisfied. Since conditions 5a and 5b are satisfied, $P_{i'}$ will add S to $T_j^{(i')}$. Since P_i is an honest leader, $\mathtt{IsLeader}^{(i')}(\mathsf{vk}_i, \mathsf{sid}||\textbf{proposal-}j, R(j, 2)) = 1$, hence condition 5d is satisfied. Since it is honest, it will not send any additional proposals in this iteration, and by Claim 4.2, it is the only leader in this round, hence no additional proposals will be sent that satisfy condition 5d for $P_{i'}$. Thus, no other set $S' \neq S$ can be added to $T_j^{(i')}$, and therefore condition 5c will also be satisfied for $P_{i'}$.

By Claim 4.7, S must satisfy condition 5h for $P_{i'}$. Thus, there are two options:

Case 1: $\mathrm{HardLocked}_{i'}^{(j)} = 1$. In this case, condition 5h implies that $L_{i'}^{(j)} = S$, hence $P_{i'}$ will call $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{commit-}j, R(j, 5), S)$.

Case 2: $\mathrm{HardLocked}_{i'}^{(j)} = 0$. In this case, since all the conditions are satisfied, $P_{i'}$ will call $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{commit-}j, R(j, 5), S)$.

Since this is true for every honest party, every honest party will call $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{commit-}j, R(j, 5), S)$ in round 5 of iteration j, and by threshold-completeness every honest party will receive $(\mathsf{sid}||\textbf{commit-}j, R(j, 5), S, 5)$ by round 6 of iteration j, and will therefore call $\mathtt{T\text{-}GradeGoss}(\mathsf{sid}||\textbf{notify-}j, R(j, 6), S)$. Since there are at least $f + 1$ active honest parties, by f-threshold-completeness, every honest party will receive $(\mathsf{sid}||\textbf{notify-}j, R(j, 6), S, 5)$ by the end of iteration j and terminate with output S at iteration $j + 1$. $\square$

Let X_j denote the event "An honest leader is elected in iteration j."

Lemma 4.10 (Expected Termination). *If the events X_j are independent, and $\Pr[X_j] \geq p$, then Protocol π_{BA} is expected to terminate in at most $7 \cdot (1 + 1/p)$ rounds.*

Proof. By Lemma 4.9, for all $j \geq 0$ the event X_j implies that all honest parties terminate in iteration $j + 1$. The number of iterations is thus bounded by a geometrically-distributed random variable with parameter p, plus an extra iteration. Hence, the expected number of iterations is at most $1 + 1/p$. Every iteration has 7 rounds (in the first iteration, there is a preround, but no soft-lock round, every subsequent iteration lacks a preround but has a soft-lock round instead), so the expected number of rounds is bounded by $7 \cdot (1 + 1/p)$. $\square$

4.4.3 Validity

Lemma 4.11 (Validity)**.** *Under the assumptions in Theorem 4.3, Protocol π_{BA} satisfies Inclusion and Exclusion Validity.*

Proof. Suppose an honest party terminated with output S, it must have received $(\mathsf{sid}\|\textbf{notify-}j, S)$ from $\mathcal{F}_{\mathsf{T\text{-}Goss}}$. By f-threshold-soundness, some honest party $P_{i'}$ must have called $\texttt{T-Gossip}(\mathsf{sid}\|\textbf{notify-}j, R(j, 6), S, g)$, which means it must have received $(\mathsf{sid}\|\textbf{commit-}j, R(j, 5), S, g')$ from $\pi_{\mathsf{thresh\text{-}gossip}}$. By Claim 4.8, there exists an honest party P_i such that

- $V_i^{(5)} \subseteq S$ and

- for every $j' > j$ and every honest party $P_{i'}$, it holds that S satisfies conditions 5b and 5f for $P_{i'}$ at iteration j'.

Inclusion validity. Let v be a value in the intersection of all honest parties' inputs. In the preround, v is in every set $S_{i'}$ for which honest parties call $\texttt{T-GradeGoss}(\mathsf{sid}\|\textbf{preround}, -1, S_{i'})$. Since there are at least $f + 1$ active honest parties, by f-threshold-completeness, $\pi_{\mathsf{thresh\text{-}gossip}}$ will output $(\mathsf{sid}\|\textbf{preround}, -1, v, 5)$ to every honest party, and hence $v \in V_i^{(5)}$, which implies $v \in S$.

Exclusion validity. Let v be a value that is not in the union of honest parties' inputs. In the preround, v is not in any set $S_{i'}$ for which honest parties call $\texttt{T-GradeGoss}(\mathsf{sid}\|\textbf{preround}, S_{i'})$, so by f-threshold-soundness no honest party will receive $(\mathsf{sid}\|\textbf{preround}, -1, v, *)$. Thus, for all i and $g > 0$, it holds that $v \notin V_i^{(g)}$. Since S satisfies condition 5b for P_i, it holds that $S \subseteq V_i^{(2)}$, and thus $v \notin S$. $\qquad\square$

4.5 Communication Complexity

Let $|\mathcal{S}|$ be a bound on the encoding of the union of the input sets, $|r|$ a bound on the encoding of the round number and λ_H a bound on the output of a collision-resistant hash. Denote $\mathsf{BCost}(L)$ the cost of 5-graded gossip for an L-bit message. Let $\mathsf{TCost}(L)$ denote the communication cost of 5-graded f-threshold gossiping a set with encoding size L, and $\mathsf{GCost}(L)$ the communication cost of gradecasting an L-bit message.

Theorem 4.12. *Let n be an upper bound on the number of* participating *public keys with non-zero grades (i.e., those eligible to generate messages in the protocol), n' an upper bound on the number of participating propose-round keys and p a lower bound on the probability that an iteration elects an honest leader. The communication complexity of the BA protocol, such that the probability of choosing an honest leader in each iteration is at least p, is at most*

$$n \cdot \mathsf{TCost}(|\mathcal{S}|) + (1 + 1/p) \cdot \big(2n \cdot \mathsf{TCost}(\lambda_H) + n' \cdot \mathsf{GCost}(|\mathcal{S}|)\big).$$

Proof. In the preround, every party threshold-gossips its input set. Since our threshold-gossip protocol's communication cost depends only on the encoding size of the input set, the cost of this round is at most $n \cdot \mathsf{TCost}(|\mathcal{S}|)$. In rounds 5 and 6, every party threshold-gossips a "pointer" to a set. This can be a collision-resistant hash of the set (which has length λ_H). Thus, the total cost of both rounds is at most $2n \cdot \mathsf{TCost}(\lambda_H)$. In round 2, at most n' keys are eligible to gradecast a set, so the cost of this round is at most $n' \cdot \mathsf{GCost}(|\mathcal{S}|)$.

Thus, in the first iteration, the total cost is bounded by

$$n \cdot \mathsf{TCost}(|\mathcal{S}|) + 2n \cdot \mathsf{TCost}(\lambda_H) + n' \cdot \mathsf{GCost}(|\mathcal{S}|)$$

and each subsequent iteration by $2n \cdot \mathsf{TCost}(\lambda_H) + n' \cdot \mathsf{GCost}(|\mathcal{S}|)$. Since the expected number of iterations is $1 + 1/p$, the total expected cost is thus bounded by

$$\mathrm{Exp}(\mathrm{Cost}) = n \cdot \mathsf{TCost}(|\mathcal{S}|) + (1 + 1/p) \cdot \big(2n \cdot \mathsf{TCost}(\lambda_H) + n' \cdot \mathsf{GCost}(|\mathcal{S}|)\big). \qquad\square$$

Denote $\lambda_{\text{overhead}} = \lambda + \lambda_{pk} + \lambda_{sig} + |r|$. Using our implementations over graded gossip, over a graph $G = (V, E)$, the communication complexity above translates into

$$
\begin{aligned}
\text{Exp}(\text{Cost}) &= n \cdot \mathsf{BCost}(|r| + |\mathcal{S}|) + (1 + 1/p) \cdot \big(2n \cdot \mathsf{BCost}(|r| + \lambda_H) + n' \cdot \mathsf{BCost}(|r| + |\mathcal{S}|)\big) \\
&= (n + (1 + 1/p) \cdot n') \cdot \mathsf{BCost}(|r| + |\mathcal{S}|) + 2n \cdot (1 + 1/p) \cdot \mathsf{BCost}(|r| + \lambda_H) \\
&= (n + (1 + 1/p) \cdot n') \cdot 2|E|(\lambda_{\text{overhead}} + |\mathcal{S}|) + 2n \cdot (1 + 1/p) \cdot 2|E|(\lambda_{\text{overhead}} + \lambda_H) \\
&= 2|E| \big((n + (1 + 1/p) \cdot n') \cdot (\lambda_{\text{overhead}} + |\mathcal{S}|) + 2n \cdot (1 + 1/p) \cdot (\lambda_{\text{overhead}} + \lambda_H)\big) \\
&= 2|E| \big((n + (1 + 1/p) \cdot (2n + n')) \cdot \lambda_{\text{overhead}} + (n + (1 + 1/p) \cdot n') \cdot |\mathcal{S}| + 2n \cdot (1 + 1/p) \cdot \lambda_H\big) .
\end{aligned}
$$

4.5.1 Concrete Complexity Comparison

To demonstrate to what extent the communication complexity is improved in practice, we compare the concrete complexity of our new protocol with several state-of-the-art protocols, using reasonably-chosen concrete parameters. We measure the complexity in terms of the total worst-case amount gossiped per peer, as the underlying gossip functionality has a similar cost in all the cases.

We compare only to protocols that are highly scalable in the permissionless setting, and security assuming simple honest majority. In all of these protocols, a small committee is sampled at random from the entire population (for the purposes of communication complexity, it does not matter if the same committee is chosen for every round). To guarantee an honest majority in the committee with probability at least $1 - 2^{-40}$ (a commonly used statistical security parameter), we need $n \approx 800$ when the population has a $2/3$ honest majority (for a smaller honest majority, the committee size would be larger, tilting the comparison even further in our favor).

Reasonable values for the security parameters are $\lambda = |r| = 64$, $\lambda_H = \lambda_{pk} = 256$ and $\lambda_{sig} = 512$ (e.g., using SHA256 for hashing and ed25519 for signatures). Thus, $\lambda_{\text{overhead}} = 896$.

For the leader-election (which is required in all the protocols), we assume we are using Algorand-style leader election using VRFs, and rely only on a simple honest majority. The committee for the propose round only needs to guarantee that at least one party is elected with overwhelming probability (rather than an honest majority on the committee); thus, we can use $n' \ll n$. Using the same 2^{-40} statistical security parameter, it is enough to use $n' = 30$.

Since at least half of the population is honest, the probability that the minimal VRF value will be honest is at least $1/2$, thus it is reasonable to use $p = 1/2$.

For an apples-to-apples comparison, we will use our protocol to agree on 256-bit scalars rather than sets; thus, $|\mathcal{S}| = 256$.

Our protocol. Plugging these values into the computation of Theorem 4.12, we have

$$
\begin{aligned}
\text{Exp}(\text{Cost}) &\leq 2|E| \big((n + (1 + 1/p) \cdot (2n + n')) \cdot \lambda_{\text{overhead}} \\
&\qquad + (n + (1 + 1/p) \cdot n') \cdot |\mathcal{S}| + 2n \cdot (1 + 1/p) \cdot \lambda_H\big) \\
&\leq 2|E| \big((7n + 90) \cdot \lambda_{\text{overhead}} + (n + 90) \cdot |\mathcal{S}| + 6n \cdot \lambda_H\big) \\
&= 2|E|(7808n + 80640) + (n + 90) \cdot |\mathcal{S}|) \\
&= |E| \cdot (15616n + 161280 + 2(n + 90) \cdot |\mathcal{S}|).
\end{aligned}
$$

Thus, per communication link, over the entire protocol:

- Less than 1.6MiB for signatures, keys and hashes (independent of the input size).

- Less than 256 bytes per bit of input.

With a 256-bit input, the total communication per link is still less than 1.6MiB.

ADDNR19 without threshold signatures [2]. Counting only the **commit** round, in which each honest party sends a message containing $f + 1 > n/2$ signatures, hence of size $(n/2) \cdot (\lambda_{pk} + \lambda_{sig}) + |\mathcal{S}|$. Assuming only 2 iterations (i.e., $p = 1/2$, just the signatures in this round alone are responsible for $n^2 \cdot (\lambda_{pk} + \lambda_{sig}) =$ bits per communication link. For our parameters, this is more than 58MiB per link.

Abraham et al. [2] propose reducing the communication by replacing the certificate of $f + 1$ signatures with a single share of a threshold signature scheme (whose size is on the order of a single signature). This does indeed reduce the communication complexity, but requires trusted setup (which is not usually feasible in permissionless settings).

ACDNPRS23 [1]. The synchronous protocol for $f < n/2$ is very similar to [2]. In this case, in expectation the protocol will need at least one Commit and one Status round, each of which requires every party to send a "certificate," of size at least $(n/2) \cdot (\lambda_{pk} + \lambda_{sig})$ (counting only public keys and signatures). Thus, the total communication for our parameters would also be more than 58MiB.

Micali-Vaikuntanathan agreement for $f < n/2$ [30]. The protocol uses one "large" round, in which each party broadcasts $n/2$ signatures. In [30], the protocol assumes a single sender (and implements broadcast). For apples-to-apples comparison, assuming we use random leader selection to select a sender, the expected number of iterations is at least 2, so the total cost would be $n^2 \cdot (\lambda_{pk} + \lambda_{sig})$. This is also more than 58MiB per communication link.

References

[1] I. Abraham, T. H. Chan, D. Dolev, K. Nayak, R. Pass, L. Ren, and E. Shi. Communication complexity of byzantine agreement, revisited. *Distributed Computing*, 36(1):3–28, 2023. doi:10.1007/s00446-022-00428-8.

[2] I. Abraham, S. Devadas, D. Dolev, K. Nayak, and L. Ren. Synchronous Byzantine agreement with expected O(1) rounds, expected o(n^2) communication, and optimal resilience. In *FC 2019*, pages 320–334, 2019.

[3] M. Andrychowicz and S. Dziembowski. Pow-based distributed cryptography with no trusted setup. In *CRYPTO 2015, Part II*, pages 379–399, 2015.

[4] C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas. Bitcoin as a transaction ledger: A composable treatment. In *CRYPTO 2017, Part I*, pages 324–356, 2017.

[5] M. Ben-Or, D. Dolev, and E. N. Hoch. Brief announcement: Simple gradecast based algorithms. In *DISC 2010*, volume 6343, pages 194–197, 2010.

[6] M. Ben-Or, D. Dolev, and E. N. Hoch. Simple gradecast based algorithms. *CoRR*, abs/1007.1049, 2010. URL: http://arxiv.org/abs/1007.1049, arXiv:1007.1049.

[7] P. Berman, J. A. Garay, and K. J. Perry. *Bit Optimal Distributed Consensus*, page 313–321. Plenum Press, USA, 1992.

[8] E. Buchman, R. Guerraoui, J. Komatovic, Z. Milosevic, D. Seredinschi, and J. Widder. Revisiting tendermint: Design tradeoffs, accountability, and practical use. In *DSN 2022*, pages 11–14, 2022.

[9] E. Buchman, J. Kwon, and Z. Milosevic. The latest gossip on BFT consensus. *CoRR*, abs/1807.04938, 2018. URL: http://arxiv.org/abs/1807.04938.

[10] S. Chaudhuri. More choices allow more faults: Set consensus problems in totally asynchronous systems. *Inf. Comput.*, 105(1):132–158, 1993.

[11] J. Chen and S. Micali. Algorand: A secure and efficient distributed ledger. *Theor. Comput. Sci.*, 777:155–183, 2019.

[12] R. Cohen, I. Haitner, E. Omri, and L. Rotem. From fairness to full security in multiparty computation. *J. Cryptol.*, 35(1):4, 2022.

[13] Concordium Foundation. Concordium white paper: An introduction to concordium. https://assets.website-files.com/64f060f3fc95f9d2081781db/64f1f9fa65a5aea515e799ef_Concordium-White-Paper-v1.6.pdf, 2023. Version 1.6, January 2023.

[14] S. Coretti, A. Kiayias, C. Moore, and A. Russell. The generals' scuttlebutt: Byzantine-resilient gossip protocols. In H. Yin, A. Stavrou, C. Cremers, and E. Shi, editors, *CCS 2022*, pages 595–608. ACM, 2022.

[15] P. Das, L. Eckey, S. Faust, J. Loss, and M. Maitra. Round efficient byzantine agreement from VDFs. *IACR Cryptol. ePrint Arch.*, page 823, 2022.

[16] C. Decker and R. Wattenhofer. Information propagation in the bitcoin network. In *IEEE P2P 2013*, pages 1–10, 2013.

[17] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart, and D. B. Terry. Epidemic algorithms for replicated database maintenance. In *PODC*, pages 1–12, 1987.

[18] D. Dolev. The byzantine generals strike again. *J. Algorithms*, 3(1):14–30, 1982. doi: 10.1016/0196-6774(82)90004-9.

[19] P. Feldman and S. Micali. An optimal probabilistic algorithm for synchronous byzantine agreement. In *Automata, Languages and Programming, 16th International Colloquium, ICALP89, Stresa, Italy, July 11-15, 1989, Proceedings*, pages 341–378, 1989.

[20] P. Feldman and S. Micali. An optimal probabilistic protocol for synchronous Byzantine agreement. *SIAM J. Comput.*, 26(4):873–933, 1997.

[21] M. J. Fischer, N. A. Lynch, and M. Merritt. Easy impossibility proofs for distributed consensus problems. In *PODC*, pages 59–70, 1985.

[22] J. A. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol: Analysis and applications. In *EUROCRYPT 2015, Part II*, pages 281–310, 2015.

[23] J. A. Garay, A. Kiayias, R. M. Ostrovsky, G. Panagiotakos, and V. Zikas. Resource-restricted cryptography: Revisiting MPC bounds in the proof-of-work era. In *EURO-CRYPT 2020, Part II*, pages 129–158, 2020.

[24] R. M. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *FOCS 2000*, pages 565–574, 2000.

[25] J. Katz and C. Koo. On expected constant-round protocols for byzantine agreement. In *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings*, pages 445–462, 2006.

[26] L. Lamport, R. E. Shostak, and M. C. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[27] C. Liu-Zhang, C. Matt, U. Maurer, G. Rito, and S. E. Thomsen. Practical provably secure flooding for blockchains. In *ASIACRYPT 2022, Part I*, pages 774–805, 2022.

[28] C. Matt, J. B. Nielsen, and S. E. Thomsen. Formalizing delayed adaptive corruptions and the security of flooding networks. In *CRYPTO 2022, Part II*, pages 400–430, 2022.

[29] S. Micali. ALGORAND: the efficient and democratic ledger. *CoRR*, abs/1607.01341, 2016. URL: http://arxiv.org/abs/1607.01341, arXiv:1607.01341.

[30] S. Micali and V. Vaikuntanathan. Optimal and player-replaceable consensus with an honest majority. Technical Report MIT-CSAIL-TR-2017-004, MIT, 2017. URL: https://dspace.mit.edu/bitstream/handle/1721.1/107927/MIT-CSAIL-TR-2017-004.pdf.

[31] A. Momose and L. Ren. Optimal communication complexity of authenticated byzantine agreement. In *DISC 2021*, pages 32:1–32:16, 2021.

[32] M. C. Pease, R. E. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.

[33] G. Tsimos, J. Loss, and C. Papamanthou. Gossiping for communication-efficient broadcast. In *CRYPTO 2022, Part III*, pages 439–469, 2022.

[34] D. Vyzovitis, Y. Napora, D. McCormick, D. Dias, and Y. Psaras. Gossipsub: Attack-resilient message propagation in the filecoin and ETH2.0 networks. *CoRR*, abs/2007.02754, 2020. URL: https://arxiv.org/abs/2007.02754, arXiv:2007.02754.