

The Generals' Scuttlebutt: Byzantine-Resilient Gossip Protocols

Sandro Coretti
IOG
sandro.coretti@iohk.io

Cristopher Moore
Santa Fe Institute
moore@santafe.edu

Aggelos Kiayias
University of Edinburgh
IOG
Aggelos.Kiayias@ed.ac.uk

Alexander Russell
University of Connecticut
IOG
acr@uconn.edu

ABSTRACT

One of the most successful applications of peer-to-peer communication networks is in the context of blockchain protocols, which—in Satoshi Nakamoto's own words—rely on the “nature of information being easy to spread and hard to stifle.” Significant efforts were invested in the last decade into analyzing the security of these protocols, and invariably the security arguments known for longest-chain Nakamoto-style consensus use an idealization of this tenet. Unfortunately, the real-world implementations of peer-to-peer gossip-style networks used by blockchain protocols rely on a number of ad-hoc attack mitigation strategies that leave a glaring gap between the idealized communication layer assumed in formal security arguments for blockchains and the real world, where a wide array of attacks have been showcased.

In this work we bridge this gap by presenting a Byzantine-resilient network layer for blockchain protocols. For the first time we quantify the problem of network-layer attacks in the context of blockchain security models, and we develop a design that thwarts *resource-restricted* adversaries. Importantly, we focus on the proof-of-stake setting due to its vulnerability to Denial-of-Service (DoS) attacks stemming from the well-known deficiency (compared to the proof-of-work setting) known as *nothing at stake*.

We present a Byzantine-resilient gossip protocol, and we analyze it in the Universal Composition framework. In order to prove security, we show novel results on expander properties of random graphs. Importantly, our gossip protocol can be based on any given bilateral functionality that determines a desired interaction between two “adjacent” peers in the networking layer and demonstrates how it is possible to use application-layer information to make the networking-layer resilient to attacks. Despite the seeming circularity, we demonstrate how to prove the security of a Nakamoto-style longest-chain protocol given our gossip networking functionality, and hence, we demonstrate constructively how it is possible to obtain provable security across protocol layers, given only bare-bone point-to-point networking, majority of honest stake, and a verifiable random function.

CCS CONCEPTS

• Security and privacy → Cryptography; Distributed systems security.

KEYWORDS

proof of stake, blockchains, gossiping, Byzantine-resilience, expander graphs, universal composability

ACM Reference Format:

Sandro Coretti, Aggelos Kiayias, Cristopher Moore, and Alexander Russell. 2022. The Generals' Scuttlebutt: Byzantine-Resilient Gossip Protocols. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22)*, November 7–11, 2022, Los Angeles, CA, USA. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3548606.3560638>

1 INTRODUCTION

1.1 Gossip Protocols and Byzantine Attackers

Gossip protocols. Gossip protocols [12, 18] provide an efficient mechanism to distribute information to a large set of parties. The key feature of such algorithms is their peer-to-peer operation that load balances the effort of information propagation in a way that individual nodes are only investing a modicum of effort when contributing to the delivery of a message network-wide.

As communication infrastructure for multiparty cryptography, gossip protocols have recently found wide-spread application in the context of blockchain protocols, notably with the introduction of the Bitcoin blockchain [26]. Among other things, gossip protocols are used by blockchain participants to diffuse newly found blocks. In the words of Nakamoto, blockchain protocols rely on the “nature of information being easy to spread and hard to stifle,” underscoring the relevance of gossip as the underlying communication layer.¹

The “security” of gossip protocols. Deploying gossip protocols as the communication layer of blockchain protocols adds a crucial new dimension to their design: their “security.”

At the very least, gossip protocols underlying blockchains must deal with the fact that the resources (e.g., network bandwidth, computation time and space) of each peer are limited, and exhausting them will lead to denial-of-service (DoS) attacks. As such, the aforementioned consideration of keeping the complexity for each peer



This work is licensed under a Creative Commons Attribution International 4.0 License.

¹Note that the security guarantees of gossip protocols are weaker than Byzantine-resilient Broadcast aka the Byzantine Generals problem [23] because they do not guarantee any kind of agreement on or consistent ordering of messages gossiped among parties.

small (sublinear, preferably constant, in the number of total participants) is therefore not only relevant for efficiency but also essential for security.

For this reason, gossip protocols currently used in practice incorporate an array of (typically ad-hoc) measures to protect information propagation against DoS attacks. In the setting of blockchain protocols in particular, the most salient feature of such DoS mitigations is the fact that the adversary is resource-bounded (e.g., has limited hashing power, in the context of the Bitcoin protocol, or stake, in the context of proof-of-stake protocols), and peers can exploit this to manage network-wide message propagation. Techniques include rejecting previously seen proof-of-work messages and skipping content downloads that will not result in a local state update (such as skipping the download of a block’s contents when the block header indicates that it cannot be adopted based on the local state of the client).

It should be stressed that such measures are far from perfect, as exemplified by a number of attacks that have been described, including eclipse attacks [17] and routing attacks [2]. Moreover, the proof-of-stake setting poses additional difficulties stemming from the possibility of reusing keys to issue numerous conflicting messages [27] and the fact that the whole stakeholder set should be at hand for proof-of-stake verification to work, in sharp contrast to proof of work, which only needs the current difficulty level.

However, the security issues extend far beyond just the per-peer complexity. Most crucially, the design, setup, and maintenance of the overlay used for gossiping must be such that it resists *Byzantine attackers*—who make participants actively and maliciously deviate from the prescribed protocol—that command a large number of peers in the network. Despite all of the above shortcomings, in practice the information-propagation guarantees of the deployed networking layers of blockchain protocols are generally postulated to be sufficient for the higher-level protocol to maintain its security and correctness.

On the theory side, blockchain consensus protocols—be they Nakamoto-style (e.g., [11, 26]) or inspired by Byzantine fault-tolerant computing (e.g., [7]), based on proof of work (PoW) or proof of stake (PoS)—all crucially rely on the reliable and timely delivery of protocol messages (blocks, votes, etc.) to achieve liveness and many also to be safe. However, while the consensus layers of all these protocols have received considerable attention, with a number of them achieving provable security against Byzantine attackers, the design and security of the network layer are usually an afterthought at best.

Specifically, all previous formal security analyses of PoS protocols (e.g., [7, 10, 11]) use (over-)idealized message-passing abstractions that essentially promise that honest messages are distributed undisturbed to all honest parties within a reasonable delay window and ignore the fact that these abstractions must be implemented in the real world.

More broadly, there exists surprisingly little published work that considers the problem of extending Byzantine resilience to the communication layer of blockchains or gossip protocols in general.

The above state of affairs highlights serious shortcomings of the approaches taken both in theory and in practice and also suggests a significant gap between the two. Given that gossip is a critical piece of the protocol stack for any permissionless distributed-ledger

protocol, the lack of a thorough, formal security treatment of its properties is a critical deficiency in the understanding of the security of these protocols. This the main motivation behind the present paper.

1.2 Our Results

This work takes a systematic and principled approach to alleviating the issues explained above and provides a novel design for Byzantine-resilient gossiping in the context of PoS blockchain protocols. The results are presented in the Universal Composability framework [5].

A Byzantine-resilient network layer for blockchains. The first main contribution of this work is a protocol for “synchronizing chains” globally among participants of PoS blockchain consensus.² The protocol is designed to work over a standard, Internet-like network with (bounded-delay) message passing.

Crucially, the security of the network layer is based on the same assumption as that of the consensus layer, namely that the majority of stake in the system is controlled by honest parties. This may seem circular at first, as the proper operation of the network layer is conditioned on agreement on the stake distribution. The way to break this “cycle” is as follows: Commonly, Nakamoto-style PoS blockchains anyway split the execution of the consensus protocol into epochs and use the stake distribution SD_{i-1} at the end of an epoch $i-1$ as a basis for consensus in epoch $i+1$ (under the assumption of bounded stake drift during epochs). The same approach can be taken for the network layer, i.e., SD_{i-1} underpins the execution of the network layer in epoch $i+1$. Specifically, in the new protocol parties use a verifiable random function (VRF) to, based on SD_{i-1} , create a stake-weighted random-graph overlay in which the degree of each party is constant in the number of participants. The use of VRFs allows parties to reject connection requests from participants that are not supposed to be in their neighbor set.³

Because the node degrees in the network protocol are constant, it is easy for an (adaptive) attacker to isolate a (bounded) fraction of the stake by corrupting all neighbors of parties making up said stake. To that end, edges in the graph have an expiration time, at which point replacements are sampled. In addition to helping parties recover from eclipses, this also allows the overlay to gradually adapt to changing stake distributions.

It may seem tempting to now simply perform run-of-the-mill block gossiping over the above overlay. Such an approach, however, seems to be unsuitable (at least) for Nakamoto blockchains: For example, it is an impractical and insecure design to ask parties to keep all received blocks around, as many of them could be adversarial in the PoS setting (an attacker may—in principle—generate as many blocks as they like, in particular as slot leader). Therefore, with block gossip, a party should delete blocks unless they extend its current local chain. During a fork event, however, the

²This work focuses on synchronizing chains; similar approaches can be taken to synchronize other protocol messages, such as votes, transactions, etc.

³Non-Nakamoto blockchains sometimes do not have the notion of an epoch built into the consensus protocol, e.g., because they *finalize* blocks and use the stake distribution of the finalized block as a basis for agreeing on the next block. However, for the purpose of equipping them with the network design presented here, such a notion can easily be added to them. Moreover, note that it seems that any reasonable design of a network layer for PoS protocols needs to assume bounded stake drift.

party may require previously deleted blocks for which the gossip is “over.” Similar issues apply if a party misses blocks, e.g., due to an eclipse. In general, which blocks are needed by a particular party to synchronize with the system is highly dependent on the party’s local state.

Therefore, stateless solutions, i.e., solutions in which neighbors do not share state, do not appear to be a good fit and will not yield DoS-resilient and scalable blockchain protocols in practice. A more suitable approach is to have each pair of neighbors run a *bilateral chain synchronization (chain sync) protocol*, which allows them to keep each other informed about their locally preferred chains. When a party discovers a better chain in one of these chain-sync instances or when it produces a better one, it informs all of its current chain-sync instances of the new chain.

This work abstracts this bilateral chain sync as a functionality $F_{\text{bilateral}}$. The actual implementation of such a protocol is outside the scope of and irrelevant for this work.⁴ It is merely important to note that chain sync is *stateful* and instances thereof take time to set up in the real world (establishing the connection, initial synchronization of blocks, etc.). Hence, bilateral chain sync is intended to run between two parties for an extended (albeit bounded) amount of time. These facts are captured by $F_{\text{bilateral}}$ in that there is an initial synchronization delay δ_{init} , which may be much larger than the delay δ_{sync} occurring once a chain sync instance has been set up.

Finally, observe that in the blockchain context particularly (but also in general), the attacker must be prevented from interrupting the propagation of any specific message; otherwise, the attacker can, e.g., prevent honest chain updates from spreading through the network. Unfortunately, with an efficient gossip protocol, in which nodes have constant degrees, the (adaptive) attacker can simply corrupt all neighbors of a block leader and thereby halt propagation. It is therefore unavoidable to consider a model in which corruption requests by the attacker only take effect after a certain amount of time.

The ideal global chain-sync functionality. The security guarantees of the protocol above are captured by a functionality F_{sync} , the second main contribution of this work. Similarly to the network functionalities assumed in prior work, F_{sync} provides global chain “propagation” within some time bound Δ_{sync} . However, there are several important differences stemming from the fact F_{sync} is implemented and not assumed. The two most crucial ones are the following:

- F_{sync} allows the attacker to “eclipse” parties and exclude them from the provided guarantees, as long as the fraction of eclipsed stake does not exceed a certain bound. Note that the adversary is allowed to be “mobile” w.r.t. which parties are eclipsed, i.e., every party can potentially be eclipsed at some point during the execution of the protocol. Note that there is an eclipse delay $\xi_{\text{ecl}} \geq \Delta_{\text{sync}}$ in F_{sync} . This ensures that the attacker cannot stop the propagation of specific chains.

- Due to the use of chain sync, F_{sync} ’s guarantees are slightly weaker than those offered by the assumed network functionalities in prior work: instead of a particular chain C “propagating” through the network within Δ , F_{sync} may also instead deliver different chains C' that are not worse (i.e., equal length or longer) than C .

Security proof. In order to show that the new network protocol securely realizes ideal functionality F_{sync} , this work derives a new result on expander graphs: Consider the stake-weighted random graph formed by parties in the protocol. Then, even after removing all corrupted nodes from the graph, leaving behind some fraction α of honest stake, there exists a subgraph of honest parties, the *backbone*, corresponding to at least an $(\alpha - \beta)$ -fraction of the total stake, for some β , and this backbone is an expander graph (with overwhelming probability). Most importantly, the result holds even if the attacker chooses which nodes to remove *with full knowledge of the entire graph*. The expander property guarantees that the diameter in the backbone is small, and therefore timely chain “propagation” is possible therein.

Fully Byzantine-resilient PoS. As a final contribution, this work demonstrates how to utilize the synchronization functionality in the context of a proof-of-stake protocol. Specifically, [11] is used to illustrate the result. First, note that the original analysis is insufficient: despite the fact that the networking model of [11] allows a Byzantine adversary controlling a minority of stake, F_{sync} permits a “mobile” eclipsing strategy that would deplete the adversarial budget of any straightforward reduction to the adversary of [11]. To circumvent this issue, a revamped analysis is presented showing that adaptive eclipsing does not disturb the forkable-string analysis of [11], which can be recovered to demonstrate that the protocol remains secure even against an adversary that exploits the enhanced capabilities of the adversarial interface of F_{sync} . This gives rise to a new analysis of [11] in an adversarial setting where in addition to party corruption, some degree of message suppression is also permitted.

1.3 Related Work

The use of gossip or epidemic algorithms in the context of distributed systems was put forth in [13] and explored at length both in networking systems [20] but also from theoretical angles [19]. The study of Byzantine fault tolerance in a network of a bounded degree was initiated in [14] and further refined in [29], where it is shown that if the adversary is bounded by $O(n/\log n)$ or $O(n)$, respectively, there exist graphs of bounded degree that facilitate broadcast.

In terms of total communication, it is known that the communication complexity of Byzantine broadcast is $\Omega(n^2)$, and assuming some type of delay in the corruption model is necessary to break the communication complexity barrier to sublinear [1].

In the context of peer-to-peer networking for blockchains, a “structured” approach in the organization of the peer-to-peer network can be used to reduce communication complexity further, but at the expense of adaptive security [28].

Our protocols, viewed from the lens of multiparty computation, exhibit a “communication locality,” which has also been studied

⁴Secure and practical implementations of chain sync exist, e.g., [9] and [8, Section 3.7].

more broadly in the context of general secure multiparty computation [6]. We also remark that message suppression as an enhanced adversarial capability was also studied in the general secure MPC setting in the context of “omission corruptions” [30].

Finally, concurrently and independently of our work, [24, 25] studied the related problem of stateless flooding amongst a set of participants, also motivated in part by the blockchain setting. Their flooding protocol has peers connecting to a bounded-size, randomly chosen neighborhood for each message transmission. This can be seen as creating a separate random graph for every message. However, as mentioned above, stateless flooding is unsuitable in practice for blockchain protocols. Moreover, their protocol does not use any mechanism to resource bound the adversary (e.g., as a VRF) and hence there is no mitigation that can prevent honest parties from being flooded with messages by the adversary.

2 PRELIMINARIES AND NOTATION

UC security. Protocols are described and proven secure in the Universal Composability (UC) framework [5]. In UC, the security of a particular protocol is captured by comparing a real-world execution of the protocol to an ideal-world execution in which the protocol is replaced by an ideal functionality. In rough terms, a protocol π securely realizes a functionality F , if for every real-world attacker \mathcal{A} attacking π , there exists an ideal-world simulator \mathcal{S} attacking F such that the real and ideal experiments become indistinguishable to all environments \mathcal{Z} .⁵

A protocol (in the real world) may itself make calls to so-called hybrid functionalities. These hybrid functionalities may serve to either model assumptions (e.g., F_{net} below) or are themselves realized by protocols. The UC framework guarantees that the security of a protocol is maintained when hybrids are replaced by the protocols that realize them.

Round structure. All functionalities/protocols proceed in rounds (not made explicit) and are assumed to have access to the current time, denoted T . Generally, the round structure is such that parties first use a fetch-type command to retrieve information, followed by a send/set-type command to distribute information.

Attacker and corruption delay. The attacker considered is polynomially bounded and may corrupt parties, thereby learning their internal state, and make them deviate from the prescribed protocol arbitrarily. The attacker is adaptive, i.e., it may choose whom to corrupt on the fly during the execution of the protocol and based on all the information observed. However, there is a corruption delay of ξ_{corr} , i.e., a corruption request issued by the attacker takes effect after a delay of ξ_{corr} rounds only.

Underlying network. This work considers parties P with so-called *relays*, identified by their IP addresses IP. Having the actual node—holding the key material—firewalled by relays is common practice as a first line of defense against intrusion attacks.⁶

Communication between relays is modeled by functionality F_{net} (cf. Figure 1), which captures a simple, Internet-like network. Parties

⁵The environment both acts as distinguisher and controls the attacker/simulator as well as the inputs to the parties.

⁶Note, however, that for the purposes of this work, the attacker corrupts parties, at which point it gains control over (all of) their relays.

Functionality F_{net}

Parameters: δ_{net} : maximum delay.

Variables: The functionality keeps track of the following variables, initialized to the following values: array $\text{IPs}[P] := \emptyset$: set of IPs owned by parties P ; array $\mathcal{M}[\text{mid}]$: message records $(\text{IP}, \text{IP}', t, m)$, indexed by message IDs (MIDs) mid.

IP addresses: Upon (GETIP) from P : Output (GETIP, P) to \mathcal{S} and ask \mathcal{S} for a unique address IP. Add IP to $\text{IPs}[P]$ and output (GETIP, IP) to P .

Malicious IPs Registration: Upon (REGIP, IP) from \mathcal{S} : if IP is unique, add it to $\text{IPs}[\mathcal{S}]$.

Message send: Upon (SEND, IP, IP', m) from P with $\text{IP} \in \text{IPs}[P]$: Pick a unique mid and store $\mathcal{M}[\text{mid}] := (\text{IP}, \text{IP}', T, m)$. Output (SEND, mid, IP, IP', m) to \mathcal{S} .

Malicious message send: Upon (SEND, IP, IP', m) from \mathcal{S} where $\text{IP} \in \text{IPs}[\mathcal{S}]$ or P with $\text{IP} \in \text{IPs}[P]$ is corrupted: Pick a unique mid and store $\mathcal{M}[\text{mid}] := (\text{IP}, \text{IP}', \infty, m)$. Output (SEND, mid) to \mathcal{S} .

Fetching: Upon (FETCH, IP) from P with $\text{IP} \in \text{IPs}[P]$:

- (1) Output (FETCH, P , IP) to \mathcal{S} and ask \mathcal{S} for a set M of MIDs.
- (2) Let M' be the set of MIDs corresponding to records $(\cdot, \text{IP}, t, \cdot)$ with $T - t \geq \delta_{\text{net}}$.
- (3) Let $\hat{M} := M[M \cup M']$ and set $\mathcal{M}[\hat{M} \cup M'] := \perp$.
- (4) Output (FETCH, \hat{M}') to P , where \hat{M}' is the set \hat{M} with the time stamp removed in each tuple.

Figure 1: Network/Internet functionality F_{net} .

can obtain (unique) IP addresses for their relays. F_{net} guarantees bounded-delay message transmission between any two relays. The attacker sees all messages sent and may send messages on behalf of any relay IP owned by a corrupted party; it is, however, not permitted to interfere with message transmission between honest relays (beyond inducing a bounded delay).

Master index. A master index MI is made up of:

- the network directory ND, which consists of tuples (P, IP, v) , where P is a party ID, IP is an IP address, and v is a VRF public key (see below);
- the stake distribution SD, which consists of tuples (P, α) , where $\alpha \in (0, 1]$ denotes P 's stake fraction; and
- a seed value R.

A master index is *valid* if (a) the same parties appear in ND and SD, and each party appears at most once, (b) values IP and v appear only once in ND, and (c) the values α in SD sum up to 1. All master indices appearing in this work are tacitly assumed to be valid.

Observe that the MI format defined above restricts each party to having only one IP address. This choice was made for simplicity; the definition of master indices as well as all protocols and functionalities can easily be adapted to allow multiple IP addresses per party, thereby modeling the fact that in practice parties often have multiple relays.

Verifiable random functions. A verifiable random function (VRF) is a cryptographic primitive that allows a party P to create key pair

Functionality F_{Vrf}

Variables: The functionality keeps track of the following variables, initialized to the values below:

- (1) an array $\text{Keys}[P] := \emptyset$: set of keys owned by P ;
- (2) an array $T[v, x] := \perp$, where v is a key and x a domain value: pair (y, S) , where y is a range value and S a set of proofs π ;
- (3) a set $E := \emptyset$: contains triples (v, x, y) to keep track of all VRF evaluations.

Keys: Upon (GETKEY) from P : Output (GETKEY, P) to S and ask S for a unique key v . Add v to $\text{Keys}[P]$ and output (GETKEY, v) to P .

Register Keys: Upon (REGKEY, v) from S : if v is unique, add it to $\text{Keys}[S]$.

Evaluation: Upon (EVAL, v, x) from P with $v \in \text{Keys}[P]$: Output (EVAL, P, v, x) to S , and upon obtaining (EVAL, π) from S :

- (1) If π is not unique, exit the procedure.
- (2) If $T[v, x]$ is undefined, pick y uniformly at random from the range and set $S := \emptyset$; otherwise, let $(y, S) := T[v, x]$.
- (3) Set $T[v, x] := (y, S \cup \{\pi\})$ and add (v, x, y) to E .
- (4) Output (EVAL, y, π) to P .

Malicious evaluation: Upon (EVAL, v, x) from S :

- *Case 1:* There exists an *uncorrupted* P with $v \in \text{Keys}[P]$: If $(y, S) := T[v, x]$ is defined, return (EVAL, y) to S ; otherwise, do nothing.
- *Case 2:* There exists a *corrupted* P with $v \in \text{Keys}[P]$ or $v \in \text{Keys}[S]$: If $T[v, x]$ is not defined, pick y uniformly at random from the range, let $S := \emptyset$ and set $T[v, x] := (y, S)$; otherwise, let $(y, S) := T[v, x]$. Return (EVAL, y) to S .
- *Else:* Do nothing.

Verification: Upon $(\text{VERIFY}, v, x, y, \pi)$ from any ITI: Send $(\text{VERIFY}, v, x, y, \pi)$ to S , and upon receiving (VERIFY, ϕ) from S :

- If $v \in \text{Keys}[\cdot]$ and $T[v, x] = (y, S)$ is defined:
 - (1) If $\pi \in S$, set $f := 1$.
 - (2) Else, if $\phi = 1$ and π is unique—i.e., if for all $(v', x') \neq (v, x)$ with $T[v', x'] = (\cdot, S)$, $\pi \notin S$ —set $T[v, x'] := (y, S \cup \{\pi\})$ and $f := 1$.
 - (3) Else, set $f := 0$.
- Else, set $f := 0$.

Output (VERIFY, f) to P .

Adversarial leakage: Upon (LEAK) from S : return (LEAK, E) to S .

Figure 2: VRF functionality F_{Vrf} .

consisting of a secret evaluation key and a public verification key such that: (a) with the secret key, P can evaluate the VRF at any input x , obtaining a random-looking output y and a proof π ; (b) given the public key of P , anyone is able to verify, using π , that y is indeed the output corresponding to x . Importantly, even a malicious P cannot bias the output of the VRF on any particular input x (for any fixed public key).

The above guarantees are abstracted by an idealized functionality F_{Vrf} (cf. Figure 2). The main commands offered by F_{Vrf} are (i) (GETKEY) , answered by (GETKEY, v) for an (adversarially chosen)

idealized public key v , (ii) (EVAL, v, x) , answered by (EVAL, y, π) , and (iii) $(\text{VERIFY}, v, x, y, \pi)$, answered by (VERIFY, ϕ) with $\phi \in \{0, 1\}$ indicating whether π is a valid proof for the input/output pair (x, y) under public key v .

3 BYZANTINE-RESILIENT NETWORKING

3.1 Overview

One of the main contributions of this work is to define and realize a *synchronization functionality* F_{Sync} that can be used by participants of proof-of-stake (PoS) consensus layers to globally synchronize their blockchains. A most crucial feature of F_{Sync} is that it is realizable *under the same assumptions* as the consensus protocol that builds on top of it.

In rough terms, the idea underpinning the protocol realizing F_{Sync} is as follows: Each consensus participant P samples $\Theta_P \cdot d$ neighbors in a stake-based fashion based on the output of a verifiable random function (VRF); d is a small constant, and Θ_P is a multiplier that depends on the amount of stake of P itself. The use of a VRF guarantees that parties can verify whether they were indeed supposed to be chosen as neighbors by other participants.

Once the neighbors are sampled, chains are synchronized globally by each party engaging in bilateral chain synchronization with all its neighbors. Due to the fact that the resulting communication overlay is essentially a *random graph*, using expander theory, one can show that even with all adversarial nodes removed, there remains a *large connected backbone* with *small diameter*. This enables *timely synchronization (TS)*, which is crucial to the security of many consensus protocols.

Of course, with a constant number of neighbors, it is unavoidable that some parties end up being connected solely to corrupted parties, which means that they are *eclipsed*. Consequently, F_{Sync} will have to grant eclipsing power to the adversary and can only guarantee TS to non-eclipsed parties.

Functionality F_{Sync} will be realized (cf. Section 3.4) from a functionality $F_{\text{Bilateral}}$ (cf. Section 3.2), which models bilateral chain synchronization, as well as from network functionality F_{Net} (cf. Section 2) and (standard) functionality F_{Vrf} (cf. Section 2).

3.2 Bilateral Chain Synchronization

Functionality $F_{\text{Bilateral}}$ (cf. Figure 3) models bilateral chain synchronization between two parties A and B . In $F_{\text{Bilateral}}$, each party has a *local chain*, and the functionality allows them to learn their counterparty's chain as it evolves. A party's local chain evolves by use of the SETC command, under the restriction that local chains are only replaced by "better" chains as determined by a strict partial order $\text{prefer}(\cdot, \cdot)$, where $\text{prefer}(C, C')$ (also denoted by $C > C'$) if and only if C is strictly preferable to C' ⁷; prefer is a parameter of the functionality.

Both parties are informed about changes to their counterparty's local chain with a delay of at most δ_{sync} . However, initially, i.e., until both parties have used SETC at least once (thereby indicating that they are ready to start chain synchronization), the delay may be up to δ_{init} ; this models the fact that in reality chain synchronization can

⁷Commonly, C is strictly preferable to C' if it is longer.

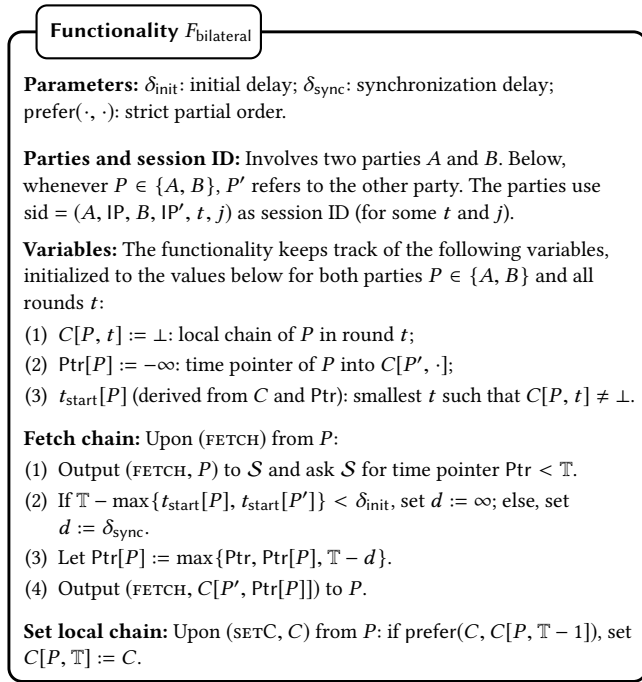


Figure 3: Bilateral chain-sync functionality $F_{\text{bilateral}}$.

take significant time between two parties who have just established a connection.

This work leaves the exact mechanism employed to realize $F_{\text{bilateral}}$ open and simply assumes $F_{\text{bilateral}}$ as a hybrid. In general, the approach to realizing $F_{\text{bilateral}}$ is along the following lines: Upon establishing a connection, two parties first determine the point at which their local chains diverge,⁸⁹ subsequently, they inform each other about changes to their local chains.

In practice, highly optimized implementations of $F_{\text{bilateral}}$ are used in order to improve synchronization time even further. For an example of such an implementation, see [9] and [8, Section 3.7].

3.3 Synchronization Functionality

Functionality F_{sync} , whose realization is the main focus of this work, allows parties to synchronize their chains with the rest of the participants. It is parametrized by an initial delay Δ_{init} , a synchronization delay Δ_{sync} , an eclipse delay $\xi_{\text{ecl}} \geq \Delta_{\text{sync}}$, a “lookback” parameter $\mu \geq \Delta_{\text{sync}}$, as well as an upper bound λ on the amount of eclipsed honest stake.

⁸For the sake of offering a concrete description: In order to figure out the common prefix, party A sends block hashes of suffixes of length 1, 2, 4, 8, . . . to party B until B finds that one of these hashes corresponds to a block on their chain. Thereafter, the exact point of divergence is located by binary search. In the worst case, this takes $O(\delta_{\text{net}} \cdot \log k)$, where k is the common prefix parameter. Note that this is independent of the length of the parties’ chains. Furthermore, on average, in a longest chain protocol, due to the exponential decay in the probability of divergence (as a function of the size of the divergence), synchronization time is much shorter (typically a small constant) between up-to-date peers.

⁹In a BFT-style blockchain like Algorand [7], one local chain is normally a prefix of the other, which simplifies this step.

Note that parties have to agree on the round number in which they start using F_{sync} . For convenience, in this section, initialization of F_{sync} begins in round $-\Delta_{\text{init}}$, and parties actually start using F_{sync} in round 0.

IP and key management. Since F_{sync} is realized from F_{net} and F_{vrf} , interfaces for IP and key registration are also provided by F_{sync} . The reason these are not abstracted away is that IPs and VRF keys must be known by the (higher-level) consensus protocol (e.g., to generate the genesis block).¹⁰

Master index and setup. As previously mentioned, in the synchronization protocol parties will sample neighbors based on their stake. The stake distribution, however, is an object that emanates from the consensus layer, which itself relies on the synchronization functionality. This apparent cycle is broken as follows: F_{sync} expects each party to input its view on the *master index* (cf. Section 2) in the beginning via command SETUP , and F_{sync} only provides guarantees if (a) all honest parties (i) agree on the master index and (ii) input chains originating from the same genesis block, and (b) all honest parties are *represented* in the master index $\text{MI} = (\text{ND}, \text{SD}, \text{R})$ they input; a party is represented in MI if one of its IP addresses IP and one of its keys v appear in ND, i.e., $(P, \text{IP}, v) \in \text{ND}$.

When all of the above conditions are satisfied, F_{sync} has *valid setup*. If the setup is not valid, F_{sync} shuts down. Note that it is the responsibility of the consensus protocol (i.e., the environment in the diffusion context) to ensure that the setup is valid. This is commonly achieved via the genesis block (which is assumed to be available to all parties) initially and later on based on the epoch-wise consensus properties of the protocol. More details on this relationship are provided in Section 5.

Eclipsing. Due to the fact that in practically feasible and scalable protocols for realizing F_{sync} , each party is connected only to a small subset of all participants, it is unavoidable that the adversary may eclipse certain honest parties by (adaptively) corrupting all of their neighbors. Consequently, F_{sync} offers an ECLIPSE command by which the attacker can exclude any honest party from the TS guarantees (see below).

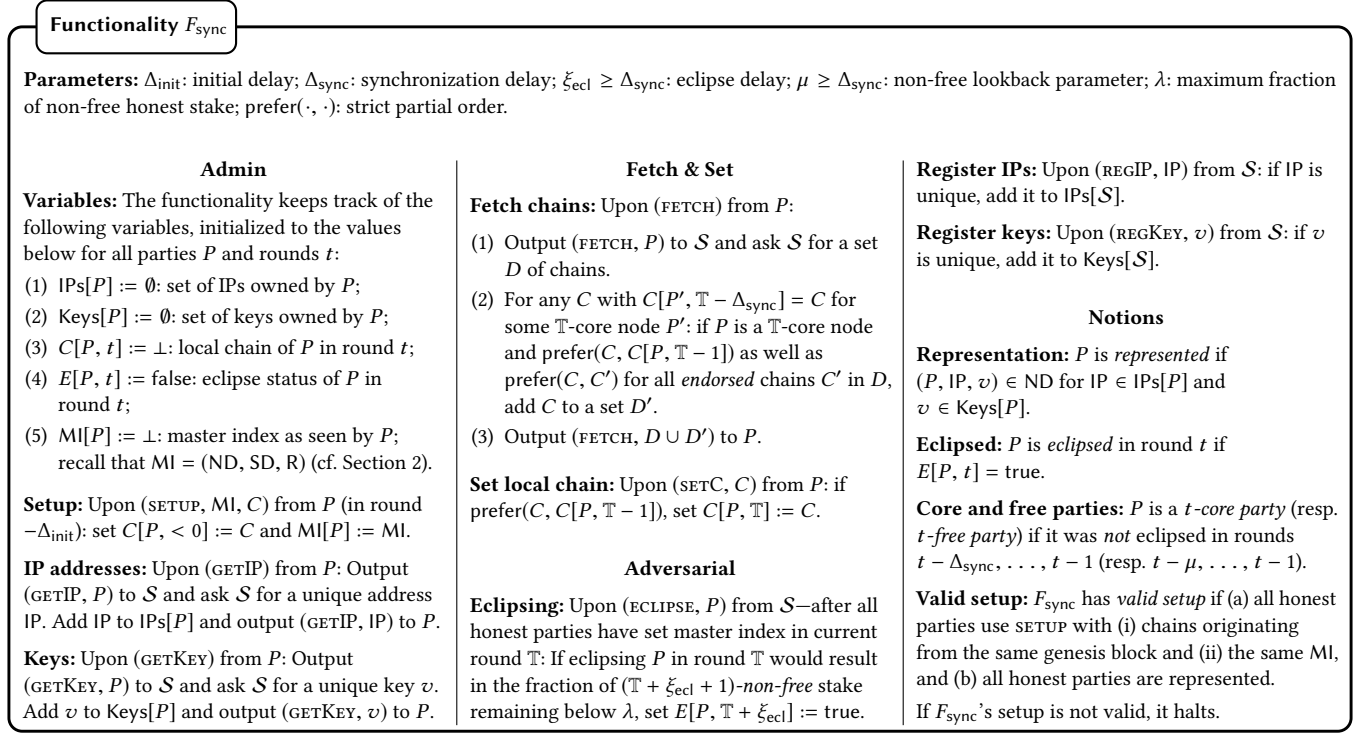
There are two limitations on the attacker’s use of the ECLIPSE command. The first one is that eclipse commands only take effect after a delay of ξ_{ecl} . The second limitation is based on the following notion of *t-free party*:

Definition 3.1. A *t-free party* is an honest party that was not eclipsed during rounds $t - 1, t - 2, \dots, t - \mu$.

The eclipse restriction is that at any time t , the fraction of stake corresponding to *t-non-free* honest parties may not exceed parameter λ .

Formulating an eclipse restriction in this way prevents the attacker from eclipsing a completely different subset of parties in each round since a party eclipsed in round t essentially blocks a fraction corresponding to its stake in the attacker’s eclipse budget λ for μ slots. Note, however, that it is absolutely possible for the adversary to eclipse a *particular* party for an *unlimited* amount of time. This

¹⁰Observe that the keys in the networking layer need not necessarily be VRF keys: any type of “public key” could potentially be used as an ID instead. F_{sync} is sufficiently general to support any such use case since the keys are only used to determine “valid setup” (see below).

Figure 4: Functionality F_{Sync} for global chain synchronization among participants of PoS consensus.

must be dealt with by the higher-level consensus protocol using F_{Sync} .

Main operation and timely synchronization (TS). Similarly to $F_{\text{bilateral}}$, a party uses SETC to switch their local chain to a strictly better one according to a predicate prefer (cf. Section 3.2). Moreover, parties use FETCH in order to receive information about chains of other participants.

The TS guarantee offered by F_{Sync} is based on the following notion of t -core:

Definition 3.2. A t -core party is an honest party that was not eclipsed during rounds $t - 1, t - 2, \dots, t - \Delta_{\text{sync}}$.

Thus, the t -core notion is very similar to that of t -free parties, except with less “lookback” (as $\Delta_{\text{sync}} \leq \mu$).¹¹

The TS guarantee is now the following: Suppose an honest party P sets its local chain to some chain C at time t . Then, by time $t + \Delta_{\text{sync}}$, provided P is in the $(t + \Delta_{\text{sync}})$ -core, the following holds for all $(t + \Delta_{\text{sync}})$ -core parties P' ,

- either P' has already switched their local chain to C' , or
- the FETCH command returns C' ,

where C' is a chain with $\neg \text{prefer}(C, C')$, i.e., incomparable to or better than C .¹²

¹¹The reason for having two parameters Δ_{sync} and μ here is that F_{Sync} is more “useful” the smaller Δ_{sync} is and the larger μ is.

¹²Note that since prefer is a strict partial order, $C' = C$ is also possible.

3.4 Synchronization Protocol

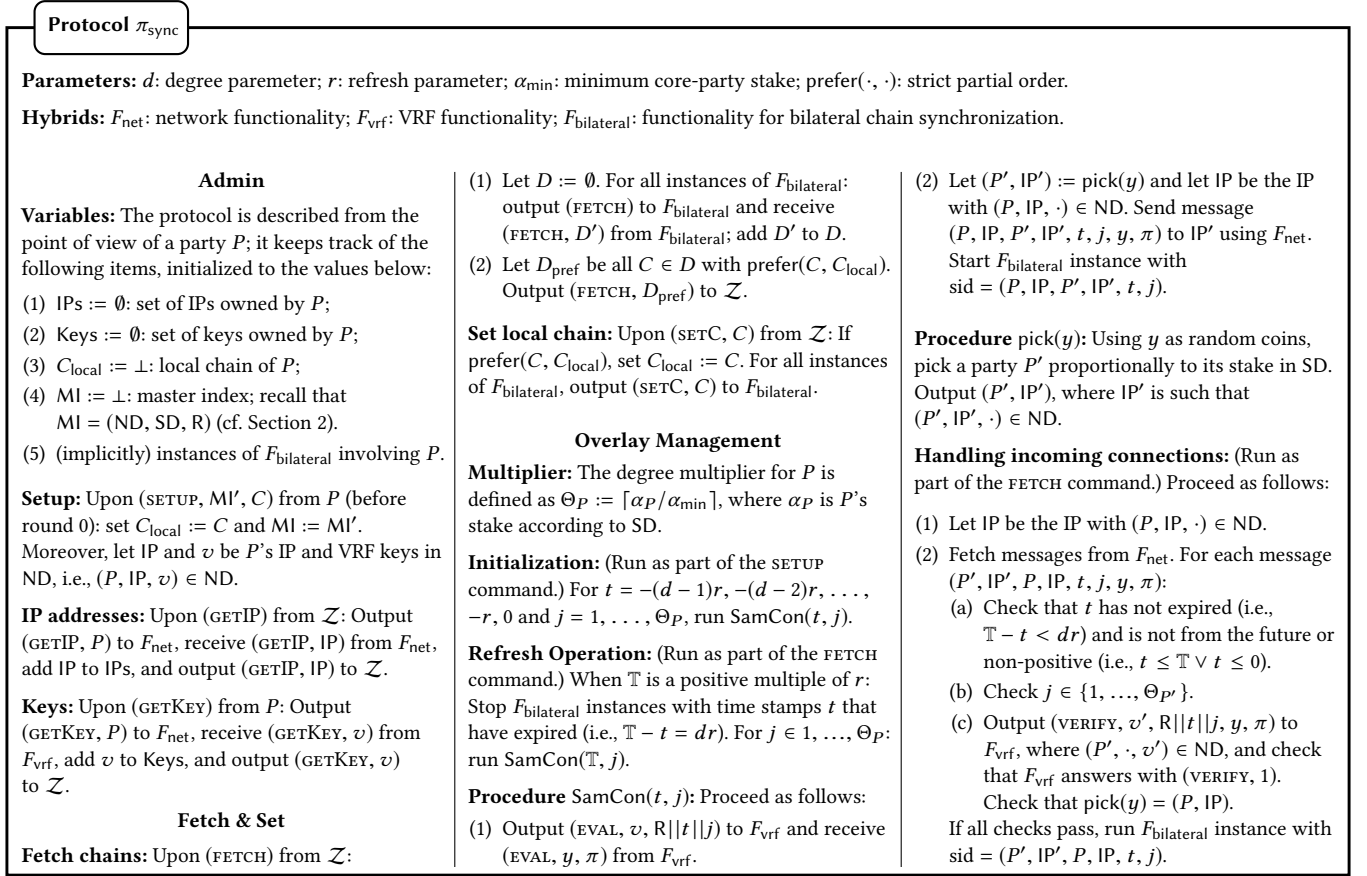
Overview. Protocol π_{sync} (cf. Figure 5), realizing F_{Sync} in the hybrid model with F_{net} , F_{vrf} , and $F_{\text{bilateral}}$, follows the stake-weighted random-graph approach outlined in Section 3.1. The protocol uses a VRF to determine the random graph and to ensure that honest parties only accept connections from parties they are neighbors of in the graph. A party runs instances of $F_{\text{bilateral}}$ with each neighbor in order to keep track of their local chain. Whenever a party learns of a *valid* chain C' better than their current local one, the idea is that they switch to C' . However, as elaborated below, validity checks and the decision to switch occur in the environment; therefore, π_{sync} will only realize F_{Sync} w.r.t. to a (very reasonably) restricted class of environments.

Admin. Protocol π_{sync} handles requests for IPs and keys by simply forwarding them to F_{net} and F_{vrf} , respectively, and the subsequent replies back to the environment.

Upon receiving $(\text{SETUP}, \text{MI}', C)$ from the environment, π_{sync} stores these values internally for later use.

Overlay. The most crucial part of π_{sync} deals with establishing and updating the random-graph overlay. A party P initially samples Θ_P neighbors for each time stamp $t = -(d-1)r, -(d-2)r, \dots, -r, 0$ independently and based on their stake, where Θ_P is a multiplier that depends on the stake α_P of P itself; specifically,

$$\Theta_P := \lceil \alpha_P / \alpha_{\min} \rceil,$$

Figure 5: Protocol π_{sync} , implementing functionality F_{sync} .

for some parameter α_{\min} .

Connections to neighbors expire after dr slots; thus, in each round that is a multiple of r , Θ_P new neighbors are sampled. Updating the overlay in this fashion helps parties recover from eclipse events in practice; in the context of this work, however, the adaptive attacker considered here may simply corrupt all new neighbors of a party. Hence, with such a powerful adversary, there is no upper bound on the duration of an eclipse for a particular party (short of the adversary exhausting its corruption budget, of course). Furthermore, refreshing neighbors also allows gradual adaptation of the overlay to changing stake distributions.

The actual sampling performed by a party P is described as a procedure $\text{SamCon}(t, j)$ in Figure 5, where t is a time stamp and $j = 1, \dots, \Theta_P$. The procedure uses a subroutine $\text{pick}(y)$ which, using VRF output y as random coins, chooses a party P' proportionally to its stake in SD . Note that the inputs to the VRF are (apart from P 's public key) the random nonce R (part of MI), t , as well as j .

Subsequent to the above sampling, P sends a connection request to P' via F_{net} and immediately begins a corresponding instance of $F_{\text{bilateral}}$. Conversely, P' will start the instance upon receiving the connection request (after performing the obvious checks).

Chain management. Upon receiving (SETC, C) from the environment, if C is better than P 's current local chain, P updates the corresponding variable and inputs (SETC, C) to all running $F_{\text{bilateral}}$ instances.

Upon receiving (FETCH) from the environment, P sends (FETCH) to all running $F_{\text{bilateral}}$ instances and collects the answer chains in a set D ; all chains in D preferable to P 's current local chain are returned to the environment.

Chain validity and switching. With the application of blockchain consensus in mind, observe that in the context of π_{sync} (or, in the ideal world, of F_{sync}) there is no notion of “chain validity.” This is a concept from the higher-level consensus layer. Consequently, π_{sync} (F_{sync}) cannot possibly check a chain for its validity and “automatically” switch to the best valid chain. This task therefore falls upon the environment.

At this point it is also important to observe that successful chain “propagation” through the network crucially depends on the environment—in each round—using SETC to make each party switch to the best *valid* chain received so far. In order to circumvent the application-specific nature of the notion of “validity,” the following proxy for it is used:

Definition 3.3. In an execution of π_{sync} (or F_{sync}), a chain C is *endorsed* if the environment at some point inputs (SETC, C) on behalf of an honest party P .

The notion of endorsed chains allows to define the class of environments \mathcal{Z} w.r.t. which π_{sync} must realize F_{sync} :

Definition 3.4. An environment \mathcal{Z} is *H-improving* if in each round t , every honest party P inputs (SETC, C) for a chain C which is maximal w.r.t. prefer among all endorsed chains received by P via FETCH in rounds up to t .

The above finally leads to the following main theorem:

THEOREM 3.5. *Let*

- $n \in \mathbb{N}$,
- $\alpha > \beta > 0$, $\delta > 0$, c_{\min} , and c_{\max} be constants satisfying $0 < c_{\min} < 1 < c_{\max}$ and $(\alpha - \beta)/2 \geq \beta + c_{\min}/n$, and
- $r \in \mathbb{N}$.

Then, for all sufficiently large d , there exists a constant $\gamma > 0$ such that protocol $\pi_{\text{sync}}[d + 1, r, c_{\min}/n]$ ϵ -securely realizes $F_{\text{sync}}[\Delta_{\text{init}}, \Delta_{\text{sync}}, \mu, \xi_{\text{ecl}}, \lambda]$ in the $\{F_{\text{net}}[\delta_{\text{net}}], F_{\text{vrf}}, F_{\text{bilateral}}[\delta_{\text{init}}, \delta_{\text{sync}}]\}$ -hybrid model w.r.t. H-improving environments that (1) input an MI with n parties in which no party owns more than a c_{\max}/n fraction of stake, (2) corrupt at most an α -fraction of stake, (3) run for at most L rounds, where

- $\Delta_{\text{init}} = \delta_{\text{net}} + \delta_{\text{init}}$,
- $\Delta_{\text{sync}} = 8\ell$ for $\ell := \lceil -\log_{1+\gamma}(2c_{\min}/n) + 1 \rceil$,
- $\Delta_{\text{sync}} \leq \xi_{\text{ecl}} \leq \min(\xi_{\text{corr}}, r)$,
- $\mu = r$,
- $\delta_{\text{init}} \leq r$,
- $\lambda = 2(\beta + c_{\min})$, and
- $\epsilon = L/r \cdot e^{-\delta n}$.

Remarks. A simple Chernoff bound can be used to show that if there are no small-stake nodes (with less than c_{\min}/n stake), the degrees of all parties are constant. In case there are many small-stake nodes, the degrees of large nodes are beyond constant. However, one can show that by having small-stake node pick their peers uniformly (instead of based on stake), all node degrees go back to constant again.

The security error ϵ is exponentially small in the number of parties n . It is, of course, important in practice to ensure by suitable means that there are sufficiently many parties participating. One possible way to achieve this by having so-called stake-pool operators (SPOs), to which parties can delegate stake, run the blockchain and use incentive mechanisms to control the total number of SPOs.

Furthermore, there are also several ways of enforcing the maximum-stake restriction in Theorem 3.5. In a system with SPOs, one may restrict the maximum delegated stake on the consensus level. Alternatively, one can again set incentives in such a way that SPOs do not attract more than a certain amount of stake. For more information, see [4].

4 SECURITY PROOF

This section presents the security proof of the synchronization protocol π_{sync} (cf. Section 3.4).

The most crucial part of the security argument is a new result on stake-based expander graphs. Specifically, given a graph G wherein

each vertex v has assigned to it some stake $\alpha_v \in (0, 1]$, where $\sum_v \alpha_v = 1$, and each vertex chooses its neighbors in the stake-based fashion adopted by protocol π_{sync} , the results in Section 4.1 show that even after removing all adversarial nodes from G , leaving behind at least some α -fraction of honest stake, there exists an honest “backbone” holding at least $\alpha - \beta$ stake, for some β , such that the backbone is an expander graph.

As shown in Section 4.2, the above translates to π_{sync} realizing F_{sync} with roughly a β -fraction of honest stake being eclipsed,¹³ while due to the expander property of the backbone, the remaining “core” of honest parties are at most $O(\log n)$ hops apart from each other, where n is the total number of parties. It should be noted that this core only approximately corresponds to the backbone above, the reason for this being that (a) the backbone cannot be efficiently computed (and thus, an approximation has to be used), and (b) backbone nodes with less than c_{\min}/n stake need to be excluded from the core (because the expander property cannot be used to bound their distance from other nodes in the backbone).

4.1 Expanders Resisting Vertex Deletion

It is a well-known fact that random graphs—with a wide variety of edge distributions—form expanders with high probability and hence have small diameter and other desirable properties. This section shows that the random graphs produced by the protocol possess such strong properties *even if an adversary with full knowledge of the graph is permitted to remove a constant fraction of the nodes*.

As outlined above, we show that for any two constants $\alpha > \beta > 0$, there is a degree parameter d for which the following holds: so long as an α fraction of nodes remain after adversarial deletion, there is a subset (the “backbone”) consisting of an $\alpha - \beta$ fraction of the nodes that is a strong expander. Typically, one would choose $\beta \ll \alpha$ so that the backbone consists of almost all of the remaining vertices.

These results are motivated by a classical theorem of Upfal [29], which uses different methods to establish a similar property of Ramanujan graphs (which achieved fixed constants α and β). Thus, our results expand on this theory by (i) handling random graphs, and (ii) establishing that *any* constants can be achieved by appropriate choice of d . A final remark before transitioning to the technical survey: we also work directly with weighted graphs (with a weighting corresponding to stake) so that we can define and treat a natural notion of “stake-weighted” expansion.

Our protocol directly motivates the following family of distributions of random directed graphs.

Definition 4.1. Let \bar{d} and n be positive integers and let \mathcal{D} be a probability distribution on $[n]$ with the property that $d_v = \bar{n}\mathcal{D}(v)$ is an integer for every vertex v . We let $\mathcal{G}_{n, \bar{d}; \mathcal{D}}$ denote the probability law on directed multigraphs with $V = [n]$ obtained by selecting, for each v , d_v outgoing neighbors w_1, \dots, w_{d_v} independently according to \mathcal{D} and defining the multiset of directed edges to be $E = \bigcup_v \bigcup_{i=1}^{d_v} (v, w_i)$.¹⁴

We let $\mathcal{G}_{n, d}$ denote the special case when \mathcal{D} is the uniform distribution, in which case $d_v = d$ for all v .

¹³More precisely, the non-free parameter λ is roughly β (cf. Section 3.3).

¹⁴In our main application, we are actually interested in the properties of the *undirected* version of such graphs; however, for analytic purposes it is convenient to distinguish the source and sinks for each edge.

We wish to show that $\mathcal{G}_{n,\bar{d};\mathcal{D}}$ is typically a “stake expander”: that is, that sets $S \subset V$ have a number of neighbors outside them, or edges leaving them, proportional to their total stake.

Definition 4.2. Let $G = (V, E)$ be a directed graph. For a subset $S \subset V$, we define the (outer) boundary of S as

$$\partial(S) = \{w \notin S \mid \exists s \in S : (s, w) \in E \text{ or } (w, s) \in E\}.$$

Let \mathcal{D} be a distribution on the set V and $\gamma > 0$. We say that G is a (\mathcal{D}, γ) -expander if for every subset of vertices $S \subset V$ for which $\mathcal{D}(S) \leq 1/2$,

$$\mathcal{D}(\partial(S)) \geq \gamma \mathcal{D}(S),$$

where we use the notation $\mathcal{D}(S)$ to denote $\sum_{s \in S} \mathcal{D}(s)$.

THEOREM 4.3. *Let $\alpha > \beta > 0$ and $\delta > 0$ be positive constants, and let c_{\min} and c_{\max} be positive constants satisfying $c_{\min} < 1 < c_{\max}$. For sufficiently large d there is a constant $\gamma > 0$ for which the following holds: Let \mathcal{D} be a distribution on $V = [n]$ for which $c_{\min}/n \leq \mathcal{D}(v) \leq c_{\max}/n$ and $d_v \triangleq nd\mathcal{D}(v)$ is an integer for each $v \in V$. Consider $G = (V, E)$ drawn according to $\mathcal{G}_{n,d;\mathcal{D}}$. Then, except with probability $p_{\text{fail}} \leq e^{-\delta n}$, for every subset $H \subseteq V$ for which $\mathcal{D}(H) \geq \alpha$, there is a subset $H' \subset H$ for which $\mathcal{D}(H') \geq \alpha - \beta$ and the subgraph induced by H' is a (\mathcal{D}', γ) -expander, where \mathcal{D}' is the distribution \mathcal{D} scaled by $1/\mathcal{D}(H')$.*

See Appendix A for the proof.

Remark. The condition in Definition 4.1 and Theorem 4.3 that $nd\mathcal{D}(v)$ is an integer for all v is a mere convenience so that the out-degree of each vertex will be a fixed integer d_v . We can also define $d_v = \lceil nd\mathcal{D}(v) \rceil$. This changes the effective stake of each player by a factor of $1 \pm O(1/d)$. In fact we only need an upper bound on the stake each vertex has, as we can hand over all the vertices with very low stake to the adversary (cf. Lemma A.5).

4.2 Security of the Synchronization Protocol

This section uses the results from Section 4.1 to finally prove the security of protocol π_{sync} .

Simulation basics. Simulator \mathcal{S} internally simulates instances of the protocol and of hybrids F_{net} , F_{vrf} , and $F_{\text{bilateral}}$ as well as the adversary \mathcal{A} (which acts as the interface to the environment); in the following, this ensemble is referred to as the *simulated real world* (SRW). Specifically, \mathcal{S} reacts as follows when receiving messages from F_{sync} :

- Upon (SETUP, P , MI, C): in the SRW, input (SETUP, MI, C) on behalf of P .
- Upon (GETIP, P): in the SRW, input (GETIP) on behalf of P ; wait to receive (GETIP, IP) on behalf of P ; return IP to F_{sync} .
- Upon (GETKEY, P): in the SRW, input (GETKEY) on behalf of P ; wait to receive (GETKEY, v) on behalf of P ; return v to F_{sync} .
- Upon (FETCH, P): in the SRW, input (FETCH) on behalf of P ; wait to receive (FETCH, \tilde{D}); return \tilde{D} to F_{sync} .
- Upon (SETC, P , C): in the SRW, input (SETC, C) on behalf of P .

Messages from \mathcal{A} are handled as follows:

- Upon (REGIP, IP): input (REGIP, IP) to F_{sync} .
- Upon (REGKEY, v): input (REGKEY, v) to F_{sync} .

Notation. Fix the master index MI = (ND, SD, R) input by the honest parties at the beginning; recall that n denotes the number of parties in MI. In the following, for a subset S of the parties in MI, denote by $\alpha_S := \text{SD}(S)$ the amount of stake held by parties in S .

Determining whom to eclipse. A crucial part of \mathcal{S} is to ensure that when handling FETCH commands, F_{sync} does not enforce delivery guarantees that contradict the SRW. To that end, \mathcal{S} determines which honest parties are “eclipsed” and relays this information to F_{sync} . Recall that F_{sync} will not offer any guarantees to non-core parties, where a party is in the core in round t if and only if it has not been eclipsed in rounds $t - 1, t - 2, \dots, t - \Delta_{\text{sync}}$.

Towards understanding which parties to eclipse at a particular time t , consider the graph G formed by all parties in MI and the connections implied by the use of the VRF at time t , but ignore the edges added to it during the latest refresh (i.e., in the largest round $t' \leq t$ that is a multiple of r). This graph follows the distribution $\mathcal{G}_{n,d;\text{SD}}$.

Let H denote the set of honest parties in at time t . By Theorem 4.3, there exists a set $H' \subseteq H$ with $\alpha_{H'} \geq \alpha_H - \beta$, such that the subgraph of G induced by H' is an (SD', γ) -expander, where SD' is SD scaled by $1/\alpha_{H'}$. This expander property allows to bound the diameter between parties in H' with a certain minimum amount of stake, a fact used to show:

CLAIM 4.4. *From each party in H' with stake at least c_{\min}/n in SD , one can reach more than $(\alpha_H - \beta)/2$ of honest stake in at most ℓ steps.*

PROOF. Consider a party $P \in H'$ with stake $\alpha_P \geq c_{\min}/n$. Its stake according to SD' is equal to

$$\alpha_P / \alpha_{H'} \geq \frac{c_{\min}}{\alpha_{H'} n}.$$

By the expander property on H' , one can reach more than $1/2$ of the stake in H' from P (according to SD') in ℓ' steps if

$$\frac{c_{\min}}{\alpha_{H'} n} (1 + \gamma)^{\ell'} > \frac{1}{2}.$$

This is satisfied for

$$\ell' := \left\lceil \log_{1+\gamma} \left(\frac{\alpha_{H'} n}{2c_{\min}} + 1 \right) \right\rceil \leq \left\lceil \log_{1+\gamma} \left(\frac{n}{2c_{\min}} + 1 \right) \right\rceil = \ell.$$

Reaching half of the stake in H' according to SD' translates to reaching

$$\frac{\alpha_{H'}}{2} \geq \frac{\alpha_H - \beta}{2}$$

of stake according to SD . \square

Since the security argument in Theorem 4.3 is non-constructive, it is unclear whether set H' can be efficiently determined. Instead, consider the set

$$I_t := \{h \in H \mid \text{can reach } \beta + c_{\min}/n \text{ HS in } \ell \text{ steps from } h\},$$

where HS stands for “honest stake.” The subscript is dropped from I whenever clear from the context.

CLAIM 4.5. $\alpha_I \geq \alpha_H - \beta - c_{\min}$.

PROOF. First, observe that $(\alpha_H - \beta)/2 \geq (\alpha - \beta)/2 \geq \beta + c_{\min}/n$, where the last inequality is by assumption. By Claim 4.4, all parties in H' with stake at least c_{\min}/n are therefore in I . The claim follows by observing that at most c_{\min} of the total stake is held by parties with more than c_{\min}/n stake. \square

As the subgraph induced by H' , that induced by I also has bounded diameter:

CLAIM 4.6. *There is a path of length at most 4ℓ between any two parties in I .*

PROOF. By definition of I , from any node in I , one can reach at least $\beta + c_{\min}/n$ stake in H in ℓ steps, and, therefore, c_{\min}/n stake in H' . The argument in the proof of Claim 4.4 suggests there is a path of length at most 2ℓ between any two nodes with stake at least c_{\min}/n in H' . One concludes that there must be a path of length at most 4ℓ between any two nodes in I . \square

Given the above, the eclipse strategy of \mathcal{S} is to, in each round t , use (ECLIPSE, P) to eclipse all parties $P \notin I_t$. However, due to the eclipse-delay property of F_{sync} , \mathcal{S} is forced to set eclipsed status based on I_t at or before time $t - \xi_{\text{ecl}}$. This is, however, possible:

- Since $\xi_{\text{ecl}} \leq r$ and the edges added during the latest round $t' \leq t$ that was a multiple of r are ignored, the topology of G in round t is known by round $t - \xi_{\text{ecl}}$.
- Since $\xi_{\text{corr}} \geq \xi_{\text{ecl}}$, which parties are honest/corrupted is also known by round $t - \xi_{\text{ecl}}$.

Bounding synchronization time. The next step in the proof consists of arguing that F_{sync} never has to add any chain C to set D' during a FETCH call by party P in some round t .

To that end, assume first that there is no multiple of r in $\{t - 4\ell, \dots, t - 1\}$, i.e., the underlying graph G does not change in this period. Consequently, $I_{t-4\ell} \supseteq I_{t-(4\ell-1)} \supseteq \dots \supseteq I_{t-1}$, where the only reason elements are dropped from I over time is corruption.

Consider now the FETCH command by P in round t . Recall that a chain C is added to D' only if

- C was held by some honest party P' in round $t - \Delta_{\text{sync}}$,
- neither P nor P' were eclipsed in rounds $t - 1, \dots, t - \Delta_{\text{sync}}$,
- C is preferable to the chain held by P' in round $t - 1$ and to all chains in set D (during the FETCH call).

Observe that $\Delta_{\text{sync}} \geq 4\ell$; thus, that the fact that (ii) holds (i.e., \mathcal{S} did not issue, in advance, eclipse commands for P or P' for rounds $t - 1, \dots, t - \Delta_{\text{sync}}$) means that both parties P and P' were in sets I during rounds $t - 1, \dots, t - 4\ell$.

Therefore, by the monotonicity of the sets I , there has been a path of length at most 4ℓ between P' and P , which, combined with the facts that \mathcal{Z} is H-improving and (again) $\Delta_{\text{sync}} \geq 4\ell\delta_{\text{sync}}$, means that by time t , either P' 's local chain will already be set to a chain $C' \neq C$ or a chain $C' \neq C$ is in set D at time t . In either case, C is not added to D' .

Finally, for intervals including changes to G , observe that since $\Delta_{\text{sync}} = 8\ell$, there are at least 4ℓ rounds either before or after the change to G .

Amount of non-free stake. The fact that the amount of non-free stake remains below $\lambda = 2(\beta + c_{\min})$ follows from Claim 4.5, noting that the extra factor of 2 is required for Δ_{sync} -sized intervals containing a multiple of r .

PROOF (OF THEOREM 3.5). In order to complete the proof based on the above, note that it remains merely to apply a union bound over all refresh periods (which are of length r). \square

5 FULLY SECURE POS CONSENSUS

This section discusses the application of running a PoS consensus protocol on top of the chain synchronization functionality F_{sync} . For concreteness, the protocol considered here is Ouroboros Praos [11] which is a longest chain Nakamoto-style blockchain PoS protocol and is described first. This is followed by a comparison of the “diffusion” functionality F_{diff} used in [11] and F_{sync} . Based on this comparison, the changes to Ouroboros Praos’ security proof are presented and discussed; in particular, as a contribution of possibly independent interest, a generalized version of the so-called *forkable-string analysis* is provided.

5.1 Ouroboros Praos

Ouroboros Praos (or, simply, *Praos*) proceeds in so-called slots. In each slot, each party P first checks whether it has “received” valid chains that are *longer* than its current local chain; if so, it switches to a longest one among those. Subsequently, based on the stake distribution used in the current epoch, P checks whether it is the slot leader. This determination is made based on the output y of a VRF evaluated on (in addition to P 's key) the slot number and the so-called epoch randomness; if y ends up below a certain threshold, which is a monotonically increasing function of P 's stake, P is a slot leader. As such, it creates a new block extending its current local chain C ; the block consists of the hash of the last block of C , the slot number, P 's identity, the VRF output and VRF proof, a data payload, as well as a signature. The resulting chain C' is then “sent to everyone.”

5.2 Diffusion vs. Synchronization

In [11], parties share a diffusion functionality F_{diff} , parametrized by a value Δ , with the simple intuitive property that any chain an honest party diffuses via F_{diff} will “arrive” at all other honest parties with a delay of no more than Δ slots.

Consider now using F_{sync} , in order to “synchronize” rather than “diffuse” chains. Figure 6 contains a description of the (static-stake) Praos protocol using F_{sync} . The protocol additionally uses hybrids F_{vrf} for slot leadership as well as the following two hybrids, which are only described on a high-level sufficient for the context of this section:

- F_{init} is used for master-index and genesis-block generation; more precisely, parties initially send information such as keys, IPs, etc. to F_{init} , which then generates the initial master index and the genesis block.
- F_{kes} is a functionality idealizing key-evolving signatures and, in particular, allows parties to create keys as well as to issue and verify signatures.

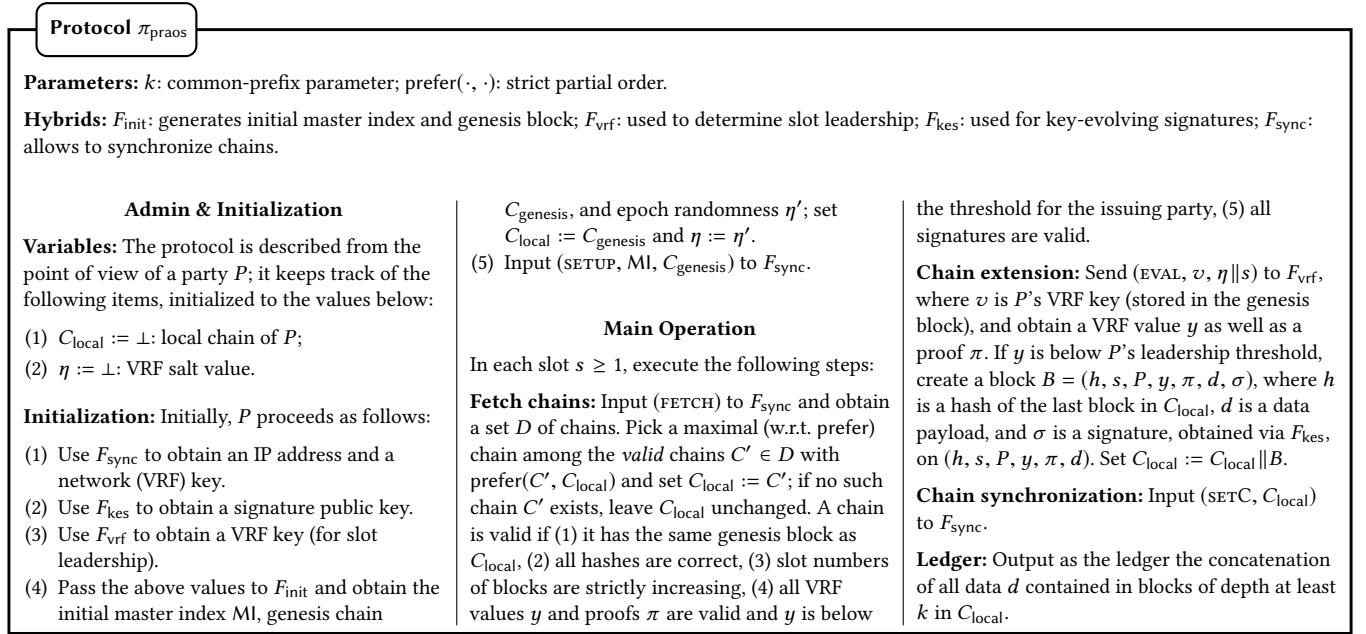


Figure 6: Protocol π_{praos} .

Due to the fact that it is implemented by a protocol—rather than merely assumed—the guarantees offered by F_{sync} are weaker in several ways:

- A λ -fraction of the total stake may be eclipsed, and the corresponding parties are excluded from the timely synchronization (TS) guarantees.
- The TS properties of F_{sync} are weaker compared to those offered by F_{diff} : instead of a particular chain C “propagating” through the network within Δ slots, F_{sync} may deliver different chains C' that are not worse (i.e., equal length or longer).
- In order to replace F_{sync} by its implementation, the environment must be H-improving (cf. Section 3.4).

The next section details how to update the security proofs of Praos in order to deal with these weaker guarantees of F_{sync} .

5.3 Updated Proofs for Praos

5.3.1 Characteristic Strings and Forks. At the heart of the Praos security proof (with functionality F_{diff}) [11] are the notions of *characteristic strings*, *forks*, and *margin*, which capture the security-relevant information about events in an execution of Praos. Characteristic strings indicate the relevant information about the sequence of elected leaders in an execution of the protocol. The analysis shows that consistency violations can be controlled by (i) distilling the family of possible blocktrees that can arise for a given sequence of leader elections (as determined by a characteristic string) into a single numeric metric of interest called *margin* and (ii) an analysis of the stochastic process that governs generation of characteristic strings and the resulting behavior of margin. The final result is articulated as a large-deviation bound for margin, which establishes

consistency with high probability. In this section, we discuss how that analysis can be adapted to our setting with a finer-grained view of networking and block delivery.

Characteristic strings in Praos have the form $w = w_1 w_2 \dots$ and record information about slot leadership in an execution. Each character in the string is a symbol $w_i \in \{A, H, \perp\}$, where

$$w_i = \begin{cases} A & \text{if slot } i \text{ has an adversarial or multiple leaders,} \\ H & \text{if slot } i \text{ has a single honest leader,} \\ \perp & \text{if slot } i \text{ has no leader.} \end{cases}$$

A Δ -fork F for a characteristic string w is a directed, rooted tree, intended to represent the topology of all chains observed during an execution of Praos: Each vertex v of F corresponds to a block in a particular chain and has a label $\ell(v) \in \mathbb{N}$, which records the block’s slot number. The genesis block is represented by the root of the tree. The edges of a fork are directed “away from” the root so that there is a unique directed path from the root to any vertex.

Based on the description of Praos, it is easy to see that a fork satisfies the following properties:

- (i) The root $r \in V$ has label $\ell(r) = 0$ and is considered honest by fiat.
- (ii) The sequence of labels $\ell(\cdot)$ along any directed path is strictly increasing. Reason: this is enforced by the protocol.
- (iii) If $w_i = H$, there is a unique vertex v for which $\ell(v) = i$. Reason: honest parties do not create multiple blocks.
- (iv) For any pair of honest vertices v, v' (i.e., $w_{\ell(v)} = w_{\ell(v')} = H$) with $\ell(v) + \Delta \leq \ell(v')$, their lengths (i.e., distance from root) $\text{len}(v)$ and $\text{len}(v')$ satisfy $\text{len}(v) < \text{len}(v')$. Reason: As per the guarantees of F_{diff} , during the creation of the block corresponding to vertex v' in slot $\ell(v')$, the block

corresponding to v was available to build on since it was created in slot $\ell(v)$, which was at least Δ slots before v' ; hence, $\text{len}(v')$ must be strictly larger than $\text{len}(v)$.

The analysis proceeds by defining a notion of *margin* for a fork, which reflects the presence of pairs of paths (chains) in the fork that diverge prior to a particular slot and exceed the length of the deepest honest block. Intuitively, such pairs of paths correspond to a consistency failure and, indeed, the precise definition of margin ensures that the quantity is positive exactly when such a consistency failure exists. This notion is extended to characteristic strings by maximizing over all forks consistent with the string. The analysis then shows that margin satisfies a recurrence relation in terms of the characteristic string (corresponding to an execution of the protocol) and, finally, that the probability of observing a positive margin is small.

5.3.2 Accounting for Eclipsed Parties. Clearly, the Δ -fork formalism above does not consider cases where (a) the leader is out of sync and potentially fails to build on some longest chain C whose last block is more than Δ slots old or (b) the leader's block takes longer than Δ to reach the next slot leader; naturally, an eclipse event can cause both (a) and (b).¹⁵ In the following, an honest leader *not* suffering from (a) is called *current*, and an honest leader *not* suffering from (b) is called *relayed*. An honest leader that is both current and relayed is called *synchronized*.

In order to update the forkable-string analysis to account for non-relayed parties, the following changes are introduced to it. First, the characteristic string is now over an alphabet $w_i \in \{A, \mathbb{C}_R, C, R, \perp\}$, where

$$w_i = \begin{cases} A & \text{if slot } i \text{ has an adversarial or multiple leaders,} \\ \mathbb{C}_R & \text{if slot } i \text{ has a single synchronized leader,} \\ C & \text{if slot } i \text{ has a single current leader,} \\ R & \text{if slot } i \text{ has a single relayed leader,} \\ \perp & \text{if slot } i \text{ has no leader.} \end{cases}$$

Second, condition (iv) for forks is **amended** as follows:

- (iv) For any pair v, v' with $w_{\ell(v)} \in \{R, \mathbb{C}_R\}$ and $w_{\ell(v')} \in \{C, \mathbb{C}_R\}$ and where $\ell(v) + \Delta \leq \ell(v')$, their lengths (i.e., distance from root) $\text{len}(v)$ and $\text{len}(v')$ satisfy $\text{len}(v) < \text{len}(v')$.

The definition of margin is essentially unchanged with this new definition of fork. However, the recursive behavior depends on the new semantics of these richer characteristic string symbols and exhibits some rather interesting properties: in particular, the effect of the C and R symbols depends on whether the worst-case margin is currently negative or positive. We discuss this in detail in Appendix B.

5.3.3 Alternative chains and H-improving. It is easy to see that (1) condition (iv) of the forks is not violated if instead of a chain C non-worse chains C' are delivered after Δ_{sync} slots and (2) Praos is an H-improving environment for F_{sync} .

¹⁵Note that there are implementations of the network layer in which the connections to peers are unidirectional. In such cases, it is possible that due to an eclipse event (a) occurs but (b) does not (or vice versa).

5.4 Multi-Epoch Praos

There are two approaches to running Praos with F_{sync} in a multi-epoch setting with an evolving stake distribution: (a) use a different instance of F_{sync} in every epoch or (b) define a version of F_{sync} that allows for changing stake distributions.

The conceptually simpler approach is, of course, (a). The idea is to, Δ_{init} before the end of each epoch, begin a new instance of F_{sync} . At that point, the master index (including the network directory, the stake distribution, and the randomness) to be used in the new epoch is determined.¹⁶ Taking approach (a) would mean that the random overlay is completely re-sampled for each new epoch. In practice, this may constitute an unacceptable overhead. Using approach (b) avoids this issue, but makes the definition, realization, and proofs of F_{sync} somewhat more cumbersome.

Observe that in either case it is important to properly deal with honest parties that have been eclipsed for more than an epoch. The issue with these parties is that once they have been eclipsed long enough for their master index to be out of date, they will face difficulties ever reconnecting to the network since the choice of neighbors via the VRF is based on an up-to-date stake distribution.

While in practice such a scenario is unlikely to occur due to the evolving nature of the random-graph overlay, in theory the adaptive attacker can keep corrupting a party's neighbors until it exhausts its corruption budget, thereby keeping it eclipsed for a long period of time. This issue can be mitigated by assuming a checkpointing functionality (realized, e.g., via some light-client infrastructure) that supplies parties, sufficiently in advance of an epoch's beginning, with the relevant master index.

6 ACKNOWLEDGEMENTS

Cristopher Moore was partially supported by National Science Foundation grant 1838251. Alexander Russell was partially supported by National Science Foundation grant 1801487.

REFERENCES

- [1] Ittai Abraham, T.-H. Hubert Chan, Danny Dolev, Kartik Nayak, Rafael Pass, Ling Ren, and Elaine Shi. 2019. Communication Complexity of Byzantine Agreement, Revisited. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 317–326. <https://doi.org/10.1145/3293611.3331629>
- [2] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. 2017. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*. IEEE Computer Society, 375–392. <https://doi.org/10.1109/SP.2017.29>
- [3] Erica Blum, Aggelos Kiayias, Cristopher Moore, Saad Quader, and Alexander Russell. 2020. The Combinatorics of the Longest-Chain Rule: Linear Consistency for Proof-of-Stake Blockchains. In *31st SODA*, Shuchi Chawla (Ed.). ACM-SIAM, 1135–1154. <https://doi.org/10.1137/1.9781611975994.69>
- [4] Lars Brünjes, Aggelos Kiayias, Elias Koutsoupias, and Aikaterini-Panagiota Stouka. 2020. Reward Sharing Schemes for Stake Pools. In *IEEE European Symposium on Security and Privacy, EuroS&P 2020, Genoa, Italy, September 7-11, 2020*. IEEE, 256–275. <https://doi.org/10.1109/EuroSP48549.2020.00024>
- [5] Ran Canetti. 2001. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd FOCS*. IEEE Computer Society Press, 136–145. <https://doi.org/10.1109/SFCS.2001.959888>
- [6] Nishanth Chandran, Wutichai Chongchitmate, Juan A. Garay, Shafi Goldwasser, Rafail Ostrovsky, and Vassilis Zikas. 2015. The Hidden Graph Model: Communication Locality and Optimal Resiliency with Adaptive Faults. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015*,

¹⁶Of course, parameters such as epoch length, common prefix, etc. need to be suitably chosen for this to be the case.

- Rehovot, Israel, January 11–13, 2015. Tim Roughgarden (Ed.). ACM, 153–162. <https://doi.org/10.1145/2688073.2688102>
- [7] Jing Chen and Silvio Micali. 2019. Algorand: A secure and efficient distributed ledger. *Theor. Comput. Sci.* 777 (2019), 155–183. <https://doi.org/10.1016/j.tcs.2019.02.001>
- [8] D. Coutts, N. Davies, K. Knutsson, M. Fontaine, A. Santos, M. Szamotulski, and A. Vieth. 2022. The Shelley Networking Protocol. <https://hydra.iohk.io/build/13272760/download/2/network-spec.pdf>.
- [9] D. Coutts, N. Davies, M. Szamotulski, and P. Thompson. 2020. Introduction to the design of the Data Diffusion and Networking for Cardano Shelley. https://hydra.iohk.io/job/Cardano/ouroboros-network/native.network-docs.x86_64-linux/latest/download/1.
- [10] Phil Daian, Rafael Pass, and Elaine Shi. 2019. Snow White: Robustly Reconfigurable Consensus and Applications to Provably Secure Proof of Stake. In *Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18–22, 2019, Revised Selected Papers (Lecture Notes in Computer Science)*, Ian Goldberg and Tyler Moore (Eds.), Vol. 11598. Springer, 23–41. https://doi.org/10.1007/978-3-030-32101-7_2
- [11] Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. 2018. Ouroboros Praos: An Adaptively-Secure, Semi-synchronous Proof-of-Stake Blockchain. In *EUROCRYPT 2018, Part II (LNCS)*, Jesper Buus Nielsen and Vincent Rijmen (Eds.), Vol. 10821. Springer, Heidelberg, 66–98. https://doi.org/10.1007/978-3-319-78375-8_3
- [12] Alan J. Demers, Daniel H. Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard E. Sturgis, Daniel C. Swinehart, and Douglas B. Terry. 1987. Epidemic Algorithms for Replicated Database Maintenance. In *6th ACM PODC*, Fred B. Schneider (Ed.). ACM, 1–12. <https://doi.org/10.1145/41840.41841>
- [13] Alan J. Demers, Daniel H. Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard E. Sturgis, Daniel C. Swinehart, and Douglas B. Terry. 1987. Epidemic Algorithms for Replicated Database Maintenance. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, Vancouver, British Columbia, Canada, August 10–12, 1987*, Fred B. Schneider (Ed.). ACM, 1–12. <https://doi.org/10.1145/41840.41841>
- [14] Cynthia Dwork, David Peleg, Nicholas Pippenger, and Eli Upfal. 1988. Fault Tolerance in Networks of Bounded Degree. *SIAM J. Comput.* 17, 5 (1988), 975–988. <https://doi.org/10.1137/0217061>
- [15] Peter Gazi, Aggelos Kiayias, and Alexander Russell. 2020. Tight Consistency Bounds for Bitcoin. Cryptology ePrint Archive, Report 2020/661. <https://eprint.iacr.org/2020/661>.
- [16] Peter Gazi, Aggelos Kiayias, and Alexander Russell. 2020. Tight Consistency Bounds for Bitcoin. In *ACM CCS 2020*, Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna (Eds.). ACM Press, 819–838. <https://doi.org/10.1145/3372297.3423365>
- [17] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. 2015. Eclipse Attacks on Bitcoin’s Peer-to-Peer Network. In *24th USENIX Security Symposium, USENIX Security 15, Washington, D.C., USA, August 12–14, 2015*, Jaeyeon Jung and Thorsten Holz (Eds.). USENIX Association, 129–144. <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/heilman>
- [18] Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. 2000. Randomized Rumor Spreading. In *41st FOCS*. IEEE Computer Society Press, 565–574. <https://doi.org/10.1109/SFCS.2000.892324>
- [19] Richard M. Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. 2000. Randomized Rumor Spreading. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12–14 November 2000, Redondo Beach, California, USA*. IEEE Computer Society, 565–574. <https://doi.org/10.1109/SFCS.2000.892324>
- [20] Anne-Marie Kermerrec and Maarten van Steen. 2007. Gossiping in distributed systems. *ACM SIGOPS Oper. Syst. Rev.* 41, 5 (2007), 2–7. <https://doi.org/10.1145/1317379.1317381>
- [21] Aggelos Kiayias, Saad Quader, and Alexander Russell. 2020. Consistency of Proof-of-Stake Blockchains with Concurrent Honest Slot Leaders. In *40th IEEE International Conference on Distributed Computing Systems, ICDCS 2020, Singapore, November 29 – December 1, 2020*. IEEE, 776–786. <https://doi.org/10.1109/ICDCS47774.2020.00065>
- [22] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. 2017. Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol. In *CRYPTO 2017, Part I (LNCS)*, Jonathan Katz and Hovav Shacham (Eds.), Vol. 10401. Springer, Heidelberg, 357–388. https://doi.org/10.1007/978-3-319-63688-7_12
- [23] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. 1982. The Byzantine Generals Problem. *ACM Trans. Program. Lang. Syst.* 4, 3 (1982), 382–401. <https://doi.org/10.1145/357172.357176>
- [24] Chen-Da Liu-Zhang, Christian Matt, Ueli Maurer, Guilherme Rito, and Søren Eller Thomsen. 2022. Practical Provably Secure Flooding for Blockchains. Cryptology ePrint Archive, Paper 2022/608. <https://eprint.iacr.org/2022/608>
- [25] Christian Matt, Jesper Buus Nielsen, and Søren Eller Thomsen. 2022. Formalizing Delayed Adaptive Corruptions and the Security of Flooding Networks. Cryptology ePrint Archive, Paper 2022/010. <https://eprint.iacr.org/2022/010>
- [26] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. <http://bitcoin.org/bitcoin.pdf>.
- [27] Joachim Neu, Srivatsan Sridhar, Lei Yang, David Tse, and Mohammad Alizadeh. 2021. Securing Proof-of-Stake Nakamoto Consensus Under Bandwidth Constraint. *CoRR* abs/2111.12332 (2021). arXiv:2111.12332 <https://arxiv.org/abs/2111.12332>
- [28] Elias Rohrer and Florian Tschorsch. 2019. Kademlia: A Structured Approach to Broadcast in Blockchain Networks. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies, AFT 2019, Zurich, Switzerland, October 21–23, 2019*. ACM, 199–213. <https://doi.org/10.1145/3318041.3355469>
- [29] Eli Upfal. 1992. Tolerating Linear Number of Faults in Networks of Bounded Degree. In *Proceedings of the Eleventh Annual ACM Symposium on Principles of Distributed Computing (PODC '92)*. Association for Computing Machinery, New York, NY, USA, 83–89. <https://doi.org/10.1145/135419.135437>
- [30] Vassilis Zikas, Sarah Hauser, and Ueli M. Maurer. 2009. Realistic Failures in Secure Multi-party Computation. In *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15–17, 2009. Proceedings (Lecture Notes in Computer Science)*, Omer Reingold (Ed.), Vol. 5444. Springer, 274–293. https://doi.org/10.1007/978-3-642-00457-5_17

A EXPANDERS RESILIENT AGAINST VERTEX DELETION

We consider the following family of distributions of random directed graphs.

Definition A.1 (Restatement of Definition 4.1). Let \bar{d} and n be positive integers and let \mathcal{D} be a probability distribution on $[n]$ with the property that $d_v \triangleq n\bar{d}\mathcal{D}(v)$ is an integer for every vertex v . We let $\mathcal{G}_{n,\bar{d};\mathcal{D}}$ denote the probability law on directed multigraphs with $V = [n]$ obtained by selecting, for each v , d_v outgoing neighbors w_1, \dots, w_{d_v} independently according to \mathcal{D} and defining the multiset of directed edges to be $E = \bigcup_v \bigcup_{i=1}^{d_v} (v, w_i)$. In our main application, we are actually interested in the properties of the *undirected* version of such graphs; however, for analytic purposes it is convenient to distinguish the source and sinks for each edge.

In the definition above, we let $\mathcal{G}_{n,\bar{d}}$ denote the special case when \mathcal{D} is the uniform distribution, in which case $d_v = \bar{d}$ for all v .

Notation warning. Note that $\mathcal{G}_{n,\bar{d}}$ does not denote the random \bar{d} -regular graph as it often does in the literature. Rather, it is a directed graph where every vertex has out-degree \bar{d} and in-degree distributed as $\text{Bin}(\bar{d}n, 1/n)$.

We wish to show that $\mathcal{G}_{n,\bar{d};\mathcal{D}}$ is typically a ‘‘stake expander’’: that is, that sets $S \subset V$ have a number of neighbors outside them, or edges leaving them, proportional to their total stake.

Definition A.2 (Restatement of Definition 4.2). Let $G = (V, E)$ be a directed graph. For a subset $S \subset V$, we define the (outer) boundary of S as

$$\partial(S) = \{w \notin S \mid \exists s \in S : (s, w) \in E \text{ or } (w, s) \in E\}.$$

Let \mathcal{D} be a distribution on the set V and $\gamma > 0$. We say that G is a (\mathcal{D}, γ) -expander if for every subset of vertices $S \subset V$ for which $\mathcal{D}(S) \leq 1/2$,

$$\mathcal{D}(\partial(S)) \geq \gamma \mathcal{D}(S),$$

where we use the notation $\mathcal{D}(S)$ to denote $\sum_{s \in S} \mathcal{D}(s)$. When \mathcal{D} is the uniform distribution, we say that G is an γ -expander.

PROPOSITION A.3 (DIAMETER BOUND FOR STAKE EXPANDERS). *Let $G = (V, E)$ be a γ -stake expander with (stake) distribution \mathcal{D} . Let $m = \min_{v \in V} \mathcal{D}(v)$ be the minimum stake of any vertex in G . Let \bar{G} be the undirected counterpart of G , obtained by replacing each directed edge with an undirected edge. Writing $m = c/n$, where $n = |V|$, the diameter of \bar{G} is no more than $2 \log_{1+\gamma}(2n/c)$.*

PROOF. Consider two vertices u and v of G . Note that the total stake reachable from u by paths of length k is at least $(1+\gamma)^k \mathcal{D}(u) \geq (1+\gamma)^k m$. If $k \geq \log_{1+\gamma}(1/(2m))$ the stake of this set exceeds $1/2$. The same can be said of the vertex v . Composing these two paths (and reversing the directions of arrows on the second) yields a path of length $2 \log_{1+\gamma}(1/(2m))$, as desired. \square

THEOREM A.4 (RESTATEMENT OF THEOREM 4.3). *Let $\alpha > \beta > 0$ and $\delta > 0$ be positive constants, and let c and C be positive constants satisfying $c < 1 < C$. For sufficiently large \bar{d} there is a constant $\gamma > 0$ for which the following holds. Let \mathcal{D} be a distribution on $V = [n]$ for which $c/n \leq \mathcal{D}(v) \leq C/n$ and $d_v \triangleq n\bar{d}\mathcal{D}(v)$ is an integer for each $v \in V$. Consider $G = (V, E)$ drawn according to $\mathcal{G}_{n,\bar{d};\mathcal{D}}$. Then, except*

with probability $p_{\text{fail}} \leq e^{-\delta n}$, for every subset $H \subseteq V$ for which $\mathcal{D}(H) \geq \alpha$, there is a subset $H' \subset H$ for which $\mathcal{D}(H') \geq \alpha - \beta$ and the subgraph induced by H' is a (\mathcal{D}', γ) -expander, where \mathcal{D}' is the distribution \mathcal{D} scaled by $1/\mathcal{D}(H')$.

Remark. The condition in Definition 4.1 and Theorem 4.3 that $n\bar{d}\mathcal{D}(v)$ is an integer for all v is a mere convenience so that the out-degree of each vertex will be a fixed integer d_v . We can also define $d_v = \lceil n\bar{d}\mathcal{D}(v) \rceil$. This changes the effective stake of each player by a factor of $1 \pm O(1/\bar{d})$.

In fact we only need an upper bound on the stake each vertex has, if we are willing to hand over all the vertices with very low stake to the adversary:

LEMMA A.5. *Let $c < 1 < C$ and let \mathcal{D} be a distribution on $[n]$ such that $\mathcal{D}(v) \leq C/n$ for all v . Then the fraction q of vertices v with $c/n \leq \mathcal{D}(v) \leq C/n$, is bounded below by*

$$q \geq \frac{1-c}{C-c}.$$

It follows that the total probability mass associated with vertices v for which $\mathcal{D}(v) \leq c/n$ is no more than $c(1-q) \leq c(C-1)/(C-c) \leq c$.

PROOF. Otherwise, the total stake would be at most $qC + (1-q)c < 1$. \square

A.1 Mixing and expansion

The Expander Mixing Lemma is a classic result asserting that if $G = (V, E)$ is a d -regular expander and $S, T \subset V$ are sufficiently large disjoint sets, then the number of edges between S and T is close to $(d/|V|)|S| \cdot |T|$ (which can be interpreted as the expected value if the graph had been drawn randomly). This property is useful to us since it is inherited by any induced subgraph of G , e.g., the subgraph H induced by the good players, albeit with modified parameters. As a partial converse, this mixing property implies expansion for sufficiently large sets, since we can take T to be the complement of S .

Here we prove a mixing lemma for graphs in our model. The proof is a simple Chernoff bound with a union bound over all S and T , and is essentially the same as for random d -regular graphs or sparse Erdős-Rényi graphs $G_{n,p=d/n}$. For simplicity, we focus here on the uniform case $\mathcal{G}_{n,\bar{d}}$ and then generalize to the stake-weighted case. We then use it to establish edge expansion and vertex expansion for sufficiently large sets.

LEMMA A.6. *Fix disjoint sets $S, T \in V$ of size $|S| = \sigma n$ and $|T| = \tau n$. Let $G \in \mathcal{G}_{n,\bar{d}}$ and let $E(S, T)$ denote the number of edges in either direction between S and T . For any $\kappa \leq 1$,*

$$\Pr_G \left[\left| \frac{E(S, T)}{2\bar{d}\sigma\tau n} - 1 \right| > \kappa \right] \leq \exp\left(-\frac{\kappa^2 \bar{d}\sigma\tau n}{3}\right). \quad (1)$$

PROOF. Recall that in $\mathcal{G}_{n,\bar{d}}$ each vertex in S chooses \bar{d} outgoing neighbors uniformly and independently from the graph. Since each of these neighbors falls in T with probability $|T|/n$, the number of directed edges from S to T is binomially distributed as $\text{Bin}(\bar{d}|S|, |T|/n)$. Similarly, the number of edges from T to S is distributed as $\text{Bin}(\bar{d}|T|, |S|/n)$. Since $E(S, T)$ is the sum of these two random variables, both of which have expectation $\bar{d}|S||T|/n = \bar{d}\sigma\tau n$,

the event $|E(S, T) - 2\bar{d}\sigma\tau n| > 2t$ implies that at least one of them deviates from its mean by t . A standard two-sided Chernoff bound for binomial variables X with mean μ states that

$$\Pr[|X - \mu| > t] \leq e^{-\frac{t^2}{3\mu}} \quad (2)$$

if $t \leq \mu$. Setting $t = \kappa\bar{d}\sigma\tau n$ and $\mu = \bar{d}\sigma\tau n$ completes the proof. \square

LEMMA A.7. *Let $0 < \varepsilon < 1/2$, $\kappa > 0$, and $\delta > 0$. Then for sufficiently large \bar{d} , with probability $1 - e^{-\delta n}$ in $G \in \mathcal{G}_{n, \bar{d}}$ for all pairs of disjoint sets $S, T \subset V$, if $|S| = \sigma n$ and $|T| = \tau n$ where $\varepsilon \leq \sigma, \tau \leq 1 - \varepsilon$, then*

$$1 - \kappa \leq \frac{E(S, T)}{2\bar{d}\sigma\tau n} \leq 1 + \kappa. \quad (3)$$

PROOF. We use Lemma A.6 and take the union bound over S and T . The number of pairs of sets of size σn and τn , disjoint or not, is

$$\binom{n}{\sigma n} \binom{n}{\tau n} \leq \exp((h(\sigma) + h(\tau))n),$$

where $h(x) = -x \ln x - (1-x) \ln(1-x)$ is the entropy function. Comparing (1), we need to set \bar{d} such that

$$\frac{\kappa^2 \bar{d} \sigma \tau}{3} > \delta + h(\sigma) + h(\tau)$$

for all $\varepsilon < \sigma, \tau < 1 - \varepsilon$. Since $h(x) \leq \ln 2$, setting

$$\bar{d} > \frac{3(\delta + 2 \ln 2)}{\kappa^2 \varepsilon^2}$$

suffices to beat the union bound over all pairs of sets S, T for any fixed σ, τ . Since the failure probability is exponentially small, it also easily covers the union bound easily covers all $O(n^2)$ possible values of σ and τ . \square

A.2 Inferring expansion of induced subgraphs

Lemma A.7 implies that, with high probability in $G \in \mathcal{G}_{n, \bar{d}}$, any sufficiently large set is “edge expanding”, with a large number of edges connecting it to its complement. The same holds in any induced subgraph of G . Consider the following definition.

Definition A.8. For $\rho > 0$ and a multigraph $G = (V, E)$, we say a subset $S \subset V$ is ρ -edge-expanding if $E(S, V \setminus S) \geq \rho|S|$. We say G is a ρ -edge-expander if all subsets $S \subset V$ with $|S| \leq |V|/2$ are ρ -edge-expanding.

Then Lemma A.7 states that in an induced subgraph of size n' , with high probability any set of size greater than εn is ρ -edge-expanding where

$$\rho = \bar{d}(1 - \kappa)n'/n \quad (4)$$

where we took $|T| \geq n'/2$ and thus $\tau \geq n'/(2n)$ in (3).

Intuitively, a graph where the only non-expanding subsets are small is close to an expander. Here we show that we can obtain an expander simply by removing the largest non-expanding subset.

LEMMA A.9. *Let $G = (V, E)$ be an undirected multi-graph of maximum degree \bar{d} . Let $\rho > 0$, and let $S \subset V$ be a subset of maximum cardinality such that $|S| \leq |V|/2$ and S is non- ρ -edge-expanding. Assume that $|S| \leq \mu|V|$ where $\mu < 1/4$. Then the subgraph G' induced by $V' = V \setminus S$ is an ρ' -edge-expander, where*

$$\rho' = \rho - 4\bar{d}\mu.$$

PROOF. To the contrary, assume there is some subset $T \subset V'$ with $|T| \leq |V'|/2$ which is not ρ' -edge-expanding: that is, $E(T, V' \setminus T) < \rho'|T|$.

There are two cases. If $|T| \leq (1/2 - \mu)|V|$ then $|S \cup T| \leq |V|/2$. However, this gives

$$E(S \cup T, V \setminus (S \cup T)) \leq E(S, V \setminus S) + E(T, V' \setminus T) < \rho|S| + \rho'|T| \leq \rho(|S \cup T|).$$

Thus $|S \cup T|$ is non- ρ -edge-expanding in V , contradicting the maximality of $|S|$.

In the other case, $|T| > (1/2 - \mu)|V| \geq |V|/4$. But this implies

$$\begin{aligned} E(T, V \setminus T) &= E(T, V' \setminus T) + E(S, T) \\ &\leq E(T, V' \setminus T) + \bar{d}|S| \\ &< \rho'|T| + \bar{d}\mu|V| \leq (\rho' + 4\bar{d}\mu)|T| \leq \rho|T|. \end{aligned} \quad (5)$$

Thus T is non- ρ -expanding in V . But since $|S| \leq \mu|V| < |T| \leq |V'|/2 \leq |V|/2$, this again contradicts the maximality of $|S|$. \square

Remark. We have chosen to write Lemma A.9 as a purely graph-theoretic result with as few assumptions as possible. In our setting, however, we can be considerably less pessimistic. First, we assumed in (5) that $E(S, T)$ could be as large as $\bar{d}|S|$, but the upper bound of Lemma A.7 shows that this is almost certainly not the case. Second, just as Lemma A.7 implies that $|S|/|V|$ is bounded by a small constant μ , the same holds for $|T|/|V'|$. Thus we can in fact exclude the second case where $|S \cup T| > |V|/2$ with high probability.

A.3 From edge expansion to vertex expansion

In order to bound the diameter of the subgraph induced by H' , we care more about vertex expansion—as in Definition 4.2—than about edge expansion. Even if a set S is edge-expanding, if many vertices outside S have large in-degree, then many of S 's outgoing edges could be incident to just a few vertices.

The next lemma shows that if $\mathcal{D}(v) \leq C/n$ for all v as in the statement of Theorem 4.3, then the set of vertices v with in-degree more than $2C\bar{d}$, and thus total degree more than $d_v + 2C\bar{d} \leq 3C\bar{d}$, is almost always a small fraction of the graph. By condemning these vertices to spend eternity with the adversary, we are left with a subgraph where no vertex has degree greater than $3C\bar{d}$. If S has m edges connecting it to its complement in that subgraph, we have $|\partial(S)| \geq m/(3C\bar{d})$. Thus if that subgraph is an α -edge-expander, it is also a γ -expander where $\gamma = \alpha/(3C\bar{d})$.

LEMMA A.10. *Let $\delta > 0$ and $\varepsilon > 0$ and assume that $\mathcal{D}(v) \leq C/n$ for all v . Then for sufficiently large \bar{d} , with probability $1 - e^{-\delta n}$ in $G \in \mathcal{G}_{n, \bar{d}; \mathcal{D}}$ at most εn vertices have in-degree greater than $2C\bar{d}$.*

PROOF. Since the outgoing neighbors of each vertex in $\mathcal{G}_{n, \bar{d}; \mathcal{D}}$ are chosen independently and the total out-degree is $\bar{d}n$, the in-degree of any given vertex v is distributed as $\text{Bin}(\bar{d}n, \mathcal{D}(v))$, which is stochastically dominated by $\text{Bin}(\bar{d}n, C/n)$. Using the Chernoff bound (2) with $t = \mu \leq C\bar{d}$, the probability this exceeds $2C\bar{d}$ is at most $e^{-C\bar{d}/3}$. Since the total in-degree is fixed, the number of such vertices is stochastically dominated by $\text{Bin}(n, e^{-C\bar{d}/3})$. Using

a crude union bound, the probability that there are more than εn such vertices is at most

$$\binom{n}{\varepsilon n} e^{-(C\bar{d}/3)\varepsilon n} \leq 2^n e^{-(C\bar{d}/3)\varepsilon n},$$

which is less than $e^{-\delta n}$ whenever $\bar{d} > (3/(\varepsilon C))(\delta + \ln 2)$. \square

A.4 Proof of Theorem 4.3

Now we put these pieces together.

PROOF OF THEOREM 4.3. We begin with the uniform case. Consider $G \in \mathcal{G}_{n, \bar{d}}$ where $c = C = 1$. Fix a subset $H \subseteq V$ such that $|H| \geq \alpha n$. Our task is to find a subset $H' \subseteq H$ of size at least $(\alpha - \beta)n$ whose induced subgraph is a γ -expander for some $\gamma > 0$.

First, we remove from H the vertices whose in-degree exceeds $2\bar{d}$. Using Lemma A.10, with high probability there are at most εn of these, and we set \bar{d} large enough so that $\varepsilon = \beta/2$. Denote the remaining set of vertices H'' and note that $|H''| \geq (\alpha - \beta/2)n$.

We now want to find a subset $H' \subseteq H''$ whose induced subgraph is an edge expander and where $|H'| \geq (\alpha - \beta)n$. Lemma A.9 tells us that we can do that by removing the largest subset S of H'' which is non- ρ -edge-expanding, as long as this set is of size at most μn where $\mu \leq \beta/2$. In fact we set $\mu = \beta/16$.

To bound the size of S we use the fact that H'' , like any induced subgraph of G , inherits the mixing property of Lemma A.7. We set \bar{d} large enough so that $\varepsilon = \mu = \beta/16$ and $\kappa = 1/2$ in that lemma. Then with high probability every set of size between εn and $|H''|/2$ is ρ -edge-expanding in H'' where (from (4))

$$\rho = (1 - \kappa)\bar{d}|H''|/n \geq \bar{d}(\alpha - \beta/2)/2.$$

Lemma A.9 then tells us that $H' = H'' \setminus S$ is a ρ' -edge-expander where

$$\rho' = \rho - 4\bar{d}\mu \geq \bar{d}(\alpha - \beta/2)/2 - \bar{d}\beta/4 = \bar{d}(\alpha - \beta)/2.$$

Since the maximum total degree of any vertex in H'' is at most $3\bar{d}$, H' is a γ -expander where

$$\gamma = \rho'/(3\bar{d}) = (\alpha - \beta)/6.$$

Finally, we set δ in both Lemmas A.7 and A.10 to be strictly larger than $\delta + \ln 2$ where δ is the desired parameter for p_{fail} in Theorem 4.3. Then the probability either of these lemmas fails is less than $2^{-n} e^{-\delta n}$. Finally, we take a union bound over all subsets $H \subseteq V$ with $|H| \geq \alpha n$, of which there are at most 2^n .

The nonuniform case of the theorem follows with minor alterations. Specifically, the mixing lemma follows because the number of independent random variables appearing in the Chernoff bounds there depend only on stake. The subgraph removal lemma translates immediately to a stake-weighted version and the remainder of the argument treats individual vertices. \square

Remark. Note that the expansion parameter γ stays constant as \bar{d} increases, since the edge expansion parameter α is proportional to \bar{d} .

B AN ANALYSIS OF NAKAMOTO CONSENSUS WITH ADVERSARIALLY SUPPRESSED MESSAGES.

Existing work on the analysis of Nakamoto consensus has focused on the setting with Δ -synchrony, where all message broadcasts are

delivered with adversarially determined delays of no more than Δ time steps. In this section, we show how to analyze a setting where the adversary may entirely eliminate messages of his choice from the protocol. We remark that this axis of adversarial activity does not have the same hard threshold at 1/2: indeed, there are natural settings where Nakamoto consensus can survive even if a majority of honestly sent messages never arrive.

B.1 The synchronous setting with message suppression.

We adopt the framework of characteristic strings and margin, developed in [22], to study the behavior of the longest chain rule under message suppression. Specifically, the analysis divides time into equal length time slots. While these time slots can be chosen to be small enough so that the possibility of multiple proof-of-work successes in a single slot can be neglected, it is convenient for our analysis to permit multiple proof-of-work discoveries in each time slot. We refer to the set of parties that discover proofs-of-work in a particular time slot as the “leaders” of that slot.

In this setting where we wish to explore the effects of message suppression, we identify two properties of interest: Δ -*current* leaders and Δ -*relayed* leaders. In the context of a particular execution of the blockchain protocol these properties—which are only applied to honest parties—have the following meaning: If a party is Δ -current at a particular time slot t then she holds a chain at least as long as any chain broadcast by a Δ -relayed party at time slot $t - \Delta$ or earlier. If a party is both Δ -current and Δ -relayed, we say that it is Δ -synchronized. When the network delay Δ can be inferred from context, we simply say *current*, *relayed*, and *synchronized*. Intuitively, any block produced by a “relayed” leader is successfully transmitted to the network layer; likewise, any “current” leader is aware of all blocks recently transmitted to the network layer.

With any fixed execution, we may associate a *characteristic string*

$$w_1 w_2 \dots$$

indicating certain properties of the sequence of parties elected to be slot leaders (and hence block producers). Specifically, each w_i is a symbol $w_i \in \Sigma = \{A, C, R, \mathbb{C}_R, \perp\}$. These symbols have the following interpretation:

- A indicates a slot with an adversarial leader, or with more than one leader;
- \mathbb{C}_R indicates a slot with a synchronized leader;
- C indicates a slot with a current leader;
- R indicates a slot with a relayed leader;
- \perp indicates a slot with no leader.

When $w_i \in \{\mathbb{C}_R, C, R\}$, we say that the index (or, alternatively, “slot”) i is *honest*; we further distinguish honest indices according to which of these symbols actually arises: if $w_i \neq C$, we say that the index is *relayed*; if $w_i = R$, we say that the index is *current*; if an index is both current and relayed, we say that it is *synchronized*. If $w_i = A$, we say that the index is *adversarial*.

We remark that this framework is “first order” modeling, in the sense that it simply treats any slot with multiple leaders as adversarial. Generalizing the model to handle more sophisticated

leadership patterns in a single slot is straightforward, but complicates the case analyses of the proofs below. Furthermore, for typical settings where slots are short with respect to network delays the major improvements in analysis arises from careful treatment of the leadership pattern *across* slots within Δ of each other (rather than careful treatment of the pattern in a single slot).

B.2 Fork notation, closure.

We adopt the modeling infrastructure developed in [11, 22] for reasoning about the consistency and liveness of Nakamoto consensus algorithms. Specifically, the relevant state of a blockchain algorithm is reflected with the following graph-theoretic object, the Δ -fork.

Definition B.1 (Δ -fork). Let Δ be a positive integer and $L \in \mathbb{N}$. A Δ -fork for the string $w \in \Sigma^L$ is a directed, rooted tree $F = (V, E)$ with a labeling

$$\ell : V \rightarrow \mathbb{N}$$

satisfying the axioms below. Edges are directed “away from” the root so that there is a unique directed path from the root to any vertex. The value $\ell(v)$ is referred to as the *label* of v , and takes values in the set $\{i \mid w_i \neq \perp\} \cup \{0\}$. We apply the same terminology to vertices as to the indices that label them; that is, the vertex v is respectively *adversarial*, *honest*, *relayed*, *current*, or *synchronized* if this is true of $\ell(v)$.

- (i) the root $r \in V$ has label $\ell(r) = 0$ and is considered honest and relayed by fiat;
- (ii) the sequence of labels $\ell()$ along any directed path is strictly increasing;
- (iii) if $w_i \in \{\mathbb{R}, C, R\}$, there is a unique vertex v for which $\ell(v) = i$;
- (iv) for any pair of vertices v, w for which v is relayed, w is current, and $\ell(v) + \Delta \leq \ell(w)$, their depths $\text{len}(v)$ and $\text{len}(w)$ satisfy $\text{len}(v) < \text{len}(w)$.

We write $F \vdash_{\Delta} w$ to indicate that F is a Δ -fork for the string w . When $\Delta = 1$, corresponding to the synchronous case, we may just write $F \vdash w$. If $F' \vdash_{\Delta} w'$ for a prefix w' of w , we say that F' is a *subfork* of $F \vdash w$, denoted $F' \sqsubset F$, if F contains F' as a consistently-labeled subgraph. We remark that with any prefix w' of w there is a unique maximal subfork $F' \vdash w'$ of a given fork $F \vdash w$ (given by all vertices labeled with the indices of w').

Tines. A path in a fork F originating at the root is called a *tine*. (Note that tines do not necessarily terminate at a leaf.) For a vertex v in F , $F(v)$ denotes the tine in F terminating in v . Given this one-to-one correspondence between vertices and tines of a fork, we routinely overload notation so that it applies to both tines and vertices. For example, we let $\text{len}(T)$ denote the *length* of the tine T , equal to the number of edges on the path; recall that $\text{len}(v)$ also indicates the depth of the vertex v . If we must identify the fork from which v is drawn, we sometimes write $\text{len}_F(v)$. We further overload $\text{len}()$ to apply to forks: $\text{len}(F)$ denotes the length of the longest tine in a fork F . A tine is called *honest* if it terminates in an honest vertex; we likewise apply the terms *adversarial*, *current*, *relayed*, and *synchronized*.

For two tines T, T' of a fork F , we write $T \sim_{\ell} T'$ if the two tines share a vertex with a label greater or equal to ℓ . Intuitively, $T \sim_{\ell} T'$ guarantees that the respective blockchains agree on the state of

the ledger up to time ℓ . Observe that \sim_{ℓ} is a “partial equivalence relation”: it is symmetric and transitive, but not necessarily reflexive (for example, $T \not\sim_{\ell} T$ if this tine has no blocks appearing in slots ℓ or larger). Looking ahead, the adversary can only make two parties disagree on the state of the ledger up to time ℓ if she makes them hold two chains corresponding to tines for which $T \not\sim_{\ell} T'$.

Fork trimming; dominance. For a characteristic string $w = w_1 \dots w_n \in \Sigma^n$ and a positive integer k , we let $w_{\uparrow k} = w_1 \dots w_{n-k+1}$ denote the string obtained by removing the last $k - 1$ symbols. When $|w| < k$, $w_{\uparrow k} = \epsilon$. For a fork $F \vdash_{\Delta} w_1 \dots w_n$ we let $F_{\uparrow k} \vdash_{\Delta} w_{\uparrow k}$ denote the fork obtained by retaining only those vertices labeled from the set $\{0\} \cup \{1, \dots, n - k + 1\}$. Observe that relayed tines appearing in $F_{\uparrow \Delta}$ are those that are necessarily visible to current players at a round just beyond the last one described by the characteristic string. We say that a tine T in F is Δ -dominant if $\text{len}(T) \geq \text{len}(\overline{F}_{\uparrow \Delta})$ and simply call it *dominant* if Δ is clear from the context.

B.3 Advantage, reach, and margin

We generally focus our analysis on forks that represent the view(s) of honest parties during the protocol. This places special emphasis on two properties defined next.

- Definition B.2 (Honest and relayed forks and subforks).*
- We say that a fork F is *honest* if every leaf is honest. In general, the *honest subfork* of a fork F is the maximal honest subfork of F .
 - We say that a fork F is *relayed* if every leaf is relayed. The *relayed subfork* of a fork F is the maximal relayed subfork of F . We let \overline{F} denote the relayed subfork of the fork F .

For a Δ -fork $F \vdash_{\Delta} w$, we define the *advantage* of a tine $T \in F$ as

$$\alpha_F(T) = \text{len}(T) - \text{len}(\overline{F}_{\uparrow \Delta}).$$

When F can be inferred from context, we remove the subscripts. Observe that $\alpha_F^{\Delta}(T) \geq 0$ if and only if T is Δ -dominant in F .

We pause to remark that our treatment differs somewhat from that developed in [22]. They consider a notion of “reserve” for a tine, which intuitively adds to that tine’s advantage any yet unused adversarial slots; with this addition, the analysis can focus entirely on *honest subforks* forks (which are called closed forks in [22]). In our current setting this mechanism is less convenient due to the more complicated role played by the various classes of honest parties. We thus opt for a treatment that considers arbitrary forks, rather than honest subforks—this somewhat complicates the analysis, but permits us to work with a simpler notion of advantage.

For $\ell \geq 1$, we define two quantities of interest

$$\rho(F) = \max_T \alpha_F^{\Delta}(T), \quad (6)$$

$$\mu_{\ell}^{\Delta}(F) = \max_{T \not\sim_{\ell} T^*} \min(\alpha_F^{\Delta}(T), \alpha_F^{\Delta}(T^*)), \quad (7)$$

this maximum extended over all pairs of tines (T, T^*) for which $T \not\sim_{\ell} T^*$. Note that there might exist multiple such pairs in F , but under the condition $\ell \geq 1$ there will always exist at least one such pair, as the trivial tine T_0 containing only the root vertex satisfies $T_0 \not\sim_{\ell} T$ for any T and $\ell \geq 1$; in particular $T_0 \not\sim_{\ell} T_0$. For this reason, we will always consider μ_{ℓ}^{Δ} only for $\ell \geq 1$. Intuitively, $\alpha_F^{\Delta}(T)$ captures the length advantage (or deficit) of the tine T against

the longest relayed tine created at least Δ slots before the upcoming slot, and hence now known to all current parties. Consequently, $\mu_\ell^\Delta(F)$ records the maximal advantage of any pair of tines F that agree only prior to ℓ , where “advantage” in this setting is treated as the minimum of the two tines.

The quantity $\rho(F)$ is the *reach* of F ; the quantity $\mu_\ell(F)$ is the *margin* of F . We say that a tine T^* is a *witness* for the reach $\rho(F)$ when it achieves the maximum of definition (6); we use the same terminology for margin: a pair of tines T and T' *witness* the margin $\mu_\ell(F)$ if they achieve the maximum of the definition (7).

We remark that one may always chose a pair of tines (T, T^*) that witness $\mu_\ell(F)$ for which one is, additionally, a witness for reach. To see this, consider any witness T^* of $\rho(F)$ and a pair (T, T') that witness $\mu_\ell(F)$; as $T \not\sim_\ell T'$, we must have either $T^* \not\sim_\ell T$ or $T^* \not\sim_\ell T'$ and hence T^* can replace one of these without reducing the margin quantity.

We overload the notation and let

$$\mu_\ell^\Delta(w) = \max_{F \vdash_\Delta w} \mu_\ell^\Delta(F).$$

We again emphasize the difference between this approach and that of [22], which maximizes only over honest forks.

The crucial property motivating these definitions is that $\mu_\ell^\Delta()$ provides explicit control over consistency failure events. This is reflected in the lemma below, which follows directly from the analogous formulation and proof in [3, 16].

LEMMA B.3. *Consider the sequence of forks $F_1 \vdash w_1, F_2 \vdash w_1 w_2, \dots$ associated with each step of an execution corresponding to a characteristic string $w = w_1 w_2 \dots$. Consider a tine T held by a current honest party in round r , which is hence Δ -dominant in F_r ; let ℓ be a round associated with a vertex (block) B that appears on T . If $\mu_\ell(w_1 \dots w_s) < 0$ for every $r \leq s \leq t$ then any Δ -dominant tine T^* of F_t contains the vertex B .*

Thus, in order to rule out consistency violations it suffices to establish that $\mu_\ell^\Delta(w) < 0$ for appropriate ℓ and $w_1 \dots w_{\ell+\Delta}$ (note that this bounds above $\mu_\ell^\Delta(F)$ for any relevant fork). Similar connections were originally established for PoS forks [22] with guaranteed network delivery.

B.4 An exact analysis in the serialized, synchronous setting.

We begin with an analysis of the quantity $\mu_\ell^\Delta(w)$ in the synchronous setting, corresponding to the case when $\Delta = 1$ and block creation is strictly serialized. In this synchronous case we abbreviate $\alpha_F^1()$ by $\alpha_F()$, and note that $\alpha_F(T) = \text{len}(T) - \text{len}(\bar{F})$. We similarly abbreviate μ_ℓ^1 by $\mu_\ell()$. In preparation for the statement of the theorem below, one final definition: we say that a characteristic string is *silent* if it contains no relayed symbols, and *relayed* otherwise.

THEOREM B.4. *Fix $\ell \geq 1$. By definition $\rho(\epsilon) = \mu_\ell(\epsilon) = 0$. In general,*

$$\begin{aligned} \rho(wA) &= \rho(w) + 1, \\ \rho(wR) &\begin{cases} = \max(\rho(w) - 1, 0) & \text{if } w \text{ is silent,} \\ = \rho(w) & \text{otherwise,} \end{cases} \\ \rho(wC) &\leq \rho(w) + 1, \\ \rho(w\mathbb{C}_R) &= \max(0, \rho(w) - 1), \end{aligned} \quad (8)$$

and, turning to margin: $\mu_\ell(w) = \rho(w)$ for $|w| < \ell$, and, for $|w| \geq \ell$,

$$\begin{aligned} \mu_\ell(wA) &= \mu_\ell(w) + 1, \\ \mu_\ell(wR) &\begin{cases} = \mu_\ell(w) + 1 & \text{if } \mu_\ell(w) < 0, \\ = \mu_\ell(w) & \text{if } \mu_\ell(w) \geq 0 \text{ and } w \text{ is relayed,} \\ = \mu_\ell(w) - 1 & \text{if } \mu_\ell(w) \geq 0 \text{ and } w \text{ is silent,} \end{cases} \\ \mu_\ell(wC) &\begin{cases} = \mu_\ell(w) & \text{if } \mu_\ell(w) < 0, \\ \leq \mu_\ell(w) + 1 & \text{if } \mu_\ell(w) \geq 0, \end{cases} \\ \mu_\ell(w\mathbb{C}_R) &= \begin{cases} 0, & \text{if } \rho(w) > 0 \text{ and } \mu_\ell(w) = 0, \\ \mu_\ell(w) - 1, & \text{otherwise.} \end{cases} \end{aligned} \quad (9)$$

The proof is presented as a sequence of lemmas, the first establishes the relationship between ρ and μ_ℓ when $|x| < \ell$. The remaining lemmas handle each input symbol separately.

LEMMA B.5. *$\rho(\epsilon) = 0$ and, for any $x \in \Sigma^*$ and $|x| < \ell$, $\mu_\ell(x) = \rho(x)$.*

PROOF. The only possible fork for the empty string ϵ contains a single honest vertex; the only tine in this fork has advantage zero; hence $\rho(\epsilon) = 0$.

Let $F \vdash x$ be a fork for which $\rho(F) = \rho(x)$; let T^* be a tine witnessing $\rho(F)$. Observe that $T^* \not\sim_\ell T^*$ because $\ell > |x|$. Therefore, $\mu_\ell(F) \geq \rho(x)$ by considering the pair (T^*, T^*) as a witness for margin. As $\mu_\ell(F)$ can clearly not exceed $\rho(x)$, it follows that $\mu_\ell(x) = \rho(x)$. \square

LEMMA B.6 (A). *For any $w \in \Sigma^*$, $\rho(wA) = \rho(w) + 1$. Additionally, for any $w \in \Sigma^*$, $\mu_\ell(wA) = \mu_\ell(w) + 1$.*

PROOF. We first establish the bounds $\rho(wA) \geq \rho(w) + 1$ and $\mu_\ell(wA) \geq \mu_\ell(w) + 1$. Consider a fork $F \vdash w$ and a pair of tines, T and T^* , for which $\alpha(T^*) = \rho(F)$, $\alpha(T) = \mu_\ell(F)$, and $T \not\sim_\ell T^*$. We may construct a new fork $\text{ext } F \vdash wA$ from F by adding two new adversarial vertices, v and v^* , both labeled with the index of the additional A symbol. Each of the two tines T and T^* of F are then extended by adding a directed edge to one of these vertices: T is extended by the vertex v ; T^* is extended by v^* . Note that $\bar{F} = \text{ext } \bar{F}$. Then $\alpha_{\text{ext } F}(\text{ext } F(v^*)) = \alpha_F(T^*) + 1$ and $\alpha_{\text{ext } F}(\text{ext } F(v)) = \alpha_F(T) + 1$. Applying this transformation to a fork F for which $\rho(F) = \rho(w)$ shows that $\rho(wA) \geq \rho(w) + 1$. Similarly, applying this transformation to a fork F for which $\mu_\ell(F) = \mu_\ell(w)$ shows that $\mu_\ell(wA) \geq \mu_\ell(w) + 1$.

To complete the proof we establish the bounds $\rho(wA) \leq \rho(w) + 1$ and $\mu_\ell(w) \leq \mu_\ell(w) + 1$. Consider a fork $F \vdash wA$ and a pair of tines, T and T^* , for which $\alpha(T^*) = \rho(F)$, $\alpha(T) = \mu_\ell(F)$, and $T \not\sim_\ell T^*$. Let $\text{res } F \vdash w$ be the maximal subfork of F for the prefix w and let $\text{res}(T)$ and $\text{res}(T^*)$ the restrictions of T and T^* to $\text{res } F$. Observe that $\text{res } \bar{F} = \bar{F}$. Furthermore, the tines T and T^* , when restricted to $\text{res } F$, lose at most one vertex: thus $\text{len}(\text{res}(T)) \geq \text{len}(T) - 1$ and $\text{len}(\text{res}(T^*)) \geq \text{len}(T^*) - 1$. It follows that $\rho(\text{res } F) \geq \rho(F) - 1$ and $\mu_\ell(\text{res } F) \geq \mu_\ell(F) - 1$. Applying this transformation to a fork F for which $\rho(F) = \rho(wA)$ proves that $\rho(w) \geq \rho(wA) - 1$; similarly, applying this transformation to a fork F for which $\mu_\ell(F) = \mu_\ell(wA)$ proves that $\mu_\ell(w) \geq \mu_\ell(wA) - 1$, as desired. \square

LEMMA B.7 (R). For any $w \in \Sigma^*$,

$$\rho(wR) = \begin{cases} \max(\rho(w) - 1, 0) & \text{if } w \text{ is silent,} \\ \rho(w) & \text{otherwise.} \end{cases}$$

Additionally, for any $w \in \Sigma^*$ with $|w| \geq \ell$,

$$\mu_\ell(wR) = \begin{cases} \mu_\ell(w) + 1, & \text{if } \mu_\ell(w) < 0, \\ \mu_\ell(w), & \text{if } \mu_\ell(w) \geq 0 \text{ and } w \text{ relayed,} \\ \mu_\ell(w) - 1, & \text{if } \mu_\ell(w) \geq 0 \text{ and } w \text{ silent.} \end{cases}$$

PROOF. We consider the quantity $\rho()$ first and begin with the setting when w is silent. It is easy to confirm that (for silent w), $\rho(w) = |w|$ while $\rho(wR) = \max(|w| - 1, 0)$. When w is not silent, we wish to show that $\rho(w) = \rho(wR)$ and we first establish that $\rho(w) \geq \rho(wR)$. If $\rho(wR) = 0$, this follows from the fact that $\rho()$ is nonnegative. Otherwise, $\rho(wR) > 0$ and we let $F \vdash wR$ be a fork for which $\rho(F) = \rho(wR)$, T a tine witnessing $\rho(F)$, and v the vertex in F associated with the last symbol (R). Let $\text{res } F$ denote the fork restricted to the string w . As $\rho(F) > 0$, the relayed vertex v cannot appear on T ; thus T appears unchanged in $\text{res } F$, where its advantage can only increase by the removal of the relayed vertex v . It follows that $\rho(w) \geq \rho(wR)$, as desired. Conversely, we establish the inequality $\rho(wR) \geq \rho(w)$. Let $F \vdash w$ be a fork for which $\rho(F) = \rho(w)$ and let T be a tine witnessing $\rho(w)$. Let $\text{ext } F \vdash wR$ denote the fork formed from F by adding a single vertex v labeled by the last symbol (R) and attached by an edge to the root. As w is not silent, $\text{len}(\overline{F}) > 0$ and we conclude that $\text{len}(\overline{\text{ext } F}) = \text{len}(\overline{F})$; then $\rho(F) = \rho(\text{ext } F)$ and hence $\rho(w) \geq \rho(wR)$.

Turning now to $\mu_\ell(w)$ (with $|w| > \ell$), we begin with the two bounds

$$\mu_\ell(w) < 0 \Rightarrow \mu_\ell(wR) \geq \mu_\ell(w) + 1, \quad (10)$$

$$\mu_\ell(w) \geq 0 \Rightarrow \mu_\ell(wR) \geq \begin{cases} \mu_\ell(w) & \text{if } w \text{ is not silent,} \\ \mu_\ell(w) - 1 & \text{if } w \text{ is silent.} \end{cases} \quad (11)$$

Let $F \vdash w$ be a fork for which $\mu_\ell(w) = \mu_\ell(F)$ and let T^* and T be two tines witnessing $\mu_\ell(w)$; as usual we assume that $\rho(F) = \alpha(T^*) \geq \alpha(T) = \mu_\ell(w)$. If $\mu_\ell(w) < 0$, define $\text{ext } F \vdash wR$ to be the fork obtained from F by adding a single vertex v to the end of the tine T labeled with the index of the last symbol (R). This is a legal fork and, moreover, as $\alpha_F(T) < 0$ the new relayed vertex appears with depth no more than $\text{len}(\overline{F})$. Thus $\text{len}(\overline{\text{ext } F}) = \text{len}(\overline{F})$ and $\alpha_{\text{ext } F}(\text{ext } F(v)) = \alpha_F(T) + 1$. The two tines T^* and $\text{ext } F(v)$ of $\text{ext } F$ establish that $\mu(\text{ext } F) = \mu_\ell(w) + 1$ and hence (10) as desired. As for the case $\mu_\ell(w) \leq 0$, we again consider a fork $F \vdash w$ and construct a fork $\text{ext } F \vdash wR$ by addition of a single new vertex v attached directly to the root vertex and labeled by the index of the last symbol (R). Observe that $\text{ext } F$ is a legal fork and that $\text{len}(\overline{\text{ext } F}) = \max(1, \text{len}(\overline{F}))$. Unless w is silent, then, $\text{len}(\overline{\text{ext } F}) = \text{len}(\overline{F})$ and the tines T and T^* in $\text{ext } F$ witness $\mu_\ell(\text{ext } F) \geq \mu_\ell(F)$, which establishes the non-silent case of (11). If w is silent, the same construction establishes $\mu_\ell(wR) \geq \mu_\ell(\text{ext } F) \geq \mu_\ell(F) - 1$ as $\text{len}(\overline{\text{ext } F}) = 1$ (and $\text{len}(\overline{F}) = 0$), concluding the argument for (11).

To complete the argument, we establish the complementary bounds:

$$\mu_\ell(wR) > 0 \Rightarrow \mu_\ell(w) \geq \begin{cases} \mu_\ell(wR) & \text{if } w \text{ is relayed,} \\ \mu_\ell(wR) + 1 & \text{if } w \text{ is silent,} \end{cases} \quad (12)$$

$$\mu_\ell(wR) \leq 0 \Rightarrow \mu_\ell(w) \geq \mu_\ell(wR) - 1. \quad (13)$$

Consider a fork $F \vdash wR$ for which $\mu_\ell(F) = \mu_\ell(wR)$; let T and T^* be a pair of tines for which $\alpha(T^*) = \rho(F)$, and $\alpha(T) = \mu_\ell(F) = \mu_\ell(wR)$. Let v be the unique vertex labeled by the last symbol (R) and let $\text{res } F \vdash w$ be the fork obtained by removal of the vertex v . Observe that if $\mu_\ell(wR) > 0$ the vertex v cannot appear on either of the tines T or T^* , as $\text{len}(\overline{F}) < \text{len}(T) \leq \text{len}(T^*)$ (and v is a leaf in the fork F). Noting that the removal of v can only reduce the length of $\overline{\text{res } F}$, and that the two tines T and T^* appear in $\text{res } F$, we conclude that $\mu_\ell(w) \geq \mu_\ell(\text{res } F) \geq \alpha_{\text{res } F}(T) \geq \alpha_F(T)$. Observe, furthermore, that if w is silent then $\text{len}(\overline{\text{res } F}) = 0$ and $\text{len}(\overline{F}) > 0$; thus $\alpha_{\text{res } F}(T) \geq \alpha_F(T) + 1$. This establishes (12). We proceed to the last case (13), in which we again consider the construction $\text{res } F \vdash w$ above with the same notation. In this case, however, there is no guarantee that the leaf node v (labeled with the index of the last symbol (R)) does not appear on one of the two tines of interest, T and T^* . We let $\text{res } T$ denote the result of removing the vertex v from T , so that $\text{res } T = T$ if v does not appear on T , and $\text{res } T$ is the resulting prefix of T otherwise; we adopt the same notation $\text{res } T^*$ for the tine T^* . Thus, $\text{res } T$ and $\text{res } T^*$ are tines of $\text{res } F$. Clearly, the lengths of T and T^* can differ from their restricted counterparts $\text{res } T$ and $\text{res } T^*$ by no more than one and

$$\min(\text{len}_{\text{res } F}(\text{res } T^*), \text{len}_{\text{res } F}(\text{res } T)) \geq \min(\text{len}_F(T^*), \text{len}_F(T)) - 1.$$

Noting, as above, that $\text{len}(\overline{\text{res } F}) \leq \text{len}(\overline{F})$, we conclude that $\mu_\ell(w) \geq \mu_\ell(\text{res } F) \geq \mu_\ell(F) - 1$, as desired. \square

LEMMA B.8 (C). For any $w \in \Sigma^*$, $\rho(wC) = \rho(w) + 1$. Additionally, for any $w \in \Sigma^*$, with $\ell \leq |w|$,

$$\begin{aligned} \mu_\ell(wC) &\leq \mu_\ell(w) + 1, & \text{if } \mu_\ell(w) \geq 0, \\ \mu_\ell(wC) &= \mu_\ell(w), & \text{if } \mu_\ell(w) < 0. \end{aligned}$$

PROOF. Note that if $F \vdash wC$ then $F \vdash wA$ and it follows that $\rho(wC) \leq \rho(wA) \leq \rho(w) + 1$ by Lemma B.6. To establish $\rho(wC) \geq \rho(w) + 1$, consider a fork $F \vdash w$ and a tine T that witnesses $\rho(w)$. Construct the fork $\text{ext } F \vdash wC$ by adding a new vertex v to the end of the tine T labeled with the additional symbol (C). This is a legal fork for which $\text{len}(\overline{\text{ext } F}) = \text{len}(\overline{F})$ as the new vertex in $\text{ext } F$ is not relayed. As $\text{len}(v) = \text{len}(T) + 1$, we conclude that $\rho(wC) \geq \rho(\text{ext } F) = \rho(w) + 1$.

Turning to margin, we first observe that the statement of the lemma is a consequence of the following three inequalities, which we treat individually below.

$$\mu_\ell(wC) \geq \mu_\ell(w), \quad (14)$$

$$\mu_\ell(wC) \leq \mu_\ell(w) + 1, \quad (15)$$

$$\mu_\ell(wC) \leq \mu_\ell(w) \quad \text{if } \mu_\ell(wC) \leq 0. \quad (16)$$

To establish inequality (14) consider a fork $F \vdash w$ for which $\mu_\ell(F) = \mu_\ell(w)$ and define $\text{ext } F \vdash wC$ to be the fork obtained from F by adding a single vertex v , labeled by the index of the new symbol (C), so that it extends the longest tine of F . As the new vertex is not relayed, $\text{len}(\overline{\text{ext } F}) = \text{len}(\overline{F})$ and it follows that $\mu_\ell(\text{ext } F) \geq \mu_\ell(F)$, considering that any pair of tines in F exist in $\text{ext } F$. Thus $\mu_\ell(wC) \geq \mu_\ell(w)$, as desired.

The inequality (15) follows from the straightforward fact mentioned above: any fork $F \vdash wC$ is also a legal fork for wA . In particular, $\mu_\ell(wC) \leq \mu_\ell(wA) \leq \mu_\ell(w) + 1$ by Lemma B.6.

It remains to show (16)—the case when $\mu_\ell(wC) \leq 0$. Let $F \vdash wC$ be a fork for which $\mu_\ell(F) = \mu_\ell(wC)$ and let T and T^* be two tines in F that witness $\mu_\ell(F)$ so that $\alpha(T^*) = \rho(F)$ and $\alpha(T) = \mu_\ell(F) = \mu_\ell(wC) \leq 0$. We construct a fork $\text{res } F \vdash w$ by deleting the vertex v associated with the final symbol (C). Considering that v is not relayed $\text{len}(\text{ext } \bar{F}) = \text{len}(\bar{F})$. As in the previous arguments, we let $\text{res } T^*$ denote the result of removing the vertex v from T^* , so that $\text{res } T^* = T$ when v does not appear on T^* . We adopt the same notation for T . As the vertex v is current, $\text{len}(v) > \text{len}(\text{res } \bar{F}) = \text{len}(\bar{F})$ and it follows that v cannot appear on the tine T (as $\text{len}(T) \leq \text{len}(\bar{F})$ by assumption); thus $\text{res } T = T$. If $\text{len}(T^*) = \text{len}(\bar{F})$, the same argument implies: v cannot appear on T^* and hence $T^* = \text{res } T^*$. In this case, the pair of tines $\text{res } T^*$ and $\text{res } T$ witness $\mu_\ell(\text{res } F) \geq \mu_\ell(F)$, as desired. If $\text{len}(T^*) > \text{len}(\bar{F})$, the restriction $\text{res } T^*$ may have length one less than T but in any case $\text{len}_{\text{res } F}(\text{res } T^*) \geq \text{len}_{\text{res } F}(\text{res } T)$. Again, the two tines $\text{res } T$ and $\text{res } T^*$ of $\text{res } F$ witness $\mu_\ell(\text{res } F) \geq \alpha_{\text{res } F}(\text{res } T) \geq \mu_\ell(w)$, as desired. \square

B.5 The case when $\Delta > 1$ and the serialization mapping

We observe that the case when $\Delta > 1$ can be handled via a direct reduction to the synchronous case ($\Delta = 1$) using a standard technique developed in [11]. We remark that one could alternatively carry out more precise accounting, such as that developed in [15], at the cost of significantly more complex analysis.

One simple analysis adopts the following transformation. Call a symbol “left isolated” if it is preceded by \perp^Δ ; similarly define the notion of “right isolated” if it is followed by \perp^Δ and “doubly isolated” if it is both left and right isolated. Then consider the transformation

T , operating on characteristic strings, so that

$$\begin{aligned} \perp &\longrightarrow \perp \\ A &\longrightarrow A \\ C &\longrightarrow \begin{cases} C & \text{if left isolated,} \\ A & \text{otherwise,} \end{cases} \\ R &\longrightarrow \begin{cases} R & \text{if right isolated,} \\ A & \text{otherwise,} \end{cases} \\ \mathring{R} &\longrightarrow \begin{cases} \mathring{R} & \text{if doubly isolated,} \\ R & \text{if right isolated,} \\ C & \text{if left isolated,} \\ A & \text{if not isolated on either side.} \end{cases} \end{aligned}$$

It is straightforward to check that any Δ -fork for a characteristic string w is also a 1-fork for the characteristic string $T(w)$. This reduces the study of Δ -forks to the synchronous case as in [11].

B.6 The asymptotics

Finally, we consider the distribution placed on characteristic strings corresponding to the executions of interest: each symbol is i.i.d. with distribution of the form

$$\begin{aligned} p_R &= \Pr[w_i = R] = p_h^1 \cdot p_{\text{send}}(1 - p_{\text{receive}}), \\ p_C &= \Pr[w_i = C] = p_h^1 \cdot (1 - p_{\text{send}})p_{\text{receive}}, \\ p_{\mathring{R}} &= \Pr[w_i = \mathring{R}] = p_h^1 \cdot p_{\text{send}}p_{\text{receive}}, \\ p_\perp &= \Pr[w_i = \perp] = (1 - p_a)(1 - p_h), \\ p_A &= \Pr[w_i = A] = 1 - p_\perp - p_R - p_C - p_{\mathring{R}}, \end{aligned}$$

where p_h is the probability that a slot has an honest slot leader, p_h^1 is the probability that a slot has a unique honest slot leader, p_a is the probability that a slot has an adversarial slot leader, p_{send} is the probability that a block sent from a stake-weighted random honest party is received by the backbone in Δ time, and p_{receive} is the probability that a block known to a member of the backbone is received by a stake-weighted random honest party in Δ time.

Then the final probabilistic bounds on margin with this distribution follow directly from the analysis of [21], yielding the following consistency result.

LEMMA B.9. *Let $w = w_1, \dots, w_n$ be chosen according to the distribution above with the added assumption that $p_a < p_{\mathring{R}}$. Then for any ℓ and k for which $\ell + k \leq n$, $\Pr[\mu_\ell(w) \geq 0] = \exp(-\Theta(k))$.*