

Randomized View Reconciliation in Permissionless Distributed Systems

Ruomu Hou^{id}, Irvan Jahja^{id}, Loi Luu, Prateek Saxena, and Haifeng Yu

Abstract—In a *sybil attack*, an adversary creates many fake identities/nodes and have them join the system. Computational puzzles have long been investigated as a possible sybil defense: nodes that fail to solve the puzzle in time will no longer be accepted by other nodes. However, a malicious node can behave in such a way that it is accepted by some honest nodes but not other honest nodes. This results in different honest nodes having different *views* on which set of nodes constitute the system. Such *view divergence*, unfortunately, breaks the overarching assumption required by many existing security protocols. Partly spurred by the growing popularity of Bitcoin, researchers have recently formalized the above *view divergence* problem and proposed interesting solutions (which we call *view reconciliation protocols*). All existing view reconciliation protocols so far have a similar $\Theta(N)$ time complexity, with N being the number of honest nodes in the system. As this paper’s main contribution, we propose a novel view reconciliation protocol whose time complexity is only $\Theta(\ln N / \ln \ln N)$. To achieve such an exponential improvement, we aggressively exploit randomization. The hidden constant factor in the asymptotic complexity of our protocol, however, is considerably larger than in previous protocols. Concrete numerical comparisons show that our protocol is more suitable for large-scale systems, while existing protocols are better for smaller-scale systems.

Index Terms—Network security, protocols, peer-to-peer-computing.

I. INTRODUCTION

SYBIL attack and computational puzzles. The *sybil attack* [2] problem has attracted tremendous amount attention from researchers over the past two decades. In a sybil attack, the adversary creates a large number of fake identities/nodes and have them join a system. The adversary can then leverage them to effectively launch various follow-up attacks. Sybil attacks exploit the fact that modern distributed

systems are typically *permissionless* and cannot rely on public key infrastructures, in the sense that they allow nodes to join the system without for example, going through an identity check based on credit card numbers to get permission to join.

Computational puzzles [3], [4] have long been investigated as an approach to defend against sybil attacks. With computational puzzles, every node is asked to complete some non-trivial computational task simultaneously. If a node fails to complete such a task within some specific time constraint, then the node’s identity will no longer be accepted by other nodes. If the attacker has limited computational power per time unit (e.g. only 10%) as compared to the aggregate computational power per time unit of all the N honest nodes, then the number of malicious nodes it can sustain in a sybil attack will also be limited (e.g., at most $0.1N$ malicious nodes). The developments of Bitcoin [5] in recent years have spurred much renewed interest in this well-known approach.

View divergence. Using computational puzzles, by itself, is yet not sufficient to offer a *security basis* for many permissionless distributed systems, for the following reason. In permissionless distributed systems, there is often no trusted central authority. Thus each honest node u will need to determine itself whether another node v should be accepted as part of the system. A malicious node v may then cooperate with u for the verification of its solution to the computational puzzle, but not with another honest node w . This results in different honest nodes having different *views* on which set of nodes should form the system. A further complication is that v (which is malicious) will also have its own view, and v can include many other sybil nodes in its view. At this point, we end up dealing with a (potentially infinite) set of nodes each with their own view about the membership of the system. An example scenario is given in Figure 1.

Such *view divergence*, unfortunately, breaks the basis of many security protocols in distributed systems: For example byzantine consensus protocols (e.g., [6]–[12]) and secure multi-party computation protocols (e.g., [13]–[15]) typically assume an agreed-upon set of nodes running the protocol. As another example, protocols for building dynamic and robust overlay networks [16], [17] typically have a bootstrapping stage during which they also assume an agreed-upon set of nodes running the protocol.


State-of-the-art protocols. Partly spurred by the growing popularity of Bitcoin [5], researchers have recently formalized [18], [19] the above view divergence problem, and proposed interesting solutions. Specifically, in CRYPTO 2015, Andrychowicz and Dziembowski [18] proposed a protocol

Manuscript received May 13, 2018; revised October 23, 2018 and February 19, 2020; accepted April 27, 2020; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor Y. Guan. Date of publication July 16, 2020; date of current version October 15, 2020. This work was supported in part by the Singapore Ministry of Education Academic Research Fund Tier-1 under Grant T1_251RES1616 and in part by DSO, Singapore, under Grant R-252-000-661-592. A preliminary conference version of this article is published in IEEE International Conference on Computer Communications (INFOCOM), April 2018. (*The authors of this paper are alphabetically ordered.*) (*Corresponding author: Irvan Jahja.*)

Ruomu Hou, Irvan Jahja, Prateek Saxena, and Haifeng Yu are with the School of Computing, National University of Singapore, Singapore 117417 (e-mail: houromu@comp.nus.edu.sg; irvan@comp.nus.edu.sg; prateeks@comp.nus.edu.sg; haifeng@comp.nus.edu.sg).

Loi Luu was with the School of Computing, National University of Singapore, Singapore 117417. He is now with Kyber Network, Singapore 059818 (e-mail: loiluu@kyber.network).

Digital Object Identifier 10.1109/TNET.2020.3004834



(a) Each node solves its computational puzzle to obtain the puzzle solution.

(b) Each puzzle solution is then sent to the other nodes.

(c) Each node constructs its view after verifying the solutions it received.

Fig. 1. An example of how using computational puzzle leads to view divergence. Consider two honest nodes A and B , and one malicious node C . Each node obtains a solution by solving a computational puzzle (Figure 1a). The puzzle solution is then sent to other nodes (Figure 1b). However, a malicious node C may instead decide to send its solution to A but not to B (Figure 1b). This results in honest nodes A and B having different views on which set of nodes should form the system (Figure 1c). Furthermore, the malicious node C will also have its own view, which it may construct arbitrarily (Figure 1c).

enabling the honest nodes to agree on a final view which is guaranteed to contain all N honest nodes and no more than a certain threshold number ($f \cdot N$) of malicious nodes. Katz *et al.* [19] proposed another protocol offering similar functionalities. For convenience, we refer to all such protocols as *view reconciliation* protocols – namely protocols that enable all the nodes, starting from divergent views, to obtain the same final reconciled view. While the existing view reconciliation protocols are quite different from each other in terms of algorithmic techniques, all of them [18], [19] require $\Theta(N)$ rounds of execution time. Here each *round* needs to be long enough to accommodate message propagation delay in the network as well as some amount of local processing.

Our central result. As our main contribution, this paper proposes and analyzes a novel randomized view reconciliation protocol, called RandomizedViewReconciliation or RVR. We prove that RVR only takes $\Theta(\frac{\ln N}{\ln \ln N})$ rounds to terminate.¹ Relatively speaking, our protocol is an exponential improvement over the state-of-the-art $\Theta(N)$ protocols [18], [19].

The hidden constant factor in the asymptotic complexity of RVR, however, is considerably larger than in previous protocols. In our concrete numerical comparisons, RVR outperforms other protocols [18], [19] in large-scale systems (e.g., with 10k nodes), and the advantage of RVR increases with the system size. On the other hand, some existing protocol [19], despite its linear round complexity, requires fewer number of rounds than RVR in smaller-scale systems. As one example of large-scale permissionless system, currently Bitcoin already has about 10k daily-active full nodes [20]–[22]. As another reference point, Algorand [23] is a recently proposed Proof-of-Stake cryptocurrency from MIT. It has been successfully commercialized [24] and attracted \$4 million dollars of initial funding [25]. In its experiments, Algorand has considered 5k to 50k users. Finally, we note that for permissionless systems relying on computational puzzles, having too small a scale (e.g., a few hundred nodes) is inherently undesirable, since it substantially reduces the malicious computation power needed for an external attacker to break the system.

Our RVR protocol currently assumes that the adversary's computational power is no more than the aggregate computational power of fN honest node, with f being any constant smaller than $\frac{1}{3}$. In comparison, [19] can tolerate all $f < 1$, while [18] can even tolerate $f > 1$ (though in most cases, to provide useful functionality later on, f needs to be below 1 anyway). Hence our protocol imposes a stronger restriction on the adversary's power than the existing protocols. It is our future work to investigate how to generalize to $f \geq \frac{1}{3}$.

On the other hand, despite requiring f to be smaller than $\frac{1}{3}$, we believe that RVR still constitutes an interesting and novel design choice in the solution space. For example, one potential application of RVR is for bootstrapping a robust overlay or for bootstrapping a genesis block in cryptocurrency protocols (see Section II). Cryptocurrency protocols often already have some restrictions on f . For example for its implementation, Algorand [23] assumes that the fraction of honest users is 80% (see Figure 3 of [23]), which corresponds to $f = \frac{1}{4}$ by our definition of f . Hence if used in Algorand, the assumption needed by RVR is already satisfied. As another example, the Elastico cryptocurrency protocol [26] uses $f = \frac{1}{3}$ as its parameter.² This is close to our requirement of f being any constant smaller than $\frac{1}{3}$.

Finally, our RVR protocol has similar communication complexity as previous protocols per round: Each honest node in our protocol sends on expectation $\Theta(N \ln N)$ bits per round,³ as compared to $\Theta(N)$ in previous protocols [18], [19]. But the total number of bits sent by each honest node in our protocol (i.e., $\Theta(N \frac{\ln^2 N}{\ln \ln N})$) is much smaller than in previous protocols (i.e., at least $\Theta(N^2)$).

Our approach. A key commonality among the existing view reconciliation protocols [18], [19] is that they are all deterministic (excluding invocations to crypto primitives). To decrease the time complexity, we aggressively exploit randomization and to allow a positive and tunable error probability δ . On the surface, having a positive error probability seems to be rather unacceptable for a security protocol. But even basic crypto primitives (such as public key crypto used in [18],

²The Elastico paper [26] states that $f = \frac{1}{4}$, but their definition of f is different from ours — when translated to our definition, their f is actually $\frac{1}{3}$.

³If the error probability δ (see definition later) is not viewed as a constant, then this will be $\Theta(N \ln \frac{N}{\delta})$.

¹If the error probability δ (see definition later) is not viewed as a constant, then this will be $\Theta(\frac{\ln N}{\ln \ln N} \ln \frac{1}{\delta})$.

[19]) have a positive probability of error: For example, the adversary might correctly guess the private key. This means that strictly speaking, even those previous protocols [18], [19] have positive error probability. The crux in our approach thus is to ensure that i) our error probability δ is tunable, and ii) δ decreases *exponentially* with the overhead of the protocol (as in the case of crypto primitives). Specifically in our RVR protocol, if δ were not viewed as a constant, then as δ approached zero, the time complexity of our RVR protocol would be $\Theta(\frac{\ln N}{\ln \ln N} \ln \frac{1}{\delta})$ rounds. Thus increasing the number of rounds decreases δ exponentially.

Allowing randomization opens up a lot of possibilities for better performance. For example, we will use computational puzzles to elect leaders, *probabilistically*. This differs from [18], [19] where the puzzles are used solely to challenge the computational power of the nodes. Our protocol will also rely on *randomized gossiping* and *randomized sampling* to achieve its end goal. While these techniques are largely commonly-used tricks in randomized algorithms, we show that combining these commonly-used techniques gives us a new view reconciliation protocol in the solution space.

Roadmap. Section II discusses related work, while Section III formalizes the view divergence problem. Section IV gives an overview of our RVR protocol, and Section V through VII elaborate on the key subroutines in RVR. Finally, Section VIII puts everything together and Section IX provides some numerical examples.

II. RELATED WORK

View reconciliation protocols. Computational puzzles [3], [4] have been investigated extensively for defending against sybil attacks in various contexts (e.g., [5], [27]–[29]). This work focuses on the view divergence problem that stems from the use of computational puzzles. Section I has already discussed the key differences between our RVR protocol and the two existing view reconciliation protocols [18], [19]. In addition, none of [18], [19] use randomization or the leader-based paradigm as we do, and none use any components similar to our three key subroutines.

Byzantine consensus protocols. Byzantine consensus protocols (e.g., [6]–[12]) are different but quite related to view reconciliation protocols. Both kinds enable the honest nodes to agree on something. But byzantine consensus protocols typically have the implicit assumption that the protocols are run over *some given set of nodes*. For example, a protocol typically assumes that there are $(1 + f)N$ nodes running the protocol, where the identities of these $(1 + f)N$ nodes are agreed upon by every node prior to the protocol. In our context, we do not have any given and agreed-upon set of nodes. In fact, establishing some agreed-upon set of $(1 + f)N$ nodes is *precisely the goal of the view reconciliation problem itself*. It is worth noting that [6] explicitly deals with sybil attacks. But the work still assumes an agreed-upon set of nodes, and explicitly leaves the implementation of such an assumption to future work.

While byzantine consensus protocols cannot solve the view divergence problem, one can nevertheless build upon ideas from them. For example, the two existing view reconciliation

protocols [18], [19] both build upon ideas from [9]. Our RVR protocol also builds upon a well-known leader-based paradigm [7], [12] widely used in byzantine consensus protocols. While this general paradigm has been well-known for decades, we implement such a paradigm under a setting where there is no agreed-upon set of nodes to start with.

Bitcoin. While Bitcoin [5] is an excellent example of permissionless large-scale distributed system, Bitcoin itself does not solve the view divergence problem. In fact, the recently discovered network-level *eclipse attack* [30] on Bitcoin is one ultimate consequence of view divergence. Using a view reconciliation protocol to bootstrap, and then dynamically maintaining a robust overlay [16], [17] will offer a provable defense against such eclipse attacks in Bitcoin.

Recently, researchers have used the Bitcoin protocol (or its variants) as a building block to solve other problems in the permissionless settings, such as for achieving consensus and electing a committee (e.g., [31]). Solving these problems in the permissionless settings, in turn, would enable one to solve the view divergence problem. However, such an approach necessarily inherits all the assumptions needed by the Bitcoin protocol. In particular, the Bitcoin protocol assumes the existence of a *genesis block*, which is unpredictable, agreed upon by all honest nodes, and released to all nodes simultaneously. As a direct contrast, our RVR protocol, as well as the previous view reconciliation protocols [18], [19], avoids such a critical assumption, by solving the view divergence problem directly rather than indirectly via Bitcoin.

Follow-up work. As an update, after the conference version [1] of this work had been published, and while this manuscript was under review for this journal, Aggarwal *et al.* [32] did some follow-up work to RVR. They sketched out an improved protocol that further reduces the time complexity to $O(1)$ rounds. But they did not provide full protocol details, formal proofs, or concrete numerical results. Their protocol also leverages randomization and builds upon some of the techniques in RVR. To achieve $O(1)$ complexity, they first elect a small committee, then reconcile the views of the committee members, and finally have other nodes adopt the view of the committee.

III. PROBLEM FORMULATION

Since the view divergence problem has already been previously formulated in [18], [19], this section will follow the existing formulation in [18], to the extent possible. The key notations used in this paper are summarized in Table I.

System model and attack model. The system has N *honest nodes* that always follows the protocol. The honest nodes do not know N , and the lack of such knowledge will be one technical difficulty that we overcome. Each honest node holds a (locally and randomly generated) public/private key pair, and the public key is the unique identity of the node. To simplify discussion, we will often simply use “node” to refer the public key of that node. Whenever a node sends a message, it includes its public key (as the sender’s “identity”) and then append a digital signature. Each node also has an IP address and a port number. Each honest node has one unit of computational power.

TABLE I
KEY NOTATIONS

N	Number of honest nodes in the system
u, v, w	Nodes in the system
f	The adversary's computational power is no more than the aggregate computational power of fN honest nodes
δ	Error probability of RVR
m	Number of hash operations that one unit of computational power can perform in one round
<code>max_hash</code>	The outputs of the hash functions will be viewed as integers in the range of $[0, \text{max_hash}]$
<code>init_view_u</code>	Initial view of node u
<code>final_view_u</code>	Final view of node u
<code>offset</code>	In RVR, the time between the first honest node starting and the last honest node starting is at most $(\text{offset} - 1)$ rounds

The adversary can have an unlimited number of *malicious nodes*, public/private key pairs, IP addresses and port numbers. Malicious nodes are byzantine and colluding. Independent of the number of malicious nodes, the adversary has total fN units of computational power, for some constant f . The adversary may spoof the IP addresses of the honest nodes, and can inject arbitrary messages into the network. It also sees all messages in the network. But the adversary cannot corrupt or remove a message from the network. The adversary cannot break crypto primitives such as digital signatures.

To allow a direct comparison, we follow [18], [19] and assume that the honest nodes are synchronized. Time is divided into rounds, where each round is sufficiently long for every honest node to do some local processing and for all messages sent at the beginning of the round to reach the corresponding receivers by the end of that round. Our protocol actually works even if the honest nodes are “out-of-sync” with each other when starting the protocol (see Section IV).

Hash functions and security parameter. We will use hash functions for various purposes in our protocol. When proving the formal guarantees, we follow [18] and model all our hash functions as random oracles [33]. The outputs of the hash functions will be viewed as integers in the range of $[0, \text{max_hash}]$. We use m to denote the number of hash operations that one unit of computational power can perform in one round. For simplicity, all asymptotic complexities in this paper treat the length of security keys, hashes, nonces, and so on, as $O(1)$.⁴

View reconciliation. In the *view reconciliation* problem, each node u has an initial view `init_viewu`, which is a set of nodes (i.e., the public keys of the nodes). The initial views are established using computational puzzles, via some existing approach (which we will review later). For all honest node u , `init_viewu` contains all the N honest nodes (since each honest node can properly solve one puzzle). Let `union_honest_init_view` be the union of the initial views of all the honest nodes. Then `union_honest_init_view`

contains at most fN malicious nodes. This constraints come from the adversary being able to solve only fN puzzles. Note that each puzzle solution corresponds to a single malicious node v , since the puzzle is tied to the public key of v . For malicious node v and honest nodes u and w , the adversary can cause v to be included only in `init_viewu` and not in `init_vieww`, thus causing view divergence.

The goal of *view reconciliation* is to give each honest node u a final view `final_viewu` such that i) the final views of all honest nodes are identical, and ii) the final view contains all the N honest nodes and at most fN malicious nodes. One may observe that the set `union_honest_init_view` already satisfies these properties. But it is difficult to determine `union_honest_init_view`, since we do not know which nodes are honest.

Establishing the initial views. The initial views on the nodes can be established using similar approaches in existing works [6], [18], [19]. This is not our contribution, not part of the view reconciliation problem, and not part of RVR. Hence, we only provide a concise review below.

First, there needs to be some mechanism for the nodes to initially learn about each other. This is addressed in [18], [19] by assuming the existence of some (insecure) *public channel*. Every node may post messages to the public channel. A message posted by an honest node will be received by all honest nodes, while a message posted by a malicious node may reach any subset (as chosen by the adversary) of the honest nodes. This public channel can be implemented by flooding.

Now if every node posts to the public channel its public key, IP address/port, together with a digital signature, then every node u will learn about all the honest nodes, together with some arbitrary number of malicious nodes. Let this set be `pre_init_viewu`. Next every u in the system will simultaneously send a random challenge to all nodes in `pre_init_viewu`. A receiving node v will combine all the challenges it receives (e.g., using a Merkle tree) into a single challenge. Node v will then solve a computation puzzle instantiated by the challenge, and then send back the puzzle solution (as well as the Merkle proofs) to each node u from which v received a challenge. Once u verifies the solution and the proof, u will include v in `init_viewu`. For simplicity, we will assume that an honest node v can always successfully

⁴If needed, one could use a security parameter λ , in which case the asymptotic time complexity of RVR and the previous protocols [18], [19] will not be affected, while the communication complexities in all cases will increase by a factor of $\Theta(\lambda)$.

Algorithm 1 RandomizedViewReconcile or RVR in short**input:** $f, \delta, \text{offset}, u, \text{init_view}_u, m$; **output:** final_view_u ;

```

1:  $\text{view}_u = \text{init\_view}_u$ ;
2: repeat  $6 \ln \frac{2}{\delta}$  times do
3:    $\text{leader} = \text{ProbLeaderElect}(f, \text{offset}, u, \text{init\_view}_u, m)$ ;
4:    $\text{score}_u[\cdot] = \text{TwoStageSample}(f, \delta, \text{offset}, u, \text{init\_view}_u, \text{view}_u)$ ;
5:   foreach  $v$  do
6:     if  $\text{score}_u[v] \geq 0.50$  then add  $v$  to  $\text{proposal}_u$ ;
7:     if  $\text{score}_u[v] \geq 0.75$  then add  $v$  to  $\text{over}_u$ ;
8:     if  $\text{score}_u[v] \leq 0.25$  then add  $v$  to  $\text{under}_u$ ;
9:    $\text{proposal} = \text{CoordinatedGossip}(f, \delta, \text{offset}, u, \text{view}_u, \text{init\_view}_u, \text{leader}, \text{proposal}_u)$ ;
10:   $\text{view}_u = \text{over}_u \cup (\text{proposal} \setminus \text{under}_u)$ ;
11:   $\text{offset} = 2$ ; //offset is 2 once gossiping is done
12: end
13:  $\text{final\_view}_u = \text{view}_u$ ;
14: return  $\text{final\_view}_u$ ;
```

solve the puzzle, while the adversary can solve total no more than fN puzzles here.

IV. OUR PROTOCOL: RANDOMIZED VIEW RECONCILIATION

This section provides an overview of our novel RandomizedViewReconcile (or RVR in short) protocol.

Messages with signatures. In RVR, each node u maintains a view (view_u) that can change over time. Initially, $\text{view}_u = \text{init_view}_u$, and at the end of the protocol, final_view_u will be set to view_u . Throughout the protocol, u only sends messages to and receives messages from node in its initial view init_view_u . A message always contains the sender, the receiver, a sequence number, and a digital signature, to enable authentication as well as to prevent replay attacks. A node u will discard all message received from nodes not in init_view_u , and also all messages whose sequence numbers are not expected. RVR has a certain number of iterations (see Algorithm 1), and the iteration number is also included in a message. In a particular iteration, a node will only use messages with the corresponding iteration number. To be concise, we will not explicitly mention these mechanisms in the pseudo-code or in the discussion.

Non-simultaneous start. Different honest nodes do not have to start running RVR at the same time. When we say that all honest nodes starts executing a certain protocol *within offset rounds of each other*, we mean that the time between the first honest node starting and the last honest node starting is at most $(\text{offset} - 1)$ rounds. Hence $\text{offset} = 1$ means simultaneous start. The value of offset will be fed into our protocol as a parameter.

Allowing $\text{offset} \geq 2$ enables our protocol to work even if the nodes are not perfectly in sync. Furthermore, even if all the nodes start simultaneously, due to the lack of knowledge of N , the nodes may get out-of-sync in the CoordinatedGossip subroutine. Hence we will need to allow offset in all other parts of our protocol anyway.

Overview of design. RVR has a certain number of iterations (see Algorithm 1). At a high level, we build upon a well-known leader-based paradigm [7], [12] (see Section II for discussion on related work): In each iteration some leader *proposes* a certain view to all the nodes. The protocol ensures that if the leader is honest, then all nodes will properly adopt the proposal. This adopted proposal (i.e., view) will already contain N honest nodes and at most fN malicious nodes, which means that it has all the desired properties. The protocol further ensures that in later iterations, even if the leader is malicious, the previously adopted proposal will not be disrupted — namely, the (bad) proposal made by a malicious leader will not be adopted (see proof of Theorem 8).

Three key subroutines. Our overarching approach in RVR to achieving small time complexity is to aggressively rely on randomization, and then carefully ensure that the error δ resulted from randomization can be decreased *exponentially*. Specifically, RVR uses three key subroutines, all of which are randomized. The first ProbLeaderElect subroutine aims to elect an honest leader in the system. Doing so perfectly will be challenging — we have unlimited number of malicious nodes in the system. But fortunately, the leader election here does not need to be always correct. Hence we elect the leader by simply having the nodes solve computational puzzles. Whoever solves the puzzle will become the leader. It is certainly possible to have multiple nodes solving the puzzles at roughly the same time, which would result in multiple leaders. This gets further complicated when the nodes do not start executing ProbLeaderElect simultaneously. Section V will show that with a proper puzzle difficulty, we can nevertheless ensure a positive constant probability of having a unique and honest leader.

The CoordinatedGossip subroutine allows the elected leader to disseminate its proposal to all honest nodes, via randomized gossiping. A difficulty here is that because the honest nodes do not know N , they may disagree on how many rounds they should gossip for. As a result, they may get out-of-sync with each other when they finish CoordinatedGossip. Making them

Algorithm 2 ProbLeaderElect**input:** f , offset , u , init_view_u , m ; **output:** leader;

```

1: generate fresh challengeu and send ⟨u, challengeu⟩ to all nodes in init_viewu;
2: spend offset rounds receiving messages in the form of ⟨w, challengew⟩;
3: form a Merkle tree with all the challengew's received as leaves, and let rootu be the root of the Merkle tree;
4: spend 6 · offset rounds trying to find x such that  $\frac{\text{hash}(x, u, \text{root}_u)}{\text{max\_hash}} \leq \frac{1}{6m(1+f)|\text{init\_view}_u| \cdot \text{offset}}$ ;
5: if such x were found, send ⟨x, u, rootu, off_path_hashes⟩ to all nodes in init_viewu;
6: spend offset rounds receiving messages in the form of ⟨y, w, rootw, off_path_hashes⟩;
7: leader = null;
8: foreach ⟨y, w, rootw, off_path_hashes⟩ received do
9:   if off_path_hashes validates properly and  $\frac{\text{hash}(y, w, \text{root}_w)}{\text{max\_hash}} \leq \frac{1}{6m|\text{init\_view}_u| \cdot \text{offset}}$  then leader = w;
10: return leader;

```

fully synchronous would require solving the classic *byzantine firing squad* problem [34]. Unfortunately, byzantine firing squad protocols typically assume that the protocols are run over *some given set of nodes* (which we do not have until after view reconciliation is complete). To overcome this difficulty, we will design a simple mechanism so that the nodes will return from CoordinatedGossip within 2 rounds (instead of 1 round, which would imply perfect synchrony) of each other.

Finally, in the leader-based paradigm, when the leader forms its proposal and when a node decides whether to adopt the leader's proposal, they need to collect the current views of all the honest nodes. Naturally in our setting, a node u will conceptually collect the views of all those nodes in init_view_u . Doing so directly entails large communication complexity, hence we instead randomly sample some of the nodes for their views. A tricky part is that the malicious nodes either may request for too many samples (hence blowing up communication complexity) or may push samples aggressively to other nodes (hence causing their views to be over-represented). We will use *two-stage sampling* in our TwoStageSample subroutine to overcome this issue.

Provable guarantees. Algorithm 1 gives the pseudo-code for the RVR. We defer a more detailed discussion of the pseudo-code and the formal analysis to Section VIII, after we elaborate the three key subroutines in Section V through VII.

V. PROBABILISTIC LEADER ELECTION

Conceptual design. Recall that RVR allows a δ probability of error. We say that ProbLeaderElect *succeeds* if it elects a unique leader and the elected leader is an honest node. As long as ProbLeaderElect succeeds with some constant probability (e.g., 0.16 as in our proof later), repeating it for $O(\ln \frac{1}{\delta})$ times can already guarantee at least one success with $1 - \delta$ probability. By our design of RVR, we only need one successful iteration anyway, to achieve our final goal.

This leads to the following design: We will exploit the limit on the adversary's computation power. We let all the nodes in the system solve computational puzzles. Whoever solves the puzzle will claim itself as the leader, by directly notifying all other nodes. ProbLeaderElect fails if i) a malicious node solves the puzzle, or ii) none of the honest nodes solve the puzzle, or iii) more than one honest nodes solve the puzzle. By

choosing a proper puzzle difficulty, we will be able to properly bound the probability of these three events away from 1.

Details of the design. Algorithm 2 gives the pseudo-code for ProbLeaderElect. To instantiate the computational puzzles, each node w generates a fresh random nonce as challenge_w at Line 1 and sends to every node in view_w . Each node u will collect challenges from all nodes in init_view_u , and then construct a standard Merkle tree with root root_u. Next u will spend $6 \cdot \text{offset}$ rounds (at Line 4) trying to find a solution x such that $\frac{\text{hash}(x, u, \text{root}_u)}{\text{max_hash}} \leq \frac{1}{6m(1+f)|\text{init_view}_u| \cdot \text{offset}}$. (We will explain this threshold later.) If u finds such x , it will notify every node w in init_view_u , while including a standard Merkle proof showing that challenge_w has been included in the Merkle tree, if any. Node w will validate (at Line 9) that the solution x is for a sufficiently hard puzzle, by checking that $\frac{\text{hash}(x, u, \text{root}_u)}{\text{max_hash}} \leq \frac{1}{6m|\text{init_view}_w| \cdot \text{offset}}$. The threshold $\frac{1}{6m|\text{init_view}_w| \cdot \text{offset}}$ here is such that regardless of the values of $|\text{init_view}_u|$ and $|\text{init_view}_w|$, a puzzle solution obtained at Line 4 on an honest node u can always be validated at Line 9 on an honest node w . Furthermore, the threshold $\frac{1}{6m|\text{init_view}_w| \cdot \text{offset}}$ is such that on expectation in the system, there will roughly be one honest node finding a puzzle solution. However, this threshold may allow the malicious nodes to solve a somewhat easier puzzle, which will be addressed when we prove the guarantees of the subroutine. **Provable guarantees.** Theorem 1 proves the complexity and success probability of ProbLeaderElect. Our proof will properly take care of the fact that the malicious nodes may i) solve easier puzzles, and ii) have more time to solve the puzzles compared to the honest nodes.

Theorem 1: Assume that $m \geq 1$, $N \geq 1000$, and $f < \frac{1}{3}$. If all honest nodes start running ProbLeaderElect within offset rounds of each other, then:

- Each honest node returns after exactly $8 \cdot \text{offset}$ rounds, while sending $\Theta(N \ln N)$ bits in each round.
- With probability at least 0.16, ProbLeaderElect returns the same leader on all honest nodes *and* the leader returned is an honest node.

Proof: We only prove the second claim since the first claim is trivial. First consider all the honest nodes. Let h denote the total number of hash operations performed at Line 4

Algorithm 3 CoordinatedGossip**input:** $f, \delta, \text{offset}, u, \text{init_view}_u, \text{leader}, \text{proposal}_u$; **output:** proposal;

```

1: if  $u \neq \text{leader}$  then proposal = null else proposal = proposalu (attached with a signature from u);
2: count = 0;
3: repeat  $\frac{3 \ln |\text{init\_view}_u|}{2 \ln \ln |\text{init\_view}_u|} + \text{offset}$  rounds do
4:   // each iteration here takes one round
5:   if proposal  $\neq$  null and  $|\text{proposal}| \leq (1+f)|\text{init\_view}_u|$  then
6:     send proposal (together with the attached signature) to  $8 \ln \frac{|\text{init\_view}_u|}{\delta}$  uniformly random nodes chosen from
       init_viewu;
7:   for each message msg received do
8:     if msg = msg_fin then increment count
9:     else if msg has leader's signature attached then proposal = msg;
10:  if count  $> \frac{f}{1+f} \cdot |\text{init\_view}_u|$  then break;
11: end
12: send msg_fin to all nodes in init_viewu;
13: wait until msg_fin is received (including msg_fin received at Line 7) from at least  $\frac{|\text{init\_view}_u|}{1+f}$  nodes;
14: return proposal;

```

across all N honest nodes. Hence we have $h = 6mN \cdot \text{offset}$. The threshold at Line 4 may range from $\frac{1}{6mN(1+f)^2 \cdot \text{offset}}$ to $\frac{1}{6mN(1+f) \cdot \text{offset}}$ on different u 's. Let p be the probability of a hash operation done by any honest node finding a solution x at Line 4. We thus have $\frac{1}{6mN(1+f)^2 \cdot \text{offset}} \leq p \leq \frac{1}{6mN(1+f) \cdot \text{offset}}$. Note that a solution x found by an honest node at Line 4 is guaranteed to pass the checking at Line 9 on every honest node w , since the threshold at Line 9 is $\frac{1}{6m|\text{init_view}_w| \cdot \text{offset}} \geq \frac{1}{6mN(1+f) \cdot \text{offset}}$.

Next consider the malicious nodes. Let h' denote the total number of hash operations done by the fN malicious nodes, in their attempts to find a solution to the puzzle. A malicious node will have at most $8 \cdot \text{offset}$ rounds to solve the puzzle.⁵ Thus $h' \leq 8fmN \cdot \text{offset}$. For a puzzle solution y from some malicious node w to be validated properly by honest node u at Line 9 of u 's execution, we need to have $\frac{\text{hash}(y, u, \text{root}_w)}{\max_hash} \leq \frac{1}{6m|\text{init_view}_u| \cdot \text{offset}}$. Hence $\frac{\text{hash}(y, u, \text{root}_w)}{\max_hash}$ needs to be no larger than $\frac{1}{6mN \cdot \text{offset}}$, in order for y to be possibly validated on *any* honest node. Let p' be the probability of a hash operation done by any malicious node finding a solution y such that $\frac{\text{hash}(y, \text{root}_w)}{\max_hash} \leq \frac{1}{6mN \cdot \text{offset}}$. We have $p' \leq \frac{1}{6mN \cdot \text{offset}}$.

Let \mathcal{E}_1 denote the event that ProbLeaderElect returns the same leader on all honest nodes with leader being an honest node, and let \mathcal{E}_2 denote the event that exactly one hash operation done by some honest node finds a solution at Line 4 and no other hash operations (done by honest or malicious nodes) find a solution. Based on the ranges of p and p' , we have: $\Pr[\mathcal{E}_1] \geq \Pr[\mathcal{E}_2] \geq h \times \left(\frac{1}{6mN(1+f)^2 \cdot \text{offset}}\right)^1 \times \left(1 - \frac{1}{6mN(1+f) \cdot \text{offset}}\right)^{h-1} \times \left(1 - \frac{1}{6mN \cdot \text{offset}}\right)^{h'}$.

Plugging in h and h' , and with some manipulation, one can eventually show that for $m \geq 1, N \geq 1000$,

and $f < \frac{1}{3}$, the above probability is at least: $\frac{1}{(1+f)^2} \cdot 0.36^{\frac{1}{1+f}} \cdot 0.36^{\frac{4f}{3}} > 0.16$. ■

VI. COORDINATED GOSSIPING

Conceptual design. CoordinatedGossip aims to disseminate a message (i.e., proposal_{leader}) from the leader to all honest nodes, after ProbLeaderElect has been invoked and succeeded. If ProbLeaderElect did not succeed, we will not care about the correctness of CoordinatedGossip (though we still care about its performance overhead).

The basic idea in CoordinatedGossip is simple and well-known: leader will first generate a signature on proposal_{leader}. Next in each round, each node u will relay proposal_{leader}, together with the signature, to $O(\log \frac{N}{\delta})$ nodes in u 's current view. Before relaying, a node needs to check proposal's size to avoid sending too many bits when the leader is malicious. A receiving node will verify leader's signature before accepting it. We will prove that $\frac{3 \ln N}{2 \ln \ln N}$ rounds is sufficient for all honest nodes to receive proposal_{leader} with probability close to 1. During such gossiping process, each node only sends $O(N \ln \frac{N}{\delta})$ bits per round. In comparison, directly having leader send proposal_{leader} to all nodes in its view would require leader to send $\Omega(N^2)$ bits. Gossiping has also been widely used for reducing the communication overhead in related contexts, for example, in distributed consensus protocols (e.g., [35]–[38]). Our approach here is similar, except that we are applying gossiping in a view reconciliation protocol under a permissionless setting.

The only difficulty in CoordinatedGossip is that the nodes do not know N , and hence do not know for how many rounds they should gossip. Each node u does have init_view_{*u*}, and can use $|\text{init_view}_u|$ in place of N for calculating $\frac{3 \ln N}{2 \ln \ln N}$. However, $|\text{init_view}_u|$ is different on different u 's. This means that two nodes u and w may spend different number of rounds in running CoordinatedGossip. A further problem is that CoordinatedGossip will be invoked multiple times in

⁵The malicious node may successfully obtain challenges from all the honest nodes at the very beginning, spend $8 \cdot \text{offset}$ rounds, and then finally send the puzzle solution right before the honest nodes stops receiving solutions.

RVR, making u and w more and more out-of-sync each time the subroutine is invoked. Making all the nodes return from CoordinatedGossip simultaneously would correspond to the classic *byzantine firing squad* problem [34]. Unfortunately, same as typical byzantine consensus protocols, byzantine firing squad protocols [34] also have the implicit assumption that the protocols are run over *some given set of nodes*.

In CoordinatedGossip, we do not make all the honest nodes return in the same round. Rather, we ensure that there exists some r , such that each honest node either returns in round r or in round $r + 1$. Achieving this will be much easier than solving the byzantine firing squad problem under our setting. Furthermore, our guarantee continues to hold even if the nodes do not start executing CoordinatedGossip simultaneously. This ensures that the “out-of-sync” will not accumulate over multiple invocations of the subroutine.

Details of the design. Algorithm 3 gives the pseudo-code for CoordinatedGossip. Our central goal here is to ensure that all honest nodes return within two rounds of each other. To achieve this, when a node u has finished its gossiping, it will send out a special `msg_fin` message to all nodes in `init_viewu`. Node u will then wait until it has received `msg_fin` from $\frac{|\text{init_view}_u|}{1+f}$ nodes in `init_viewu` (at Line 13), before u finally returns from the subroutine. Note that $\frac{|\text{init_view}_u|}{1+f} \leq N$ and `init_viewu` always contains N honest nodes. Hence as long as all N honest nodes send out `msg_fin`, u is guaranteed to eventually receive `msg_fin` from a sufficient number of nodes and return.

Let r be the round during which u returns. Consider any other honest node v . We want v to return no later than round $r + 1$. To achieve this, we let v break from the gossiping process and move on (at Line 10), as long as v has received `msg_fin` from more than $\frac{f}{1+f} \cdot |\text{init_view}_v|$ nodes in its view. We will be able to prove later that at least one node out of all these $\frac{f}{1+f} \cdot |\text{init_view}_v|$ nodes must be honest. Hence, the malicious nodes by themselves will not be able to trick v into breaking prematurely. At the same time, note that u 's returning in round r (at Line 13) implies that u has received `msg_fin` from $\frac{|\text{init_view}_u|}{1+f}$ nodes. We will be able to prove later that, out of these $\frac{|\text{init_view}_u|}{1+f}$ nodes, at least $(1-f)N$ nodes are honest. Hence v will also receive `msg_fin` from at least $(1-f)N$ nodes in its view in round r . Since $(1-f)N > fN \geq \frac{f}{1+f} \cdot |\text{view}_v|$ for all constants $f < \frac{1}{3}$, these `msg_fin` messages will be sufficient for v to break from the gossiping process (at Line 10). We will later prove that v will then return in the next round.

Provable guarantees. The following theorem presents the complexity of CoordinatedGossip, and further proves that i) all honest nodes will return within two rounds of each other from CoordinatedGossip, and ii) there will be a sufficient number of rounds during which all honest nodes are gossiping.

Theorem 2: Assume $f < \frac{1}{3}$. If all honest nodes start executing CoordinatedGossip within `offset` rounds of each other, then:

- 1) Each honest node sends $\Theta(N \ln \frac{N}{\delta})$ bits in each round of the execution;

- 2) Each honest node spends at most $\frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + (2 \cdot \text{offset}) + 1$ rounds before returning from the subroutine;
- 3) All honest nodes will return within 2 round of each other;
- 4) There are at least $\frac{3 \ln N}{2 \ln \ln N}$ rounds during which all honest nodes are executing the loop at Line 3.

Proof:

- 1) Obvious from the pseudo-code.
- 2) Every honest node u will eventually reach Line 12 and send `msg_fin` to all nodes, no later than $\frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + \text{offset} + 1$ rounds after u starts executing CoordinatedGossip. For any other given honest node v , u will send its `msg_fin` message no later than $\frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + (2 \cdot \text{offset}) + 1$ rounds after v starts executing CoordinatedGossip. There will be total N such u 's. Receiving finished messages from all of them will enable v to return.
- 3) Let round r be the first round during which one or more honest nodes return. Let u be one such honest node, where `init_viewu` contains N honest nodes and aN malicious nodes, for some a where $0 \leq a \leq f$. Consider any other honest node v . We claim that if v does not return in round r , then it must return in round $r + 1$. To see why, note that u must have received `msg_fin` (at Line 13) from $\frac{|\text{init_view}_u|}{1+f} = \frac{1+a}{1+f}N$ nodes by round r . At least $\frac{1+a}{1+f}N - aN$ such messages must be from honest nodes. For all $f < \frac{1}{3}$, we have $\frac{1+a}{1+f}N - aN = \frac{1-af}{1+f}N \geq (1-f)N > fN \geq |\text{view}_v| \frac{f}{1+f}$. This means that v must satisfy the condition at Line 10 by round r , and thus break from the loop. Once v breaks, in round $r + 1$, it will send out `msg_fin`. In fact, such argument applies to all honest nodes, and they will all have sent out `msg_fin` by round $r + 1$. All these $N \geq \frac{|\text{view}_v|}{1+f}$ `msg_fin` messages will be received by v and enable v to return in round $r + 1$.
- 4) For all honest node u , we have $N \leq |\text{init_view}_u| \leq (1+f)N$, and hence the number of iterations in the loop at Line 3 will be from $\frac{3 \ln N}{2 \ln \ln N} + \text{offset}$ to $\frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + \text{offset}$. Since all honest nodes starts executing CoordinatedGossip within `offset` rounds of each other, there will be at least $\frac{3 \ln N}{2 \ln \ln N}$ rounds during which all honest nodes are executing the loop, unless the condition at Line 10 is satisfied during these rounds. To prove that the condition is not satisfied, we need to show that those `msg_fin` messages from the malicious nodes will never be sufficient to satisfy the condition. Consider any honest node u where `init_viewu` contains N honest nodes and aN malicious nodes, for some a where $0 \leq a \leq f$. For the condition at Line 10, it is easy to verify that $\frac{f}{1+f} \cdot |\text{init_view}_u| = \frac{f}{1+f} \cdot (1+a)N \geq aN$. This means that the condition for u will never be satisfied by those `msg_fin` messages from those aN malicious nodes. ■

The next theorem proves the success probability of CoordinatedGossip:

Theorem 3: Assume that $N \geq 1000$, $\delta \leq 0.1$, and $f < \frac{1}{3}$. If i) all honest nodes start executing CoordinatedGossip within offset rounds of each other, ii) all honest nodes invoke CoordinatedGossip using the same leader parameter, iii) leader is an honest node, and iv) $|\text{proposal}_{\text{leader}}| \leq (1+f)N$, then with probability at least $1 - \frac{1}{4000}$, all honest nodes return $\text{proposal}_{\text{leader}}$.

Proof: First, because of the signature verification at Line 9, only the leader's proposal will ever be adopted at Line 9. Furthermore, because $|\text{proposal}_{\text{leader}}| \leq (1+f)N \leq (1+f)|\text{init_view}_u|$, a node that has adopted the leader's proposal will always send the proposal in Line 6. In any given round, we call an honest node as a *black node* if it has adopted the leader's proposal at Line 9. Otherwise the honest node is a *white node*. A malicious node is neither black nor white. Theorem 2 has proved that there are at least $\frac{3 \ln N}{2 \ln \ln N}$ rounds during which all honest nodes are executing the loop at Line 3. For convenience, in this proof, number these rounds as round 1 through $\frac{3 \ln N}{2 \ln \ln N}$. For any given round r between 1 and $\frac{3 \ln N}{2 \ln \ln N}$, define b_r to be the number of black nodes at the end of that round. We will separately prove the following two key equations:

$$\Pr[b_{\frac{\ln N}{\ln \ln N}} \geq 0.2N] \geq 1 - \frac{1}{8000} \quad (1)$$

$$\Pr[b_{\frac{3 \ln N}{2 \ln \ln N}} = N \mid b_{\frac{\ln N}{\ln \ln N}} \geq 0.2N] \geq 1 - \frac{1}{8000} \quad (2)$$

Combining these two via a union bound will directly lead to the theorem.

We first prove Equation 2, which is easier. Since $b_{\frac{\ln N}{\ln \ln N}} \geq 0.2N$, starting from the beginning of round $\frac{\ln N}{\ln \ln N} + 1$, in each round we have at least $0.2N$ black nodes. We call a black node's sending view' at Line 6 as a *push*. Each black node u does $8 \ln \frac{|\text{init_view}_u|}{\delta} \geq 8 \ln \frac{N}{\delta}$ pushes per round. Hence from round $\frac{\ln N}{\ln \ln N}$ to round $\frac{3 \ln N}{2 \ln \ln N}$, there will be at least total $0.2N \cdot 8 \ln \frac{N}{\delta} \cdot \frac{\ln N}{2 \ln \ln N} = 0.8 \cdot N \ln N \cdot \frac{\ln N}{\ln \ln N} > 3.6N \ln N$ pushes (since $N \geq 1000$ and $\delta \leq 0.1$). There are at most $0.8N$ white nodes at the end of round $\frac{\ln N}{\ln \ln N}$. For any such white node v and any given push, v is chosen as the target of the push with probability at least $\frac{1}{(1+f)N} > \frac{2}{3N}$. Hence the probability that v is never chosen as the target of any of the $3.6N \ln N$ pushes is at most $(1 - \frac{2}{3N})^{3.6N \ln N} = ((1 - \frac{2}{3N})^{\frac{3N}{2}})^{2.4 \ln N} < e^{-2.4 \ln N} = \frac{1}{N^{2.4}}$. Taking a union bound across all $0.8N$ possible v 's, we know that with probability at least $1 - \frac{0.8}{N^{1.4}}$, all white nodes have been chosen as the target of some push at least once by the end of round $\frac{3 \ln N}{2 \ln \ln N}$. Hence:

$$\begin{aligned} \Pr[b_{\frac{3 \ln N}{2 \ln \ln N}} = N \mid b_{\frac{\ln N}{\ln \ln N}} \geq 0.2N] \\ \geq 1 - \frac{0.8}{N^{1.4}} > 1 - \frac{1}{8000} \quad (\text{since } N \geq 1000) \end{aligned}$$

We next move on to Equation 1. We will later prove the following key equation for all r :

$$\Pr[b_{r+1} \geq \min(0.2N, b_r \ln \frac{N}{\delta})] \geq 1 - \frac{\delta^{1.125}}{N^{1.125}} \quad (3)$$

It is easy to prove Equation 1 via Equation 3. Consider rounds 1 through $\frac{\ln N}{\ln \ln N}$. With probability at least $1 - \frac{\delta^{1.125}}{N^{1.125}} \times \frac{\ln N}{\ln \ln N}$,

we have $b_{r+1} \geq \min(0.2N, b_r \ln \frac{N}{\delta})$ for all $1 \leq r \leq \frac{\ln N}{\ln \ln N}$. Since $(\ln \frac{N}{\delta})^{\frac{\ln N}{\ln \ln N}} > 0.2N$, this will immediately imply $b_{\frac{\ln N}{\ln \ln N}} \geq 0.2N$. Hence we have:

$$\begin{aligned} \Pr[b_{\frac{\ln N}{\ln \ln N}} \geq 0.2N] \\ \geq 1 - \frac{\delta^{1.125}}{N^{1.125}} \times \frac{\ln N}{\ln \ln N} \\ > 1 - \frac{1}{8000} \quad (\text{since } N \geq 1000 \text{ and } \delta \leq 0.1) \end{aligned}$$

Now the only missing part is the proof for Equation 3. If $b_r \geq 0.2N$, then b_{r+1} must be at least $0.2N$ and Equation 3 holds trivially. Now consider the case where $b_r < 0.2N$. At the beginning of round $r+1$, there are total b_r black nodes, doing total at least $8b_r \ln \frac{N}{\delta}$ pushes in round $r+1$. Order all these pushes into an arbitrary sequence, and it will be convenient to imagine that these pushes are done one by one sequentially. For each push, we say it is *effective* if the target of the push is a white node at the time of the push. Before the number of black nodes reaches $\min(0.2N, b_r \ln \frac{N}{\delta})$, the probability of a push being effective is at least $\frac{N-0.2N}{(1+f)N} \geq \frac{0.8}{1+f} > \frac{0.8}{1.5} > 0.5$. Having $b_r \ln \frac{N}{\delta}$ effective pushes is sufficient for the number of black nodes (and hence b_{r+1}) to reach $\min(0.2N, b_r \ln \frac{N}{\delta})$. Let random variable x denote the number of pushes needed (starting from the beginning of the above sequence) for the number of black nodes to reach $\min(0.2N, b_r \ln \frac{N}{\delta})$. If the number of black nodes never reaches this quantity after all the pushes in the sequence, x is defined to be ∞ .

In the next, we will reason about x via a simple coupling argument. Consider a separate experiment where we flip a fair coin for $8b_r \ln \frac{N}{\delta}$ times. Let random variable y denote the number of flips done, in order to encounter the first $b_r \ln \frac{N}{\delta}$ heads. (We define $y = \infty$ if we do not encounter so many heads even after all the flips.) We can easily couple x and y in such a way that for all a , we have $\Pr[x \leq a] \geq \Pr[y \leq a]$. Putting everything so far together, we have:

$$\begin{aligned} \Pr[b_{r+1} \geq \min(0.2N, b_r \ln \frac{N}{\delta})] &= \Pr[x \leq 8b_r \ln \frac{N}{\delta}] \\ &\geq \Pr[y \leq 8b_r \ln \frac{N}{\delta}] \end{aligned}$$

Now define random variable z to be the total number of heads observed when flipping a fair coin for $8b_r \ln \frac{N}{\delta}$ times. We draw a trivial connection between y and z , and then invoke a standard Chernoff bound on z :

$$\begin{aligned} \Pr[y > 8b_r \ln \frac{N}{\delta}] &= \Pr[z < b_r \ln \frac{N}{\delta}] \\ &= \Pr[z < (1 - 0.75)(4b_r \ln \frac{N}{\delta})] \\ &\leq \exp(-\frac{1}{2}(0.75)^2(4b_r \ln \frac{N}{\delta})) = \frac{\delta^{1.125}}{N^{1.125}} \end{aligned}$$

Hence we have $\Pr[b_{r+1} \geq \min(0.3N, b_r \ln \frac{N}{\delta})] \geq \Pr[y \leq 8b_r \ln \frac{N}{\delta}] \geq 1 - \frac{\delta^{1.125}}{N^{1.125}}$. ■

VII. TWO-STAGE SAMPLING

Conceptual design. Consider any given honest node u . For each node v , define $\text{frac}_u[v] = \frac{|\{w \mid w \in \text{init_view}_u \text{ and } v \in \text{view}_w\}|}{|\text{init_view}_u|}$ — namely, $\text{frac}_u[v]$ is

Algorithm 4 TwoStageSample()**input:** $f, \delta, \text{offset}, u, \text{init_view}_u, \text{view}_u$; **output:** $\text{score}_u[\cdot]$;

```

1: generate fresh  $\text{nonce}_u$  and send  $\langle u, \text{hash}(\text{nonce}_u) \rangle$  to all nodes in  $\text{init\_view}_u$ ;
2: spend  $\text{offset}$  rounds receiving messages in the form of  $\langle w, \text{hash}(\text{nonce}_w) \rangle$ ;
3: send  $\langle u, \text{nonce}_u \rangle$  to all nodes in  $\text{init\_view}_u$ ;
4: spend  $\text{offset}$  rounds receiving messages in the form of  $\langle w, \text{nonce}_w \rangle$ ;
5: validate each  $\text{nonce}_w$  received if  $\text{nonce}_w$  matches the previously received hash;
6: for each validated  $\text{nonce}_w$  do
7:   if  $\left( \frac{\text{hash}(u, \text{nonce}_u, w, \text{nonce}_w)}{\text{max\_hash}} \leq \frac{1+f}{|\text{init\_view}_u|} \left( \frac{30}{1-3f} \right)^2 \ln \frac{3(1+f)|\text{init\_view}_u|}{\delta} \right)$  and  $(|\text{view}_u| \leq (1+f)|\text{init\_view}_u|)$  then
8:     send  $\langle u, \text{view}_u \rangle$  to  $w$ ;
9: spend  $\text{offset}$  round receiving messages in the form of  $\langle w, \text{view}_w \rangle$ ;
10: validate each  $\text{view}_w$  received if  $\text{nonce}_w$  was validated and if  $\frac{\text{hash}(w, \text{nonce}_w, u, \text{nonce}_u)}{\text{max\_hash}} \leq \frac{1}{|\text{init\_view}_u|} \left( \frac{30}{1-3f} \right)^2 \ln \frac{3|\text{init\_view}_u|}{\delta}$ ;
11: for each  $v$  do
12:    $\text{votes}_u[v] = |\{w \mid v \in \text{view}_w \text{ and } \text{view}_w \text{ was validated}\}|$ ;
13:    $\text{score}_u[v] = \text{votes}_u[v] / \left( \left( \frac{30}{1-3f} \right)^2 \ln \frac{3|\text{init\_view}_u|}{\delta} \right)$ ;
14: return  $\text{score}_u[\cdot]$ ;
```

the fraction of the nodes in init_view_u that have v in their current views. Our TwoStageSample() subroutine enables u to estimate $\text{frac}_u[v]$ for all node v . We only obtain an estimate since otherwise u will need to retrieve $O(N)$ views where each view is of $O(N)$ size, which would result in $\Omega(N^2)$ communication complexity. Later we will prove that obtaining a (good enough) estimate is nevertheless still sufficient for everything to work.

The simplest way to estimate $\text{frac}_u[v]$ is perhaps for u to choose t random nodes from init_view_u . For each w chosen, node u will then *pull* from w , by requesting w to send view_w (which is considered one sample) to u . Node u can then estimate $\text{frac}_u[v]$ (for all v), based on the view_w 's received. While such a way of estimating does work, all the malicious nodes may pull from w and cause w to send $\Omega(N^2)$ bits, breaking the desirable guarantee on communication complexity. Instead of u *pulling* from w , an alternative approach is for w to choose t random nodes from init_view_w . For each u chosen, w will *push* to u , by sending view_w to u (without u requesting for it). This gives w control over the communication complexity incurred. But now a malicious node w will aggressively push view_w to all nodes, causing view_w to be over-represented.

Our TwoStageSample() overcomes the above problem via two stages. The first stage is the same as the push-based design above, where w pushes view_w to t other nodes. In the second stage, each receiving node u will decide whether to validate and use the received view_w , and use it for the estimation later. Node u validates view_w if it feels that w should indeed have pushed view_w to u . Two things are still needed to enable such design. First, w and u need to determine, in a consistent way, whether view_w should be a sampled (i.e., pushed) to u . We achieve this by letting w and u generate a shared random number. Second, the probability of view_w being a sample for u depends on N . Neither w nor u knows N , and they can only use $|\text{init_view}_u|$ and $|\text{init_view}_w|$ in place of N . But $|\text{init_view}_u|$ and $|\text{init_view}_w|$ may be different.

To resolve this issue, w will need to slightly over-push, by using a somewhat larger probability.

Details of the design. Algorithm 4 gives the pseudo-code for TwoStageSample(). For Line 1 to Line 5, each node u sends a fresh nonce nonce_u to all other nodes. To prevent a malicious node from carefully constructing its nonce after seeing other nodes' nonces first, the protocol requires each node to publish a commitment of its nonce first. Next, any given ordered pair $w \rightarrow u$ corresponds to two nonces, nonce_w and nonce_u . Node w will push to u iff i) w knows both nonces, and ii) $\frac{u, \text{hash}(\text{nonce}_u, w, \text{nonce}_w)}{\text{max_hash}} \leq \frac{1+f}{|\text{init_view}_u|} \left(\frac{30}{1-3f} \right)^2 \ln \frac{3(1+f)|\text{init_view}_u|}{\delta}$. Similarly, node v validates w 's push iff i) u knows both nonces, and $\frac{\text{hash}(\text{nonce}_w, w, \text{nonce}_u, u)}{\text{max_hash}} \leq \frac{1}{|\text{init_view}_u|} \left(\frac{30}{1-3f} \right)^2 \ln \frac{3|\text{init_view}_u|}{\delta}$. The two conditions on w and u controls how many samples are taken. Furthermore, they are designed in such a way that w always tries to over-push: Namely, the condition on w is always no harder to satisfy than the condition on u .

Provable guarantees. The following theorem summarizes the complexity of TwoStageSampling():

Theorem 4: Assume $f < \frac{1}{3}$. If all honest nodes start executing TwoStageSampling() within offset rounds of each other, then each honest node returns after exactly $3 \cdot \text{offset}$ rounds, while sending on expectation $\Theta(N \ln \frac{N}{\delta})$ bits per round.

Proof: The time complexity is obvious. For the communication complexity, note that for sending the nonce and the hashes of the nonce, it takes $O(N)$ bits per round. For sending the samples at Line 8, each sample (i.e., view_u) has size of $O(N)$ bits if they are sent. On expectation, each honest node sends $|\text{init_view}_u| \cdot \frac{1+f}{|\text{init_view}_u|} \left(\frac{30}{1-3f} \right)^2 \ln \frac{3(1+f)|\text{init_view}_u|}{\delta} = O(\ln \frac{N}{\delta})$ samples. Hence the total expected number of bits is $O(N \ln \frac{N}{\delta})$. ■

The next theorem proves that for any pair $w \rightarrow u$, u controls the probability of view_w being sampled. In particular, malicious nodes cannot force a higher probability on u .

Theorem 5: Assume $f < \frac{1}{3}$. Consider any given node w and any given honest node u . Define $p_u = \frac{1}{|\text{init_view}_w|} \left(\frac{30}{1-3f}\right)^2 \ln \frac{3|\text{init_view}_u|}{\delta}$. Define $p(w \rightarrow u)$ to be the probability that view_w is received and then validated by u at Line 10. If all honest nodes start `TwoStageSampling()` within `offset` rounds of each other, then $p(w \rightarrow u) = p_u$ for honest w and $p(w \rightarrow u) \leq p_u$ for malicious w .

Proof: First, a nonce sent by an honest node will always be validated by all other honest nodes. Let event \mathcal{E}_1 denote $\frac{\text{hash}(\text{nonce}_w, w, \text{nonce}_u, u)}{\text{max_hash}} \leq \frac{1+f}{|\text{init_view}_w|} \left(\frac{30}{1-3f}\right)^2 \ln \frac{3(1+f)|\text{init_view}_w|}{\delta}$, which corresponds to Line 8 in w 's invocation of the subroutine. Let event \mathcal{E}_2 denote $\frac{\text{hash}(\text{nonce}_w, w, \text{nonce}_u, u)}{\text{max_hash}} \leq \frac{1}{|\text{init_view}_u|} \left(\frac{30}{1-3f}\right)^2 \ln \frac{3|\text{init_view}_u|}{\delta}$, which corresponds to Line 10 in u 's invocation of the subroutine. Since $\frac{1}{|\text{init_view}_u|} \leq \frac{1+f}{|\text{init_view}_w|}$ and $\frac{3|\text{init_view}_u|}{\delta} \leq \frac{3(1+f)|\text{init_view}_w|}{\delta}$ always hold, we know that if \mathcal{E}_2 happens, \mathcal{E}_1 must happen.

If w is honest, then u will receive and validate view_w if and only if \mathcal{E}_2 happens. We thus have $p(w, u) = \Pr[\mathcal{E}_2] = \frac{1}{|\text{init_view}_u|} \left(\frac{30}{1-3f}\right)^2 \ln \frac{3|\text{init_view}_u|}{\delta}$. If w is malicious, for u to receive and validate view_w , \mathcal{E}_2 is a necessary (but not sufficient) condition. Hence $p(w, u) \leq \Pr[\mathcal{E}_2] = \frac{1}{|\text{init_view}_u|} \left(\frac{30}{1-3f}\right)^2 \ln \frac{3|\text{init_view}_u|}{\delta}$. ■

The next theorem proves that the estimates ($\text{score}_u[\cdot]$) returned by `TwoStageSampling()` satisfy some critical properties.

Theorem 6: Assume $f < \frac{1}{3}$. For any honest or malicious node v , define $\text{count}[v] = |\{u | v \in \text{init_view}_u \text{ and } u \text{ is honest}\}|$. If for all honest nodes u , $|\text{view}_u| \leq (1+f)N$ and they all start executing `TwoStageSampling()` within `offset` rounds of each other, then with probability at least $1 - \frac{16\delta^3}{81}$, none of the following (bad) events will happen:

- 1) There exists some node v and some honest node u such that $\text{count}[v] = N$ and $\text{score}_u[v] \leq 0.75$.
- 2) There exists some node v and some honest node u such that $\text{count}[v] = 0$ and $\text{score}_u[v] \geq 0.25$.
- 3) There exists some node v , some honest nodes u and w , such that $\text{score}_u[v] \geq 0.5$ and $\text{score}_w[v] \leq 0.25$.
- 4) There exists some node v , some honest nodes u and w , such that $\text{score}_u[v] \leq 0.5$ and $\text{score}_w[v] \geq 0.75$.

Proof: Define $s_u = \left(\frac{30}{1-3f}\right)^2 \ln \frac{3|\text{init_view}_u|}{\delta}$ and $s_w = \left(\frac{30}{1-3f}\right)^2 \ln \frac{3|\text{init_view}_w|}{\delta}$. Let $a_u \cdot N$ and $a_w \cdot N$ be the number of malicious nodes in init_view_u and init_view_w , respectively. Here a_u and a_w are both (unknown) values between 0 and f . Since for all honest nodes u , $|\text{view}_u| \leq (1+f)N \leq (1+f)|\text{init_view}_u|$, the second condition of Line 8 will always be satisfied.

- 1) Consider any given node v where $\text{count}[v] = N$ and any given honest node u . By Theorem 5, each of the N honest nodes in init_view_u will be sampled by u with probability $\frac{s_u}{|\text{init_view}_u|}$, while each of the $a_u N$ malicious nodes will be sampled by u with probability at most $\frac{s_u}{|\text{init_view}_u|}$. We thus have $E[\text{votes}_u[v]] \geq \frac{s_u}{|\text{init_view}_u|} \cdot$

$\text{count}[v] = \frac{s_u}{|\text{init_view}_u|} N \geq \frac{1}{1+f} \cdot s_u$. Hence:

$$\begin{aligned} \Pr[\text{score}_u[v] \leq 0.75] &\leq \Pr[\text{score}_u[v] \cdot s_u \leq (1 - \frac{1-3f}{4}) \cdot (\frac{1}{1+f} \cdot s_u)] \\ &= \Pr[\text{votes}_u[v] \leq (1 - \frac{1-3f}{4}) \cdot (\frac{1}{1+f} \cdot s_u)] \\ &\leq \exp(-(\frac{1-3f}{4})^2 \cdot (\frac{1}{1+f} \cdot s_u) \cdot \frac{1}{2}) \\ &< \exp(-3 \ln \frac{3N}{\delta}) = \frac{\delta^3}{27N^3} \end{aligned}$$

One the steps above invoked a Chernoff bound on $\text{votes}_u[v]$, which is the sum of independent Poisson trials. Finally, there are at most $(1+f)N^2 \leq \frac{4}{3}N^2$ combinations of different v and u . Taking a union bound shows that the first bad event in the theorem happens with probability at most $\frac{4\delta^3}{81}$.

- 2) Consider any given node v where $\text{count}[v] = 0$ and any given honest node u . By Theorem 5, we have $E[\text{votes}_u[v]] \leq \frac{s_u}{|\text{init_view}_u|} \cdot \text{count}[v] + \frac{s_u}{|\text{init_view}_u|} \cdot a_u N = \frac{s_u}{|\text{init_view}_u|} \cdot \frac{a_u}{1+a_u} \cdot |\text{init_view}_u| \leq \frac{f}{1+f} \cdot s_u$. Hence:

$$\begin{aligned} \Pr[\text{score}_u[v] \geq 0.25] &\leq \Pr[\text{score}_u[v] \cdot s_u \geq (1 + \frac{1-3f}{4f}) \cdot (\frac{f}{1+f} \cdot s_u)] \\ &= \Pr[\text{votes}_u[v] \geq (1 + \frac{1-3f}{4f}) \cdot (\frac{f}{1+f} \cdot s_u)] \\ &\leq \exp(-(\frac{1-3f}{4f})^2 \cdot (\frac{f}{1+f} \cdot s_u) \cdot \frac{1}{3}) \\ &< \exp(-3 \ln \frac{3N}{\delta}) = \frac{\delta^3}{27N^3} \end{aligned}$$

One the steps above invoked a Chernoff bound on $\text{votes}_u[v]$. Finally, there are at most $(1+f)N^2 \leq \frac{4}{3}N^2$ combinations of different v and u . Taking a union bound shows that the second bad event in the theorem happens with probability at most $\frac{4\delta^3}{81}$.

- 3) Consider any given node v and any given honest nodes u and w . We consider two separate regions for $\text{count}[v]$. First, for $\text{count}[v] \leq \frac{3N-fN}{8}$, by Theorem 5, we have $E[\text{votes}_u[v]] \leq \frac{s_u}{|\text{init_view}_u|} \cdot \text{count}[v] + \frac{s_u}{|\text{init_view}_u|} \cdot a_u N \leq (\frac{3-f}{8} + a_u) \frac{s_u N}{|\text{init_view}_u|} = \frac{3-f+a_u}{1+a_u} \cdot s_u \leq \frac{3+7f}{8+8f} \cdot s_u$. Hence:

$$\begin{aligned} \Pr[\text{score}_u[v] \geq 0.5] &\leq \Pr[\text{score}_u[v] \cdot s_u \geq (1 + \frac{1-3f}{3+7f}) \cdot (\frac{3+7f}{8+8f} \cdot s_u)] \\ &= \Pr[\text{votes}_u[v] \geq (1 + \frac{1-3f}{3+7f}) \cdot (\frac{3+7f}{8+8f} \cdot s_u)] \\ &\leq \exp(-(\frac{1-3f}{3+7f})^2 \cdot (\frac{3+7f}{8+8f} \cdot s_u) \cdot \frac{1}{3}) \\ &< \exp(-3 \ln \frac{3N}{\delta}) = \frac{\delta^3}{27N^3} \end{aligned}$$

One the steps above invoked a Chernoff bound on $\text{count}_u[v]$. The second case is for $\text{count}[v] > \frac{3N-fN}{8}$.

By Theorem 5, we have $E[\text{votes}_w[v]] \geq \frac{s_w}{|\text{init_view}_w|} \cdot \text{count}[v] > \frac{s_w}{|\text{init_view}_w|} \frac{3N-fN}{8} \geq \frac{3-f}{8+8f} \cdot s_w$. Hence:

$$\begin{aligned} \Pr[\text{score}_w[v] \leq 0.25] &\leq \Pr[\text{score}_w[v] \cdot s_w \leq (1 - \frac{1-3f}{3-f}) \cdot (\frac{3-f}{8+8f} \cdot s_w)] \\ &= \Pr[\text{votes}_w[v] \leq (1 - \frac{1-3f}{3-f}) \cdot (\frac{3-f}{8+8f} \cdot s_w)] \\ &\leq \exp(-(\frac{1-3f}{3-f})^2 \cdot (\frac{3-f}{8+8f} \cdot s_w) \cdot \frac{1}{2}) \\ &< \exp(-3 \ln \frac{3N}{\delta}) = \frac{\delta^3}{27N^3} \end{aligned}$$

One the steps above invoked a Chernoff bound on $\text{votes}_w[v]$. Finally, there are at most $(1+f)N^3 \leq \frac{4}{3}N^3$ combinations of different v , u , and w . Taking a union bound shows that the third bad event in the theorem happens with probability at most $\frac{4\delta^3}{81}$.

- 4) Consider any given node v and any given honest nodes u and w . We consider two separate regions for $\text{count}[v]$. First, for $\text{count}[v] \leq \frac{5N+fN}{8}$, by Theorem 5, we have $E[\text{votes}_w[v]] \leq \frac{s_w}{|\text{init_view}_w|} \cdot \text{count}[v] + \frac{s_w}{|\text{init_view}_w|} \cdot a_w N \leq (\frac{5+f}{8} + a_w) \frac{s_w N}{|\text{init_view}_w|} = \frac{5+f+a_w}{1+a_w} \cdot s_w \leq \frac{5+9f}{8+8f} \cdot s_w$. Hence:

$$\begin{aligned} \Pr[\text{score}_w[v] \geq 0.75] &\leq \Pr[\text{score}_w[v] \cdot s_w \geq (1 + \frac{1-3f}{5+9f}) (\frac{5+9f}{8+8f} \cdot s_w)] \\ &= \Pr[\text{votes}_w[v] \geq (1 + \frac{1-3f}{5+9f}) (\frac{5+9f}{8+8f} \cdot s_w)] \\ &\leq \exp(-(\frac{1-3f}{5+9f})^2 \cdot (\frac{5+9f}{8+8f} \cdot s_w) \cdot \frac{1}{3}) \\ &< \exp(-3 \ln \frac{3N}{\delta}) = \frac{\delta^3}{27N^3} \end{aligned}$$

One the steps above invoked a Chernoff bound on $\text{votes}_w[v]$. The second case is for $\text{count}[v] > \frac{5N+fN}{8}$. By Theorem 5, we have $E[\text{votes}_u[v]] \geq \frac{s_u}{|\text{init_view}_u|} \cdot \text{count}[v] > \frac{5+f}{8} \frac{s_u N}{|\text{init_view}_u|} = \frac{5+f}{8+8a_u} s_u \geq \frac{5+f}{8+8f} s_u$. Hence:

$$\begin{aligned} \Pr[\text{score}_u[v] \leq 0.5] &\leq \Pr[\text{score}_u[v] \cdot s \leq (1 - \frac{1-3f}{5+f}) (\frac{5+f}{8+8f} \cdot s_u)] \\ &= \Pr[\text{votes}_u[v] \cdot s_u \leq (1 - \frac{1-3f}{5+f}) (\frac{5+f}{8+8f} \cdot s_u)] \\ &\leq \exp(-(\frac{1-3f}{5+f})^2 \cdot (\frac{5+f}{8+8f} \cdot s_u) \cdot \frac{1}{2}) \\ &< \exp(-3 \ln \frac{3N}{\delta}) = \frac{\delta^3}{27N^3} \end{aligned}$$

One the steps above invoked a Chernoff bound on $\text{votes}_u[v]$. Finally, there are at most $(1+f)N^3 \leq \frac{4}{3}N^3$ combinations of different v , u , and w . Taking a union bound across all these combinations shows that the fourth bad event in the theorem happens with probability at most $\frac{4\delta^3}{81}$.

Finally, combining the probabilities of the four bad events via a union bound directly leads to the theorem. \blacksquare

VIII. PUT EVERYTHING TOGETHER

We are now ready to put everything together and prove the formal guarantees of RVR (Algorithm 1). All line numbers in this section refers to Algorithm 1. We first provide some additional comments on Algorithm 1. On the leader, the $\text{proposal}_{\text{leader}}$ contains a node v if the score for v is at least 0.5 (at Line 6). A node u will overrule the leader's proposal with respect to v , if u 's score for v is either overwhelming (at Line 7) or underwhelming (at Line 8). Line 11 sets $\text{offset} = 2$ since offset must be 2 after the invocation of CoordinatedGossip. The following theorem proves the complexity of our protocol:

Theorem 7: Assume that $m \geq 1$, $N \geq 1000$, $\delta \leq 0.1$, and $f < \frac{1}{3}$. If all honest nodes start executing RVR within offset rounds of each other, then each honest node will return within $\Theta(\frac{\ln N}{\ln \ln N} \ln \frac{1}{\delta})$ rounds, while sending on expectation $\Theta(N \ln \frac{N}{\delta})$ bits per round.

Proof: The communication complexity of $O(N \ln \frac{N}{\delta})$ bits per round follows directly from Theorem 1, 2, and 4. For the total number of rounds, by Theorem 1, 2, and 4, each iteration in RVR takes at most $8 \cdot \text{offset} + 3 \cdot \text{offset} + \frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + 2 \cdot \text{offset} + 1 = \frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + 1 + 13 \cdot \text{offset}$ rounds. Note that starting from the second iteration in RVR, offset will always be 2. Since there are total $6 \ln \frac{2}{\delta}$ iterations, the total number of rounds needed is at most:

$$\begin{aligned} &(\frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + 1 + 13 \cdot \text{offset}) \\ &+ (\frac{3 \ln((1+f)N)}{2 \ln \ln((1+f)N)} + 1 + 13 \cdot 2) \times (6 \ln \frac{2}{\delta} - 1) \\ &= (\frac{9 \ln((1+f)N)}{\ln \ln((1+f)N)} + 162) \cdot \ln \frac{2}{\delta} + 13 \cdot \text{offset} - 26 \\ &= \Theta(\frac{\ln N}{\ln \ln N} \ln \frac{1}{\delta}) \end{aligned} \quad (4)$$

The next theorem proves that with probability of at least $1 - \delta$, our protocol achieves the intended goal. \blacksquare

Theorem 8: Assume that $m \geq 1$, $N \geq 1000$, $\delta \leq 0.1$, and $f < \frac{1}{3}$. If all honest nodes start executing RVR within offset rounds of each other, then with probability at least $1 - \delta$, both of the following hold:

- 1) For all honest nodes u , final_view_u is the same.
- 2) For all honest nodes u , final_view_u contains all the N honest nodes and at most fN malicious nodes.

Proof: RVR invokes TwoStageSample exactly $6 \ln \frac{2}{\delta}$ times. Let \mathcal{E}_1^r be the event that none of the four bad events in Theorem 6 happens immediately after Line 4 in the r -th iteration of RVR. We claim that $\Pr[\cap_{1 \leq r \leq 6 \ln \frac{2}{\delta}} \mathcal{E}_1^r] \geq 1 - 0.04\delta$. In the first iteration immediately before Line 4 we have $|\text{view}_u| = |\text{init_view}_u| \leq (1+f)N$ for all honest nodes u , which satisfies the condition required to invoke Theorem 6. Conditioned on \mathcal{E}_1^r , immediately before Line 9 in the r -th iteration of RVR, $|\text{proposal}_u| \leq (1+f)N$ for any honest node u since $v \in \text{proposal}_u$ implies that $\text{count}[v] > 0$ and thus is must be included in init_view_w for some honest node w . Recall the definition of $\text{union_honest_init_view}$ from Section III.

Thus $|\text{proposal}_u| \leq |\text{union_honest_init_view}| \leq (1 + f)N$. Conditioned on \mathcal{E}_1^r , it can be similarly shown that immediately before Line 10 in the r -th iteration of RVR, $|\text{view}_u| \leq (1 + f)N$. The claim then follows trivially via induction on r by repeatedly invoking Theorem 6. Let $\mathcal{E}_1 = \bigcap_{1 \leq r \leq 6 \ln \frac{2}{\delta}} \mathcal{E}_1^r$.

Next, we say that an iteration in RVR is *good* if during that iteration, i) ProbLeaderElect returns the same leader on all honest nodes and the leader returned is an honest node, and ii) CoordinatedGossip returns $\text{proposal}_{\text{leader}}$ on all honest nodes. By Theorem 1 and 3 and a union bound, an iteration is good with probability at least $0.16 - \frac{1}{4000} = 0.15975$. Hence with probability at least $1 - (1 - 0.15975)^{6 \ln \frac{2}{\delta}} > 1 - \delta^{1.03} > 1 - 0.94\delta$, there exists some good iteration. Let such event be \mathcal{E}_2 . A union bound immediately tells us that with probability at least $1 - \delta$, both \mathcal{E}_1 and \mathcal{E}_2 occur. It thus suffices to prove the two claims in the theorem while condition upon that \mathcal{E}_1 and \mathcal{E}_2 occur. The following proves them one by one.

First claim. We first prove that at the end of the good iteration, view_u is the same on all honest nodes u . The good iteration must have a unique and honest leader. Consider any node v . We will show that for all honest nodes u , at the end of that iteration, $v \in \text{view}_u$ iff $v \in \text{proposal}_{\text{leader}}$. To prove this, note that if $v \in \text{proposal}_{\text{leader}}$, then $\text{score}_{\text{leader}}[v] \geq 0.5$. By event \mathcal{E}_1 , $\text{score}_u[v] \not\leq 0.25$ and hence $v \notin \text{underwhelm}_u$. By Line 10, v will be in view_u . The case for $v \notin \text{proposal}_{\text{leader}}$ is similar. Now because $\text{proposal}_{\text{leader}}$ is the same on all honest nodes u , view_u is the same on all honest nodes u .

Next we prove that in any iteration after the good iteration, on all honest nodes u , view_u will no longer change. Immediately after the good iteration, for all nodes v , we have either $\text{count}[v] = N$ or $\text{count}[v] = 0$. If $\text{count}[v] = N$, then event \mathcal{E}_1 ensures that $v \in \text{over}_u$ in all future iterations. If $\text{count}[v] = 0$, then \mathcal{E}_1 ensures that $v \in \text{under}_u$ in all future iterations. Together with the condition at Line 10, we know that view_u will no longer change in future iterations.

Finally, since at the end of the good iteration, view_u is the same on all honest nodes u , and since view_u does not change in all iterations after the good iteration, we know that final_view_u is the same for all honest nodes u .

Second claim. Consider any given honest node v . We know that $\text{count}[v] = N$ at the beginning of the first iteration. By the same argument as earlier, v will continue to be in view_u (and thus $v \in \text{final_view}_u$) for all honest node u in all iterations. Next, $\text{union_honest_views}$ contains at most fN malicious nodes. For all malicious node $v \notin \text{union_honest_views}$, we know that $\text{count}[v] = 0$ at the beginning of the first iteration. By the same argument as earlier, we know that $v \notin \text{view}_u$ (and thus $v \notin \text{final_view}_u$) for all honest nodes u in all iterations. ■

IX. NUMERICAL COMPARISONS

The previous section has formally proved the correctness and performance of our RVR protocol. In particular, we have shown that the time complexity of RVR is $\Theta(\frac{\ln N}{\ln \ln N} \ln \frac{1}{\delta})$ rounds. This section aims to supplement such asymptotic results with some concrete numerical examples. In particular, our numerical examples will capture the various hidden

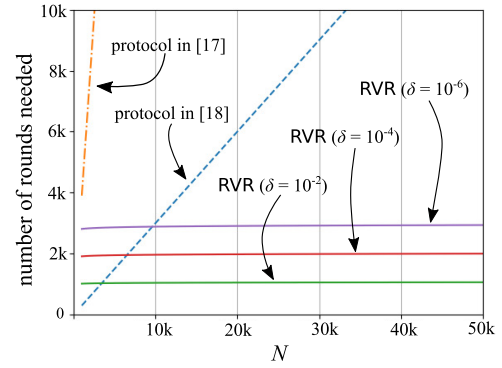


Fig. 2. Number of rounds needed by our RVR protocol and the existing view reconciliation protocols [18], [19].

constants in the asymptotic terms. Since this paper focuses on formal and provable guarantees, our goal here is not to do large-scale real-world deployment/evaluation of RVR, which is well beyond the scope of this paper. (Similarly, the previous efforts [18], [19] also focused on formal/provable guarantees, instead of deployment/evaluation.)

For results on RVR in this section, we compute the number of rounds needed by RVR by following the exact form of its time complexity analysis in Equation 4, with $\text{offset} = 1$. For the view reconciliation protocol from [18] (based on Lemma 2 in [18] and Fig. 4 in [18]), one can verify that their protocol takes at least $(2l + 3) + (l + 1)$ rounds. Here the value of l , when translated to our setting, is $(f + 1)N$. Hence we use $3(f + 1)N + 4$ as the total number of rounds needed by that protocol. Finally, one can verify (based on Figure 1 in [19] and Theorem 1 in [19]) that the protocol in [19] takes $fN + 1$ rounds.⁶

Figure 2 compares the number of rounds needed by the three protocols, under $f = 0.3$ and N from 1k to 50k. Results under other f values (e.g., $f = 0.1$) are qualitatively similar. For RVR in the figure, we consider three different δ values: 10^{-2} , 10^{-4} , and 10^{-6} . As expected, the number of rounds needed by the two existing protocols increase linearly with N , while the number of rounds for RVR increases only logarithmically. (In the figure, such logarithmic curves appear almost flat.) When N is large (e.g., above 4k for $\delta = 10^{-2}$, above 7k for $\delta = 10^{-4}$, and above 10k for $\delta = 10^{-6}$), RVR always outperforms both existing protocols. The benefit of RVR increases with N : For $N = 50k$, RVR (even under $\delta = 10^{-6}$) is over 400% faster than the existing protocols (not shown in the figure). On the other hand, due to the larger constant factor in RVR, when N is smaller, the existing protocol in [19] requires fewer number of rounds than RVR.

X. CONCLUSIONS

This paper proposes a novel randomized RVR protocol for solving the view divergence problem. Compared with the state-of-the-art protocols [18], [19] with time complexity of $\Theta(N)$ rounds, our protocol only takes $\Theta(\frac{\ln N}{\ln \ln N} \ln \frac{1}{\delta})$ rounds to terminate. The hidden constant factor in the asymptotic complexity of our protocol, however, is considerably larger

⁶In [19], the number of rounds is shown to be “ $f + 1$ ”. But their definition of “ f ” actually equals fN in this paper’s setting.

than in previous protocols. Concrete numerical comparisons show that our protocol is more suitable for large-scale systems, while existing protocols are better for smaller-scale systems.

ACKNOWLEDGMENT

The authors thank the IEEE/ACM TRANSACTIONS ON NETWORKING anonymous reviewers and the IEEE INFOCOM anonymous reviewers for their helpful feedback. The authors of this article are alphabetically ordered.

REFERENCES

- [1] R. Hou, I. Jahja, L. Luu, P. Saxena, and H. Yu, "Randomized view reconciliation in permissionless distributed systems," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Apr. 2018, pp. 2528–2536.
- [2] J. Douceur, "The sybil attack," in *Proc. IPTPS*, 2002, pp. 251–260.
- [3] C. Dwork and M. Naor, "Pricing via processing or combatting junk mail," in *Proc. CRYPTO*, 1992, pp. 139–147.
- [4] C. Dwork, M. Naor, and H. Wee, "Pebbling and proofs of work," in *Proc. CRYPTO*, 2005, pp. 37–54.
- [5] S. Nakamoto. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. [Online]. Available: <http://bitcoin.org/bitcoin.pdf>
- [6] J. Aspnes, C. Jackson, and A. Krishnamurthy, "Exposing computationally-challenged Byzantine impostors," Yale Univ. Dept. Comput. Sci., New Haven, CT, USA, Tech. Rep. YALEU/DCS/TR-1332, Jul. 2005.
- [7] P. Berman and J. Garay, "Cloture votes: $n/4$ -resilient distributed consensus in $t + 1$ rounds," *Math. Syst. Theory*, vol. 26, no. 1, pp. 3–19, Jan. 1993.
- [8] M. Castro and B. Liskov, "Practical byzantine fault tolerance and proactive recovery," *ACM Trans. Comput. Syst.*, vol. 20, no. 4, pp. 398–461, Nov. 2002.
- [9] D. Dolev and H. R. Strong, "Authenticated algorithms for byzantine agreement," *SIAM J. Comput.*, vol. 12, no. 4, pp. 656–666, Nov. 1983.
- [10] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, p. 382–401, 1982.
- [11] M. Pease, R. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *J. ACM*, vol. 27, no. 2, pp. 228–234, Apr. 1980.
- [12] M. Raynal, "Consensus in synchronous systems: A concise guided tour," in *Proc. Pacific Rim Int. Symp. Dependable Comput.*, Dec. 2002, pp. 221–228.
- [13] M. Ben-Or and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation," in *Proc. 20th Annu. ACM Symp. Theory Comput. (STOC)*, 1988, pp. 1–10.
- [14] O. Goldreich, S. Micali, and A. Wigderson, "How to play ANY mental game," in *Proc. 19th Annu. ACM Conf. Theory Comput. (STOC)*, 1987, pp. 218–229.
- [15] A. C. Yao, "Protocols for secure computations," in *Proc. 23rd Annu. Symp. Found. Comput. Sci. (FOCS)*, Nov. 1982, pp. 160–164.
- [16] B. Awerbuch and C. Scheideler, "Towards a scalable and robust DHT," *Theory Comput. Syst.*, vol. 45, no. 2, pp. 234–260, Aug. 2009.
- [17] R. Guerraoui, F. Huc, and A.-M. Kermerrec, "Highly dynamic distributed computing with Byzantine failures," in *Proc. ACM Symp. Princ. Distrib. Comput. (PODC)*, 2013, pp. 176–183.
- [18] M. Andrychowicz and S. Dziembowski, "Pow-based distributed cryptography with no trusted setup," in *Proc. Int. Cryptol. Conf. (CRYPTO)*, Aug. 2015, pp. 379–399.
- [19] J. Katz, A. Miller, and E. Shi, "Pseudonymous secure computation from time-lock puzzles," Cryptology ePrint Arch., Tech. Rep. 2014/857, 2014. [Online]. Available: <https://eprint.iacr.org/contact.html>
- [20] *Global Bitcoin Nodes Distribution–Bitnodes*. Accessed: Jul. 6, 2020. [Online]. Available: <https://bitnodes.earn.com>
- [21] *Coin Dance | Bitcoin Nodes Summary*. Accessed: Jul. 6, 2020. [Online]. Available: <https://coin.dance/nodes>
- [22] S. Park, S. Im, Y. Seol, and J. Paek, "Nodes in the bitcoin network: Comparative measurement study and survey," *IEEE Access*, vol. 7, pp. 57009–57022, 2019.
- [23] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling byzantine agreements for cryptocurrencies," in *Proc. 26th Symp. Operating Syst. Princ.*, 2017, pp. 51–68.
- [24] *The First Pure Proof of Stake Blockchain Platform | Algorand*. Accessed: Jul. 6, 2020. [Online]. Available: <https://www.algorand.com/>
- [25] K. O'Brien. (2018). *MIT Professor Raises \$4M to Solve Common Blockchain Problems*. Boston Business Journal. [Online]. Available: <https://www.bizjournals.com/boston/news/2018/02/15/mit-professor-raises-4m-to-solve-common-blockchain.html>
- [26] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena, "A secure sharding protocol for open blockchains," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Oct. 2016, pp. 17–30.
- [27] N. Borisov, "Computational puzzles as sybil defenses," in *Proc. 6th IEEE Int. Conf. Peer-to-Peer Comput. (P2P06)*, Sep. 2006, pp. 171–176.
- [28] F. Li, P. Mittal, M. Caesar, and N. Borisov, "SybilControl: Practical Sybil defense with computational puzzles," in *Proc. 7th ACM Workshop Scalable Trusted Comput.*, 2012, pp. 67–78.
- [29] H. Rowaihy, W. Enck, P. McDaniel, and T. L. Porta, "Limiting sybil attacks in structured P2P networks," in *Proc. INFOCOM*, May 2007, pp. 1–5.
- [30] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on Bitcoin's peer-to-peer network," in *Proc. USENIX Secur. Symp.*, 2015, pp. 129–144.
- [31] R. Pass and E. Shi, "Hybrid consensus: Efficient consensus in the permissionless model," in *Proc. DISC*, Oct. 2017, pp. 39:1–39:16.
- [32] A. Aggarwal, M. Movahedi, J. Saia, and M. Zamani, "Brief announcement: Bootstrapping public blockchains without a trusted setup," in *Proc. PODC*, Jul. 2019, pp. 366–368.
- [33] M. Bellare and P. Rogaway, "Random oracles are practical: A paradigm for designing efficient protocols," in *Proc. 1st ACM Conf. Comput. Commun. Secur.*, Nov. 1993, pp. 62–73.
- [34] J. Burns and N. Lynch, "The Byzantine firing squad problem," *Adv. Comput. Res., Parallel Distrib. Comput.*, vol. 4, pp. 147–161, Apr. 1987.
- [35] B. Chlebus and D. Kowalski, "Time and communication efficient consensus for crash failures," in *Proc. DISC*, Sep. 2006, pp. 314–328.
- [36] B. S. Chlebus and D. R. Kowalski, "Robust gossiping with an application to consensus," *J. Comput. Syst. Sci.*, vol. 72, no. 8, pp. 1262–1281, Dec. 2006.
- [37] C. Georgiou, S. Gilbert, R. Guerraoui, and D. Kowalski, "Asynchronous gossip," *J. ACM*, vol. 60, no. 2, p. 11:1–11:42, May 2013.
- [38] S. Gilbert, R. Guerraoui, and D. Kowalski, "On the message complexity of indulgent consensus," in *Proc. DISC*, Sep. 2007, pp. 283–297.