Tiny Groups Tackle Byzantine Adversaries

Mercy O. Jaiyeola[*], Kyle Patron[†], Jared Saia[‡], Maxwell Young[*], Qian M. Zhou[*]

[*]Mississippi State University,
Dept. of Computer Science and Eng.,
Miss. State, MS, USA
{moj31,my325,qz70}@msstate.edu

[†]Palantir Technologies
New York, NY, USA
kyle.patron@gmail.com

[‡]University of New Mexico,
Department of Computer Science,
Albuquerque, NM, USA
saia@cs.unm.edu

Abstract—A popular technique for tolerating malicious faults in open distributed systems is to establish small *groups* of participants, each of which has a non-faulty majority. These groups are used as building blocks to design attack-resistant algorithms.

Despite over a decade of active research, current constructions require group sizes of $O(\log n)$, where n is the number of participants in the system. This group size is important since communication and state costs scale polynomially with this parameter. Given the stubbornness of this logarithmic barrier, a natural question is whether better bounds are possible.

Here, we consider an attacker that controls a constant fraction of the total computational resources in the system. By leveraging proof-of-work (PoW), we demonstrate how to reduce the group size *exponentially* to $O(\log \log n)$ while maintaining strong security guarantees. This reduction in group size yields a significant improvement in communication and state costs.

I. Introduction

Byzantine fault tolerance addresses the challenge of performing useful work when participants in the system are malicious. Participants, or ***identifiers (IDs)*** in the system, may be malicious; these malicious IDs can discard or corrupt information that is routed through them or stored on them.

A popular technique for overcoming these challenges is to arrange IDs into sets called ***groups***,[1] where each has a non-faulty majority. We can then ensure the following:

- Secure routing is possible. For groups G_1 and G_2 along a route, all members of G_1 transmit messages to all members of G_2. This all-to-all exchange, followed by majority filtering by each non-faulty ID in G_2, guarantees correctness of communication between groups despite malicious IDs.

- Computation is performed by all members of a group via protocols for ***Byzantine agreement (BA)*** [28], or more general secure multiparty computation [49], to guarantee that tasks execute correctly. In this way, each group simulates a ***reliable processor*** upon which jobs can be run.

The use of groups provides a *scalable* approach to designing an attack-resistant distributed system, by avoiding the need to have all n IDs perform BA in concert.[2]

Designing attack-resistant systems using groups has been an active topic of research for over a decade, with many results [8]–[10], [12], [18], [21], [23], [26], [27], [37], [39], [45], [51]. Yet, despite this progress, an enduring requirement is that each group contains $O(\log n)$ members; that is, the group size grows logarithmically in n.

Why does this logarithmic size matter? At first glance, it is an unlikely bottleneck. However, since groups are building blocks for the system, their size, $|G|$, impacts costs:

(i) *Cost of Group Communication.* Group members must often act in concert; for example, executing distributed key generation [51], or generating random numbers [8], [18]. Such protocols require messages be exchanged between all members. We label this as ***group communication*** and it has $\Theta(|G|^2)$ message cost.

(ii) *Cost of Secure Routing.* Routing via all-to-all exchange between two groups incurs $\Omega(|G|^2)$ message complexity. Given a route of length D, communication between any two groups requires $O(D|G|^2)$ messages.[3]

(iii) *Cost of State Maintenance.* Each ID w must maintain state on its neighbors; this includes both the members of all groups to which w belongs <u>and</u> the members of neighboring groups. This requires storing link information, as well as periodically testing links for liveness.[4]

In each case above, reducing $|G|$ would directly reduce cost. Unfortunately, in prior results, $|G| \approx \log n$ is key to ensuring that all groups have a non-faulty majority with high probability (w.h.p.).[5] Without this property, all previous group constructions succumb to adversarial attack. A natural question is: *Are there new ideas that allow us to decrease $|G|$ while maintaining strong security guarantees?*

In response, we consider an attacker that controls a constant fraction of the total computational resources in the system. By employing proof-of-work (PoW), we obtain our main result: **Group size can be reduced exponentially while still allowing routing in all but a vanishingly small portion (an $o(1)$-fraction) of the network**.

---

[1]Such sets have appeared under different names in the literature, such as "swarms" [18], "clusters" [21], and "quorums" [51]. Our choice of "groups" aligns with pioneering work in this area [7].

[2]Groups also improve robustness in other ways. Members may agree to ignore an ID if it misbehaves too often, hence reducing spamming. Data may also be redundantly stored at multiple group members.

[3]Improvements are possible, but they come with caveats. Results in [18], [45] lower the cost to $O(D|G|)$ in expectation but require a non-trivial (expander-like) construction, and [51] further reduces this to $O(D)$ in expectation but with a poly($|G|$) message cost each time routing tables are updated which is expensive even with moderate churn.

[4]Constructions in [18], [39] have each ID belonging to $\eta > 1$ groups for a state overhead of $\Omega(|G|\eta)$. Also, if each group links to Δ neighboring groups, then $O(|G|\Delta)$ links must be maintained; typically, $\Delta = O(\log n)$.

[5]With probability at least $1 - 1/n^{c'}$ for a tunable constant $c' \geq 1$.

1

A. Defining ε-Robustness

Consider a system of n IDs and n groups, where a β-fraction of IDs are malicious; such IDs may deviate arbitrarily from any prescribed protocol to derail operations in the network. The following defines our notion of *ε-robustness*:

For a small $\varepsilon > 0$, at least $(1-\varepsilon)n$ groups have a non-faulty majority and can securely route messages to each other.

The parameter β is a sufficiently small positive constant less than $1/2$, and $\varepsilon = o(1)$. We consider the following questions:

Is this a useful concept? Consider decentralized storage and retrieval of data. This definition guarantees all but an ε-fraction of data is reachable and maintained reliably. Example applications include distributed databases, name services, and content-sharing networks. Alternatively, consider n jobs in an open computing platform that are run on individual machines. This definition guarantees that all but an ε-fraction of those jobs can be correctly computed.[6]

Is satisfying this definition trivial? Given $\Theta(n)$ non-faulty IDs, this definition characterizes simulating $(1-\varepsilon)n$ reliable processors <u>and</u> the ability to route information between them. If we ignore the use of groups or, equivalently, consider groups each consisting of a single ID, then we trivially have $(1-\beta)n$ reliable processors, but routing between them is challenging. For example, establishing links between all pairs of IDs will give secure routing, but this is hardly scalable.

Do previous solutions solve this problem? Prior results using groups are subsumed by this definition and they address $\varepsilon = 1/\text{poly}(n)$. In this case, routing is possible – albeit, costly – because w.h.p. *all* groups have a non-faulty majority.

To reduce cost, we consider $\varepsilon = 1/\mathbf{poly(\log n)}$. We show how this allows us to reduce the group size exponentially, but at the price of having a small fraction of the groups with a majority of malicious IDs.

B. Related Work

Tolerating Byzantine Faults via Groups. The use of groups for building attack-resistant distributed systems has received significant attention, and all address the case of $\varepsilon = 1/\text{poly}(n)$.

Early results obtain a $\text{poly}(\log n)$ factor increase in costs, assuming constraints on the amount of system dynamism [7], [17], [18], [23], [38].

Full dynamism was achieved by Awerbuch and Scheideler in a series of breakthrough results [8]–[10]; they propose a *cuckoo rule* that w.h.p. preserves a non-faulty majority in all groups over $n^{\Theta(1)}$ joins/departures when the system size remains $\Theta(n)$. More recently, Guerraoui et al. [21] give similar guarantees when the system size can vary polynomially.

Simulations of the cuckoo rule are conducted in [47]. The trade-off between group size and security is examined, and findings suggest that $|G|$ must be fairly large. For $n = 8,192$ (the largest size examined) and $\beta \approx 0.002$, $|G| = 64$ preserves a non-faulty majority in each group for 10^5 joins/departures; $\beta \approx 0.07$ is possible with suggested improvements in [47].

Several results have focused on reducing communication costs when the good majority of all groups is guaranteed via an algorithm like the cuckoo rule [45], [51]. Here too, group size impacts performance and $|G| = 30$ incurs significant latency in PlanetLab experiments [51]. Groups have also been used in conjunction with quarantining malicious IDs [27], [43] with limited churn.

Finally, we observe that none of these results explicitly uses PoW, with the possible exception of [7], where computational challenges or Turing tests are briefly discussed as a means for throttling the join rate of Byzantine IDs. These prior results assume a model where the fraction of Byzantine IDs is always limited to strictly less than $1/2$. In contrast, the use of PoW provides a plausible mechanism by which to enforce this limit (see below for more discussion).

Attack-Resistance Without Groups. Other distributed constructions exist that do not explicitly use groups [13], [17], [44]. However, the associated techniques retain some form of $O(\log n)$ redundancy with regards to data placement or route selection and, therefore, incur the typical $\text{poly}(\log n)$ cost.

In [12], [26], [37], malicious faults are tolerated by routing along multiple diverse routes. However, it is unclear that these systems can provide theoretical guarantees on robustness.

Byzantine resistance when $O(\sqrt{n}/\text{poly}(\log n))$ IDs may depart and join *per time step* is examined in [2], [3].[7] In this challenging model, (roughly) $O(\sqrt{n})$ Byzantine IDs can be tolerated. Our result addresses more moderate churn while tolerating $\Theta(n)$ Byzantine IDs.

Central authorities (CAs) have been used in prior results [12], [42] to achieve robustness. While our results can be used in conjunction with a CA, it is not always plausible to assume such an authority is available and immune to attack. For this reason, our work does not depend on a CA.

Computational Puzzles. Proof-of-work (PoW) via computational puzzles has been used to mitigate the Sybil attack [14], whereby an adversary overwhelms a system with a large number of malicious IDs. We note that such PoW schemes have been proposed in decentralized settings; for examples, see [22], [30]. However, such PoW schemes only *limit the number of Sybil IDs* — typically commensurate with the amount of computational power available to the adversary — and the problem of tolerating these adversarial IDs must still be addressed by other means; for examples, see [46], [52].

A prominent example of using PoW to provide security is Bitcoin [36]. However, note that analyses of Bitcoin and related systems typically assume a communication primitive that allows an ID to disseminate a value to all other IDs within a known bounded constant amount of time despite an adversary [20], [31], [33]. In contrast, our results do not assume the existence of such a primitive.

We note that PoW imposes a computational overhead on the system participants. Nonetheless, examples such as Bitcoin and emerging blockchain technologies (for example, Ethereum [16]) illustrate the success of PoW in practice, and exemplify that computational overhead from PoW may be tolerable given the security guarantees received in exchange.

---

[6] While this may not be sufficient for general computation, it is valuable for tasks where an $o(1)$ error rate or bias can be tolerated; for example, obtaining statistics on a group of machines to measure network performance.

[7] High churn without Byzantine fault tolerance is considered in [1], [4], [5].

C. Our Model and Preliminaries

We consider a system of n IDs. An ID is **good** if it obeys the protocol; otherwise, the ID is Byzantine or **bad**.

The Adversary. We assume that our adversary controls a β-fraction of the computational power in the network, where β is a sufficiently small positive constant less than $1/2$.[8] This is a common assumption when using PoW to design attack-resistant, open systems [22], [30], [41]. For simplicity, throughout Sections II and III, we assume there is always at most a β fraction of bad IDs. In Section IV, this assumption is justified by proving that the adversary is constrained, via proof-of-work assumptions, to generating (roughly) at most a β-fraction of IDs at any time.

We assume that a single adversary controls all the bad IDs. This is a challenging model since a single adversary allows the bad IDs to perfectly collude and coordinate their attacks. The adversary also knows the network topology and all message contents; however, the adversary does not know the random bits generated locally by any good ID.

Groups. In our system, each group has size $\Theta(\log \log n)$. Each ID w has its own group G_w and w is referred to as the **leader**. A group G is **good** if (i) $d_1 \ln \ln n \leq |G| \leq d_2 \ln \ln n$ for appropriate constants $d_1 < d_2$, and (ii) the number of bad IDs in G is at most $(1 + \delta)\beta|G|$ for some tunably small constant $\delta > 0$ depending only on sufficiently large n; otherwise, the group is **bad**. Note that groups are not necessarily disjoint; in addition to being the leader of G_w, ID w may belong to other groups. Group construction is described in Section III-A.

Input Graph. Our result builds off an **input graph H** on N vertices, where each vertex corresponds to an ID.[9] Each ID is a virtual participant, and each ID is represented as a value in the interval $[0, 1)$ known as the **ID space**; this is viewed as a unit ring where moving **clockwise** along the ring corresponds to moving from away from 0 towards 1. The **successor** of a point x in $[0, 1)$ is the first ID encountered by moving clockwise from x on the unit ring; this is denoted by **$suc(x)$**. Assuming there is no adversary and that IDs are distributed independently and uniformly at random (u.a.r.) in $[0, 1)$, the following properties hold for H with probability at least $1 - N^{-c'}$ for a tunable constant $c' \geq 1$.

P1 – Search Functionality. There exists a **search** (or **routing**) **algorithm** that, for any **key value** in $[0, 1)$, returns contact information for the ID responsible for the corresponding **resource** (i.e., data item, computational job, network printer, etc.). A search requires traversing $D = O(\log N)$ IDs.[10]

P2 – Load Balancing. A randomly chosen ID is responsible for at most a $(1 + \delta'')/N$-fraction of the key values (and the corresponding resources) for an arbitrarily small $\delta'' > 0$ depending on sufficiently large N.

P3 – Linking Rules. Each ID w links to IDs in a set of neighbors S_w; $|S_w| = O(\log^\gamma n)$ for some constant $\gamma > 0$. Any ID may determine the elements in S_w by performing searches.[11]

There are also $O(\text{poly}(\log n))$ IDs each of which has w in its respective set of neighbors. Again, any ID may verify this by performing searches. The number of links on which ID w is incident is the **degree** of w, and every ID has the same degree asymptotically.

P4 – Congestion Bound. The **congestion** is $C = O(\log^c n/n)$ for a constant $c \geq 0$, where congestion is the maximum probability (over all IDs) that a ID is traversed in a search initiated at a randomly chosen ID for a randomly chosen point in $[0, 1)$.

We emphasize that H is not assumed to provide any security or performance guarantees if there are bad IDs present. Rather, any such H provides a viable topology that, using our result, can be made to tolerate bad IDs. Note that many constructions for H exist such as Chord [48], the distance-halving construction [39], Viceroy [32], Chord++ [6], D2B [19], FISSIONE [29], and Tapestry [53].

IDs and PoW in Our Construction. An ID is a virtual participant in the network, and each ID is represented as a value in $[0, 1)$ in our construction. Note that adequate precision is obtained using $O(\log n)$ bits. Important properties that our system guarantees are:

- IDs expire after a period of time that can be set by the system designers.
- A claim to own an ID can be verified by any good ID.
- The adversary possesses (roughly) at most βn IDs, and these IDs are u.a.r. from the ID space $[0, 1)$.

We emphasize that our construction does not take these properties for granted; rather they are enforced via a PoW scheme. However, given space constraints and that the bulk of our results are proved without the need to reference these details, we will assume these properties in Sections II and III. We remove these assumptions in Section IV.

We make the *random oracle assumption* [11]: there exist hash functions, h, such that $h(x)$ is uniformly distributed over h's range, when any x in the domain of h is input to h for the first time. We assume that both the input and output domains are the real numbers $[0, 1)$. In practice, h may be a cryptographic hash function, such as SHA-2 [40], with inputs and outputs of sufficiently large bit lengths.

Joins and Departures. Our work addresses a dynamic system where IDs may join and depart. We delay our description of this aspect until Section III.

We use the following concentration results.

Theorem 1. *(Chernoff Bounds [35]) Let $X_1, \ldots, X_N$ be independent indicator random variables such that $\Pr(X_i) = p$*

---

[8]For simplicity in our proofs, we let β be small since we do not try to optimize this quantity. However, larger values of β are likely possible.

[9]For clarity, we use N to denote the number of IDs in H, and we use n for the number of IDs in our attack-resistant construction.

[10]An ID is "traversed" if the ID lies on the path of search between the initiator of the search and the ID responsible for the resource being sought. For more discussion, please see our online version [24], Appendix VI.

[11]For example, in Chord [48], w's neighbors are (1) the ID counter-clockwise and the ID clockwise to w on the unit ring, and (2) the respective successors of points $w + \Delta(i)$, where $\Delta(i)$ is an exponentially increasing distance in the ID space for $i = 1, ..., O(\lg N)$. Any ID may verify that u is a neighbor of w via a search on $w + \Delta(i)$ and checking that the result is u.

and let $X = \sum_{i=1}^{N} X_i$. For any δ, where $0 < \delta < 1$, the following holds:

$$\Pr(X > (1 + \delta) E[X]) \leq e^{-\delta^2 E[X]/3}$$

$$\Pr(X < (1 - \delta) E[X]) \leq e^{-\delta^2 E[X]/2}$$

Theorem 2. *(Method of Bounded Differences, Corollary 5.2 in [15]) Let f be a function of the variables $X_1, ..., X_N$ such that for any b, b' it holds that $|f(X_1, ..., X_i = b, ..., X_N) - f(X_1, ..., X_i = b', ..., X_N)| \leq c_i$ for $i = 1, ..., N$. Then, the following holds:*

$$Pr(f > E[f] + t) \leq e^{-2t^2/(\sum_i c_i^2)}$$

$$Pr(f < E[f] - t) \leq e^{-2t^2/(\sum_i c_i^2)}$$

Finally, all of our results hold given that n is sufficiently large; we assume this throughout.

D. Overview and Our Main Result

As discussed above, reducing group size is desirable, but gives rise to the possibility of bad groups. In Section II, we demonstrate how to achieve $1/\text{poly}(\log n)$-robustness with groups of size $\Theta(\log \log n)$ when there is no churn. This argument leverages the bound on congestion given by the input graph, along with a careful tallying of the fraction of ID space that cannot be securely searched.

This result is applied in Section III where we show that $O(1/\text{poly}(\log n))$-robustness can be maintained with churn. A key component of our construction is the use of two graphs (composed of groups) that, when used in tandem, limit the number of bad groups that can be formed.

Finally, in Section IV, we describe how PoW is used to provide the guarantees on IDs discussed in Subsection I-C. The main challenge is defending against an adversary that wishes to store a large number of IDs for use in a massive future attack (i.e., a pre-computation attack).

Our main result is the following:

Theorem 3. *Assume an input graph H that satisfies P1 - P4, and that the adversary has at most a β-fraction of the computational power, for some sufficiently small but positive constant β. Then our construction using $|G| = O(\log \log n)$ provides the following guarantees w.h.p. over a polynomial number of join and departure events.*

- *All but an $O(1/\text{poly}(\log n))$-fraction of groups are good.*
- *All but an $O(1/\text{poly}(\log n))$-fraction of IDs can successfully search for all but an $O(1/\text{poly}(\log n))$-fraction of the resources.*

That is, we achieve $O(1/\text{poly}(\log n))$-robustness.

To illustrate our cost improvements, we also establish:

Corollary 1. *Using any of the constructions in [19], [32], or [39] as an input graph H, our result gives the following bounds on cost.*

- *Group communication incurs $O(\text{poly}(\log \log n))$ messages.*
- *Secure routing incurs $O(D \, \text{poly}(\log \log n))$ messages.*
- *State maintenance is expected $O(\text{poly}(\log \log n))$ per ID.*

Note that these are substantial improvements over the costs described in Section I.

Can we do better? We offer some intuition for why significantly improving on our result seems unlikely. With $|G| = \Theta(\log \log n)$, the probability of a bad group is roughly $1/\text{poly}(\log n)$. All constructions with $o(\log n)$ degree requires $D = \omega(\log n / \log \log n)$ IDs to be traversed in a search. Thus, the probability of encountering a bad group along the path of search is (roughly) at most $\sum_1^D 1/\text{poly}(\log n)$ by a union bound, and this can be *less than* 1.

Now, consider a smaller group size of $o(\log \log n / \log \log \log n)$. Then, the probability of a bad group is $\omega(\log \log n / \log n)$, and over D IDs traversed, a union bound no longer bounds the probability of a failed search by less than 1.

In this sense, our choice of $|G|$ appears to be pushing the limits of what is possible when using groups to design attack-resistant systems.

II. THE STATIC CASE

We first prove a result for searches without dealing with ID joins and departures.

A. The Group Graph

Given an input graph H, our ϵ-robust construction is a **group graph** $\mathcal{G}$ where each ID w in H corresponds to group G_w.

We refer to each group in $\mathcal{G}$ as **blue** or **red**. A blue group corresponds to a good group with its neighbors correctly established, while a red group corresponds to a bad group <u>or</u> a group that has an incorrect neighbor set. When addressing the group graph, we use the notation G_w to denote a vertex in $\mathcal{G}$; however, G_w may also refer to the group with leader w, and this will be made clear from the context.

In $\mathcal{G}$, each blue group G_w has a **neighbor set**, L_w, where for each u that is a neighbor of w in H, the group G_u exists in L_w; each red group has an arbitrary neighbor set determined by the adversary. A search in $\mathcal{G}$ proceeds over edges in $\mathcal{G}$ as it would in H with the corresponding group members participating in the search; this is illustrated in Figure 1. For an edge (G_w, G_v) between two blue groups in $\mathcal{G}$, there are all-to-all links between (at least) the good members in G_w and G_v.

The following properties of $\mathcal{G}$ are useful for our analysis:

- *S1.* For each ID v in H, there is a group G_v in $\mathcal{G}$ where the leader v of $\mathcal{G}$ has the same ID in both graphs.

- *S2.* Each group in $\mathcal{G}$ is red independently with probability $p_f \leq 1/\log^k n$ for a tunably large constant $k > 0$ depending only on d_1; and blue otherwise.

- *S3.* Edges incident to blue groups are set according to the topology of H. All other edges are set by the adversary.

The value of p_f in S2 corresponds to the probability that a group is bad or does not have the correct neighbor set. To provide intuition for our bound on p_f, note that if we select $\Theta(\log \log n)$ IDs u.a.r., then the probability that a majority are bad is $O(1/\text{poly}(\log n))$ by a Chernoff bound. A similar bound can be derived on the probability of incorrectly setting up neighbors; there is a subtlety with respect to bounding this

Fig. 1: Left: An input graph H with IDs w, u, v, and y. Illustrated is a search initiated at w and terminating at y; this search also traverses u and v. Right: A group graph with corresponding groups G_w, G_u, G_v, and G_y. Red groups are marked with a "B". Large dashed arrows represent all-to-all links between (at least) good members of the corresponding groups.

probability, and this is discussed in our online version [24], Appendix X. Keeping p_f upper bounded by $1/\log^k n$ when IDs can join and depart is non-trivial. We show how to do this in Section III-B.

Since the adversary controls all red groups, it is free to insert or delete edges between red groups (cf. $S3$). However, edges involving at least one blue group are not modified. This is because the adversary cannot modify a good group's knowledge of who its neighbors are, since that knowledge is kept consistent by the good majority, although the neighboring bad group may certainly ignore or corrupt incoming messages from that good group.

Overview of Analysis. A search in group graph $\mathcal{G} = (V, E)$ is said to **fail** if it traverses any red group. Otherwise, the search will **succeed**.

In $\mathcal{G}$, consider the path of a search that begins at the initial group and halts either upon succeeding or encountering the first red group (in which case the search fails); we call this a **search path**.

For any group G_v, we define the **responsibility** of G_v to be the probability that a search path in $\mathcal{G}$ from a random group to a random point in $[0, 1)$ traverses G_v. We denote the responsibility of group G_v by $\rho(G_v)$.

Why is responsibility defined in terms of search paths? The issue is that responsibility is not well-defined after a search encounters the first red group since the adversary may redirect a search to any red group after this point. For example, the adversary may have the same red group traversed by multiple different searches, thus arbitrarily inflating the number of searches that traverse this red group. This motivates the notion of a search path.

B. Analysis

Lemma 1. *The following holds, w.h.p, for any group G_v:* $\rho(G_v) = O(\log^c n/n)$.

Proof. By property P4, for any vertex v in H, w.h.p. any search initiated at a random vertex for a random point traverses v with probability $C = O(\log^c n/n)$. The corresponding search path in $\mathcal{G}$ terminates either when it is successful or when the first red group is encountered; therefore, the search path is always

a subpath of the corresponding path of a search in H. Also, note that any extra edges added between red groups in the $\mathcal{G}$ (due to $S3$) do not affect how the search path proceeds given that the search path terminates at the first red group. Consequently, w.h.p., the corresponding group G_v is traversed with probability $O(\log^c n/n)$. $\square$

Let X be a random variable that is the probability that a search that begins at a randomly-chosen group for a random point in $[0, 1)$ fails. The randomness of X depends on which groups are red.

Lemma 2. *With high probability, $E(X) = O(p_f \log^c n)$.*

Proof. For any group G_v, let the random variable $X_v = \rho(G_v)$ if G_v is red, and 0 otherwise. By the definition of X, it is the case that $X \leq \sum_v X_v$. By linearity of expectation, $E(\sum_v X_v) = \sum_v E(X_v)$. Then, by Lemma 1, and the fact that each group is red with probability at most p_f, $\sum_v E(X_v) = O(p_f \log^c n)$. $\square$

Lemma 3. *With high probability, $X = O(p_f \log^c n)$.*

Proof. For any group G_v, let $X_v = \rho(G_v)$ if G_v is red, and 0 otherwise. We will bound $\sum_v X_v$, which we again note is always at least as large as X. Let $f(X_1, ..., X_n) = \sum_v X_v$, where we index the groups by 1 to n. By Lemma 1, for any X_i, $|f(..., X_i = x, ...) - f(..., X_i = x', ...)| = O(\log^c n/n)$. Thus, we can apply Theorem 2 with $c_i^2 = O(\log^{2c} n/n^2)$ for all $1 \leq i \leq n$. We have that:

$$Pr(|X - E(X)| \geq \lambda) \leq e^{-2\lambda^2 n/(d \log^{2c} n)}$$

for some constant $d > 0$. Let $\epsilon > 0$ be an arbitrarily small constant, and set $\lambda = \epsilon p_f \log^c n$, which is at least $\epsilon/\log^{k-c} n$, by S2. Then, we have:

$$Pr(|X - E(X)| \geq \lambda) \leq e^{-2\epsilon^2 n/(\log^{2k} n)}$$

By Lemma 2, w.h.p, $E(X) \leq d' p_f \log^c n$ for some constant $d' > 0$, which gives the result. $\square$

Lemma 4. *With high probability, any search from a random group to a random point in $[0, 1)$ succeeds with probability $1 - O(1/\log^{k-c} n)$.*

Proof. By Lemma 3, for n sufficiently large, w.h.p. $X = O(p_f \log^c n) = O(1/\log^{k-c} n)$, given that $p_f \leq 1/\log^k n$ by S2. $\square$

III. THE DYNAMIC CASE

We now consider the case where IDs can join and depart. We still make the assumptions about IDs described in Section I-C. Time is divided into disjoint consecutive windows of T **steps** called **epochs** indexed by $j \geq 1$.

Model of Joins and Departures. We assume that n IDs are always present even under churn; that is, when an ID departs, another is assumed to join; this is a popular model considered in much of the previous literature on tolerating $\Theta(n)$ Byzantine faults; for example, [7]–[10], [13], [17], [18], [23], [38]. Additionally, our results hold when the system size

is $\Theta(n)$ – that is, the size changes by a constant factor – but we omit these details in this extended abstract.

Recall that a good group G contains at most a $(1 + \delta)\beta$-fraction of bad IDs where $\beta > 0$ is a sufficiently small constant, and $\delta > 0$ is a tunably small constant depending only on sufficiently large n. We assume the following: in any epoch, at most an $(\epsilon'/2)$-fraction of good IDs depart any group, where $\epsilon' = 1 - 2(1 + \delta)\beta$. This value of ϵ' ensures that a good group retains a good majority over its lifetime. For ease of exposition in our analysis of the dynamic case, we revise our definition of a good group to be: a group G that begins with size $d_1 \ln \ln n \leq |G| \leq d_2 \ln \ln n$ and with at most a $(1 + \delta)\beta$-fraction of bad IDs, and retains a good majority. In practice, the length of an epoch, T, may be set appropriately by the system designers based on the expected rate of departures, and the value of ϵ' can be increased by increasing d_1.

Algorithmic Overview. In any epoch j, there are:
- two **old** group graphs $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$, each with n IDs.
- two **new** group graphs $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$, each with $\leq n$ IDs.

We emphasize that the use of two group graphs per epoch is critical. A naive approach is to use a *single* group graph in the current epoch in order to build a new group graph in the next epoch. However, this approach will fail because errors from bad groups will accumulate over time. Below we give some intuition for why.

Informally, in epoch j, we have a process where (1) bad groups build new bad groups, and (2) good groups build bad groups with some failure probability $p_f^j > 0$ that depends on the current number of bad groups. Therefore, in the next epoch $j + 1$, the population of bad groups has increased and so has p_f^{j+1}. Left unchecked, this increasing error probability will surpass the desired value of $1/\log^k n$. By using two group graphs, we can upper bound p_f^j by this value.

The new group graphs are built using the old group graphs over the n deletions and additions that occur in the current epoch j; we describe this in Subsection III-A. By the end of epoch j, the old group graphs $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$ are no longer needed, and the new ones $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$ are complete.

A. Building New Group Graphs

We describe how the new group graphs $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$ are created. We assume the correctness of two initial group graphs $\mathcal{G}_1^0$ and $\mathcal{G}_2^0$ with neighbor sets for blue groups correctly established. This aligns with prior literature in the area and more discussion is given in our online version [24] in Appendix X. Later, in Section III-B, we prove that w.h.p. this construction preserves ε-robustness.

Preliminaries. Assume the system is in epoch $j \geq 1$. Each good ID v already in the system uses the same ID in both $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$.

Any ID (already participating in the system or a newcomer) that wishes to participate in the next epoch $j + 1$ must begin generating an ID by the halfway point of the current epoch j. Generating this ID requires an expenditure of computational power as described in Section IV.

Recall from Section I-C that IDs expire after a tunable period of time. Upon creation, the new ID will be ***active*** throughout epoch $j + 1$ allowing v to initiate searches via G_v and for v to be added to other groups. When v's ID expires, the group G_v (this includes v) should remain in both old graphs for an additional T steps. During these steps, G_v will forward communications, but v cannot initiate searches using G_v, nor can v be added to new groups; we say that v's ID is ***passive***.

For any group G_v, if the leader v departs the system, G_v remains. That is, the members of G_v still persist as a the group G_v in their respective active or passive states. We discuss departures further below in the context of updating links.

IDs are assumed to know when the system came online (i.e. step 0).[12] Since T is set when the system is designed, any ID that wishes to join knows when the current epoch ends and the next one begins. Some synchronization between devices is implicit. In practice, this is rarely a problem given the near-ubiquitous Internet access (see the Network Time Protocol [34]) available to users.

A new ID joins the new group graph by a ***bootstrapping group*** denoted by G_{boot}.[13] Throughout, we assume that a joining ID knows a good bootstrapping group; we discuss this further in our online version [24] in Appendix IX.

Making a Group-Membership Request. In epoch j, group graphs $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$ are built using searches in $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$. An ID w uses the same ID in $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$.

G_w is added to $\mathcal{G}_1^j$ as follows. The i^{th} member of G_w is $\text{suc}(h_1(w, i))$ (recall the notation in Section I-C) in the old group graphs for $i = 1, ..., d_2 \ln \ln n$ where h_1 is a secure hash function and d_2 is defined with respect to group size in Section I-C. That is, in <u>both</u> $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$, a search for each successor of $h_1(w, i)$ is performed; this is executed by the bootstrapping group and $\text{suc}(h_1(w, i))$ is solicited for membership in G_w. Note that if different IDs are returned by the two searches, the successor to $h_1(w, i)$ is selected.

How is $\ln \ln n$ estimated? A standard technique for estimating $\ln n$ to within a constant factor is as follows. For u.a.r. IDs, the distance $d(u, v)$ between any two IDs u and v satisfies $\frac{\alpha''}{n^2} \leq d(u, v) \leq \frac{\alpha' \ln n}{n}$ w.h.p., depending only on sufficiently large positive constants α', α'' Therefore, w.h.p. $\ln \ln(\frac{1}{d(u,v)}) = \ln \ln(n) + O(1)$; this approach works even when an adversary decides to omit some (or all) of its IDs (see Chapter 4 in [50]), which is considered in Section III-B.

During epoch j, all IDs in $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$ are active — and will remain in a passive state over the next epoch $j + 1$ — and so can be used as members for new groups in $\mathcal{G}_1^j$.

Finally, a similar process occurs to build G_w in $\mathcal{G}_2^j$, except that a different secure hash function, h_2, is used. That is, a search for the successors of $h_2(w, i)$ occurs in both $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$. Note that the membership of G_w is likely different in each group graph.

---

[12]This is a fixed parameter included as part of the application, along with T, the hash functions, and various constants.

[13]As with much of the literature, we do not address concurrency. Since a join or departure requires updating only $\text{poly}(\log n)$ links in a group graph, we assume that there is sufficient time between events to do so.

Making a Neighbor Request. If w and u are neighbors in the input graph, then G_w and G_u should be neighbors in the group graph; recall that this entails all-to-all links between the good members of both groups. To set up the neighbors of G_w, G_{boot} performs a search on behalf of w to locate each such neighbor u in both old group graphs (again, favoring the successor if the results differ). In this way, G_{boot} allows u (and G_u) to learn about w and agree to set up a link in the respective group graph.

Verifying Requests. The adversary may attempt to have many good IDs join as neighbors or members of a bad group. This attack is problematic since good IDs will have resources consumed by maintaining too many neighbors or joining too many groups; this increases the state cost (see Section I). To prevent this attack, any such request must be verified:

Verifying a Group-Membership Request. When ID u in $\mathcal{G}_1^{j-1}$ is asked to become a member of group G_w in $\mathcal{G}_1^j$, ID u must verify that this request aligns with the linking rules; recall that this is assumed possible by property P3 of the input graph.

To do this verification, u performs a search on $h_1(w,i)$ in both group graphs $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$. If either returns u, then the request is considered verified and u becomes a neighbor of w; otherwise, the request is rejected. Note that u may erroneously reject a membership request; the impact on establishing good groups is addressed in Lemma 7. Conversely, u may erroneously accept a membership request, the impact of this on expected state cost is addressed in Lemma 10.

Verifying a Neighbor Request. An ID u that is asked to become a neighbor of ID w, and thus establish links between the members of G_u and G_w, must also verify this request. Similar to a group-membership request, u will determine via a search in $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$ whether u should indeed be a neighbor of w. If either search returns u, then the request is verified and u becomes a neighbor of w; otherwise, the request is rejected.

Note that u may erroneously reject a neighbor request; this is addressed in Lemma 8. Also, u may erroneously accept a request; the impact on expected state cost is addressed in Lemma 10.

Performing a Search. Throughout epoch j, each new ID w performs searches only in the old group graphs $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$. Searches are performed by forwarding the request to G_{boot}, and executing the search from that position. Since G_{boot} was active when w joined, then G_{boot} should remain – even if in a passive state – to facilitate searches for another T steps.

Over the duration of epoch j, a group G_w may not be able to reliably perform searches in the new group graphs $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$ since they are still under construction. For example, w might be the first ID to join $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$. Once epoch $j+1$ starts, the new group graphs $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$ are to be used. At this point, group G_w will initiate any search using its own links in these graphs, rather than relying on G_{boot} which may no longer be present in the system.

Updating Links. When a new ID (and its corresponding group) is added to the group graphs $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$ that are under construction, a group G_w must update its neighbor links in $\mathcal{G}_1^j$ and $\mathcal{G}_2^j$ if this new ID is a better match as a neighbor under the linking rules. This update is done via searches by a bootstrapping group in the old group graphs $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$.

Conversely, if G_w links to some group G_v whose members all depart — note that G_v must consist entirely of bad IDs given our model of churn — then G_w treats that link as null (until perhaps a join operation requires an update).[14]

Groups in the old group graphs $\mathcal{G}_1^{j-1}$ and $\mathcal{G}_2^{j-1}$ do not update links. For an old group graph, if all members of G_v depart, G_w treats that link as null until the group graph expires.

B. Analysis

In this section, we analyze the construction of new group graphs. Due to space constraints, some of our proofs are provided in our online version [24] in Appendix VII.

Properties P1-P4 of input graph H play an important role in the design of the corresponding group graph. However, a prerequisite to these properties is that *all* IDs are selected uniformly at random (see Section I-C), which is untrue if the adversary chooses to add only some of its bad IDs; for example, maybe only bad IDs in $[0, \frac{1}{2})$ are added by the adversary. Intuitively, this should not interfere with any of the properties; we now formalize this intuition.

In the following, we consider H$'$ to be a modified input graph which uses the same construction as H, but is subject to an adversary that only includes a subset of its IDs, from a larger set of u.a.r. IDs. We note that H$'$ is not necessarily a subgraph of H; the omission of bad IDs can result in a different topology.

Lemma 5. *Consider a graph H$'$, where the IDs are formed from two sets:*
- $\mathcal{N}_1$ *consists of at least $(1 - \beta)n$ IDs selected u.a.r. from $[0, 1)$.*
- $\mathcal{N}_2$ *is an arbitrary subset of at most βn IDs selected u.a.r. from $[0, 1)$.*

W.h.p., under the same construction as the input graph H, graph H$'$ has properties P1 - P4.

Throughout, the above result is assumed — that properties P1-P4 continue to hold if the adversary includes only a subset of its IDs — even if we do not always make it explicit. For example, P1 is important throughout our arguments/construction, P2 is used in Lemma 6, P4 in Lemma 9, and P3 in Lemma 10.

As described above, for an ID u, there are searches on random key values, via hashing under the random oracle assumption, in order to find members for group G_u. But if a key value maps to a bad ID, then this results in a bad member added to the group. We can bound the probability of this event as follows.

Lemma 6. *W.h.p. a random key value in an old group graph maps to a bad ID with probability at most $(1 + \delta'')\beta$ for an arbitrarily small constant $\delta'' > 0$ depending only on sufficiently large n.*

Let $q_f = O(1/\log^{k-c} n)$ be the probability that a search for a random key in an old group graph fails; recall Lemma 4.

---

[14]In practice, neighbors will periodically ping each other in order to check that neighbors are still alive.

Lemma 7. *Each group in a new group graph is bad with probability at most $O(q_f^2 d_2 \log \log n + 1/\log^{d'} n)$ for a tunable constant $d' > 0$ depending on d_2.*

Proof. For a new ID w, there are two ways in which a search for a member of G_w may result in a bad ID. First, a search for a group member may fail; that is, the search encounters a bad group. Given a point $h_1(w, i)$, the probability that both searches in the old group graphs fail is at most q_f^2. By a union bound, the probability of such a dual failure occurring over $d_2 \ln \ln n$ searches is $O(q_f^2 d_2 \log \log n)$.

Second, the search succeeds but returns $\text{suc}(h(w, i))$ where $\text{suc}(h(w, i))$ is a bad ID, even though its group is good. By the random oracle assumption, $h(w, i)$ is a random point. Thus, this event occurs with probability at most $(1 + \delta'')\beta$ by Lemma 6, for an arbitrarily small constant $\delta'' > 0$ given sufficiently large n. Over $d_2 \ln \ln n$ searches, the expected number of such events is at most $(1 + \delta')\beta d_2 \ln \ln n$. The probability of exceeding this expectation by more than a small constant factor (and adding too many bad IDs) is $O(1/\log^{d'} n)$ by a Chernoff bound, where the constant $d' > 0$ is tunable depending only on sufficiently large d_2.

Finally, the ID being asked to join may reject the request. This occurs if both searches used to verify the request fail. By a union bound, this happens with probability at most $O(q_f^2 d_2 \log \log n)$. □

A blue group G_w should link to all groups in the neighbor set L_w. Recall that $|L_w| = O(\log^\gamma n)$ for some constant $\gamma > 0$. If G_w (1) links to any group not in L_w, or (2) fails to link to any group in L_w, then G_w is said to be ***confused***. We can bound the probability of a confused group; the proof is provided in our online version [24] in Appendix VII.

Lemma 8. *Each group in a new group graph is confused independently with probability at most $O(q_f^2 \log^\gamma n)$.*

We now prove that w.h.p. each new group graph is ε-robust.

Lemma 9. *Assume the old group graphs are ε-robust and that the adversary has at most βn u.a.r. IDs. Then, for $k \geq 2c + \gamma$, w.h.p., each new group graph is ε-robust.*

Proof. In our analysis, a group that is bad or confused is a red group; otherwise, the group is good and not confused, and this a blue group. To show ε-robustness of the new group graph, we prove that the probability of creating a red group in the new group graph is at most $p_f \leq 1/\log^k n$ for a tunable constant $k > 0$. By Lemmas 7 and 8, each group is red independently with probability at most:

$$\leq O\left(q_f^2 \log^\gamma n\right) + O\left(q_f^2 d_2 \log \log n + 1/\log^{d'} n\right)$$
$$\leq O\left(\frac{\log \log n}{\log^{2(k-c)-\gamma} n} + \frac{1}{\log^{d'} n}\right) \leq \frac{1}{\log^k n}$$

The last line follows by setting d_2 to be sufficiently large such that d' exceeds k. Note that d_2 is fixed at the beginning and never needs to be changed throughout the lifetime of the network. Then, setting $k > 2c + \gamma$ to be a sufficiently large constant yields the necessary value of p_f to establish the inequality.

This implies that all but an $o(1)$-fraction of groups are good and not confused. Furthermore, once we have this bound on p_f, the remaining proof is equivalent to that of Lemma 4. □

The next lemma bounds the amount of state a good ID maintains due to (1) membership in groups, and (2) being a neighbor of a group. This is done by analyzing the verification process described in Section III-A; see [24] for the proof.

Lemma 10. *In expectation, each good ID w in a group graph is a member of $O(\log \log n)$ groups and maintains state on $O(|L_w|)$ groups that are either neighbors or have w as a neighbor.*

We can now prove Theorem 3.

Proof. Lemma 9 guarantees w.h.p. that in the new group graphs, all but a $1/\text{poly}(\log n)$-fraction of groups are good, and all but a $1/\text{poly}(\log n)$-fraction of IDs can search for all but a $1/\text{poly}(\log n)$-fraction of the resources.

Given that groups have size $O(\log \log n)$, it follows that group communication incurs $O((\log \log n)^2)$ messages. Recall that secure routing proceeds via all-to-all communication between members of groups and that searches have maximum length D (P1 in Section I-C). Thus, the message complexity is $O(D(\log \log n)^2)$.

To bound the expected state cost, we invoke Lemma 10. Each good ID w belongs to $O(\log \log n)$ groups in expectation. This implies $O((\log \log n)^2)$ expected state cost to keep track of the members of these groups.

In terms of links to and from other groups, w maintains state on $O(|L_w|)$ groups in expectation. The constructions for H defined in [39], [19], or [32] provide the properties P1-P4, but with a bound of $O(1)$ expected degree. Using any of these constructions, the state cost incurred by these neighboring groups is $O(\log \log n)$ in expectation. Thus, the total expected state cost is $O((\log \log n)^2) + O(\log \log n) = O((\log \log n)^2)$. Corollary 1 follows immediately. □

IV. COMPUTATIONAL PUZZLES

Up to this point, we have assumed that the adversary can inject into each new group graph at most βn bad IDs with u.a.r. values, and that these IDs can be verified and forced to expire after a period of time; recall our discussion in Section I-C. We now remove these assumptions. Given space constraints, we limit our discussion to the main ideas of how to use computational puzzles to guarantee these properties.

A. Generating an ID

All participants are assumed to know two secure hash functions, f and g, with range and domain $[0, 1)$ and that both hash functions satisfy the random oracle assumption.

In the current epoch i, ID w is assumed to possess a "globally-known" random string r_{i-1} of $\ell \ln n$ bits. By "globally-known", we mean known to all good IDs except the $1/\text{poly}(\log n)$-fraction from our earlier analysis. We motivate r_{i-1} and describe how it is generated in Subsection IV-B.

Starting at step $T/2$ in the current epoch, each good ID begins generating a new ID for use in the next epoch, as described below.

Description of ID Generation. To generate an ID, a good ID w selects a value σ_w of $\ell \ln n$ random bits (matching the length of r_{i-1}). Then, w XORs these two strings to get $\sigma_w \oplus r_{i-1}$, and checks if $g(\sigma_w \oplus r_{i-1}) \leq \tau$; if so, then $f(g(\sigma_w \oplus r_{i-1}))$ is a valid ID. We assume the value τ is set small enough such that w.h.p. $(1 \pm \epsilon)T/2$ steps are required to find a σ_w that satisfies this inequality, where $\epsilon > 0$ is a tunable (small) positive constant and $T > 0$ is a parameter set when the system is initialized.

The value of T can be large to amortize the cost of forcing IDs to depart (and possibly rejoin) over a long period of time; for example, $T > n$, since new group graphs are being built over T steps. Given an application domain, designers may estimate the rate of churn for their application and set a (loose) upper bound on n, then they can set T accordingly.

Why Use Two Hash Functions? Consider using a single secure hash function f to assign IDs; that is, if $g(x) < \tau$, then x is a valid ID. Then, for example, the adversary may restrict itself to small inputs x in order to confine its solutions to yielding small IDs. In other words, the IDs obtained by the adversary will not be u.a.r. from $[0, 1)$. This can be solved via composing two secure hash functions, f and g, as described above. See [24] for the proof of the following:

Lemma 11. *W.h.p., the adversary generates at most $(1+\epsilon)\beta n$ IDs over $(1 \pm \epsilon)(T/2)$ steps and these IDs are u.a.r. in $[0, 1)$.*

We note that Lemma 11 implies that the adversary might be able to generate up to $3(1 + \epsilon)\beta n$ IDs for use in the next epoch: computing over the $T/2$ steps in the last half of the previous epoch and the T steps prior to the end of the current epoch. However, we can revise the adversary's power from β to $\beta/3$, and results in Sections II and III hold. Note that all such IDs will be invalidated when the next random string is created.

ID Verification. Upon receiving a message from some ID w, a good ID u *verifies* w's ID. This could be done naively by having w send σ_w to u who checks that $g(\sigma_w \oplus r_{i-1}) \leq \tau$ and that $f(g(\sigma_w \oplus r_{i-1}))$ evaluates correctly to the claimed ID (note that u already has r_i since it is globally-known). Unfortunately, this allows u to steal σ_w if u is bad.

To avoid this issue, we can use a zero-knowledge scheme for revealing the pre-image of the hashing; such a scheme is provided for the SHA family [25]. This allows w to prove the validity of σ_w without revealing it.

If w's ID fails verification, then u simply ignores w going forward. Note that w's current ID will not be valid in the next epoch since it is signed by the older string r_{i-1} (rather than the next globally-known random string r_i); that is, w's ID will have expired. IDs that are not verified are effectively removed from the system; they may consort with bad IDs, but they have no interactions with good IDs.

B. Generating Global Random Strings

Imagine if no random string was used in the creation of IDs described above in Subsection IV-A. The adversary would know the format of the ID-generation puzzles, and so could spend time computing a large number of IDs, and then use these IDs all at once to overwhelm the system at some future point. This is a ***pre-computation attack***.

Signing IDs with a random string prevents such an attack as it is impossible for the adversary to know far in advance how to generate IDs. We provide a protocol where random strings are generated and propagated in the system to be used in ID generation. Due to space constraints, this content is provided in our online version [24] in Appendix VIII. We show the following:

Lemma 12. *W.h.p., the protocol for propagating strings (i) guarantees that, for each good ID w, its string used for generating an ID is known to each good ID, (ii) the number of strings stored by each ID is $O(\ln n)$, and (iii) has message complexity $\tilde{O}(n \ln T)$.*

Note that, averaged over a sufficiently large epoch, this message cost is low.

V. CONCLUSION AND FUTURE WORK

We showed that groups of size $O(\log \log n)$ can be used to tolerate a powerful Byzantine adversary. Our result utilizes PoW to limit the number of IDs the adversary controls; however, this imposes a computational overhead on participants. An open question is whether the computational costs can be reduced. Might there be a way to avoid the continual solving of puzzles? Is there an approach that would only utilize puzzle solving when malicious IDs are present?

Another problem that deserves attention is providing a detailed mechanism for bootstrapping in the presence of a Byzantine adversary. Such a result would likely benefit prior work as well as our own; we discuss this in our online version [24] in Appendix IX.

Finally, is it possible to show a lower bound on group size of $\Omega(\log \log n)$? We provided intuition for why this may be the smallest group size that admits strong security guarantees, but proving this appears challenging.

REFERENCES

[1] J. Augustine, A. R. Molla, E. Morsy, G. Pandurangan, P. Robinson, and E. Upfal. Storage and Search in Dynamic Peer-to-peer Networks. In *Proceedings of the Twenty-fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 53–62, 2013.

[2] J. Augustine, G. Pandurangan, and P. Robinson. Fast Byzantine Agreement in Dynamic Networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pages 74–83, 2013.

[3] J. Augustine, G. Pandurangan, and P. Robinson. *Fast Byzantine Leader Election in Dynamic Networks*, pages 276–291. 2015.

[4] J. Augustine, G. Pandurangan, P. Robinson, S. Roche, and E. Upfal. Enabling Robust and Efficient Distributed Computation in Dynamic Peer-to-Peer Networks. In *Proceedings of the IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 350–369, 2015.

[5] J. Augustine, G. Pandurangan, P. Robinson, and E. Upfal. Towards Robust and Efficient Computation in Dynamic Peer-to-peer Networks. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 551–569, 2012.

[6] B. Awerbuch and C. Scheideler. Chord++: Low-Congestion Routing in Chord. www.cs.jhu.edu/~scheideler/papers/chord4.ps.gz, 2004.

[7] B. Awerbuch and C. Scheideler. Group Spreading: A Protocol for Provably Secure Distributed Name Service. In *Proceedings of the 31^{st} International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 183–195, 2004.

[8] B. Awerbuch and C. Scheideler. Robust Random Number Generation for Peer-to-Peer Systems. In *Proceedings of the 10^{th} International Conference On Principles of Distributed Systems (OPODIS)*, pages 275–289, 2006.

[9] B. Awerbuch and C. Scheideler. Towards a Scalable and Robust DHT. In *Proc. of the 18^{th} ACM Symposium on Parallelism in Algorithms and Architectures*, pages 318–327, 2006.

[10] B. Awerbuch and C. Scheideler. Towards Scalable and Robust Overlay Networks. In *Proc. of the 6^{th} International Workshop on Peer-to-Peer Systems (IPTPS)*, 2007.

[11] M. Bellare and P. Rogaway. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *Proc. of the 1st ACM Conference on Computer and Communications Security*, pages 62–73, 1993.

[12] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach. Secure Routing for Structured Peer-to-Peer Overlay Networks. In *Proceedings of the 5^{th} Usenix Symposium on Operating Systems Design and Implementation (OSDI)*, pages 299–314, 2002.

[13] M. Datar. Butterflies and Peer-to-Peer Networks. In *Proceedings of the 10^{th} European Symposium on Algorithms (ESA)*, pages 310–322, 2002.

[14] J. Douceur. The Sybil Attack. In *Proceedings of the Second Internation Peer-to-Peer Symposium (IPTPS)*, pages 251–260, 2002.

[15] D. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 1st edition, 2009.

[16] Ethereum. Ethereum: Blockchain App. Program, 2016. https://www.ethereum.org/.

[17] A. Fiat and J. Saia. Censorship Resistant Peer-to-Peer Content Addressable Networks. In *Proceedings of the 13^{th} ACM Symposium on Discrete Algorithms (SODA)*, pages 94–103, 2002.

[18] A. Fiat, J. Saia, and M. Young. Making Chord Robust to Byzantine Attacks. In *Proceedings of the 13^{th} European Symposium on Algorithms (ESA)*, pages 803–814, 2005.

[19] P. Fraigniaud and P. Gauron. D2B: A De Bruijn Based Content-addressable Network. *Theor. Comput. Sci.*, 355(1):65–79, Apr. 2006.

[20] J. Garay, A. Kiayias, and N. Leonardos. The Bitcoin Backbone Protocol: Analysis and Applications. In *Proceedings of the Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 281–310, 2015.

[21] R. Guerraoui, F. Huc, and A.-M. Kermarrec. Highly Dynamic Distributed Computing with Byzantine Failures. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 176–183, 2013.

[22] D. Gupta, J. Saia, and M. Young. Proof of Work Without All the Work. In *Proceedings of the 19^{th} International Conference on Distributed Computing and Networking (ICDCN)*, page To appear, 2018.

[23] K. Hildrum and J. Kubiatowicz. Asymptotically Efficient Approaches to Fault-Tolerance in Peer-to-Peer Networks. In *Proceedings of the 17^{th} Intl. Symposium on Distributed Computing*, pages 321–336, 2004.

[24] M. O. Jaiyeola, K. Patron, J. Saia, M. Young, and Q. M. Zhou. Tiny Groups Tackle Byzantine Adversaries. *CoRR*, abs/1705.10387, 2017.

[25] M. Jawurek, F. Kerschbaum, and C. Orlandi. Zero-knowledge Using Garbled Circuits: How to Prove Non-algebraic Statements Efficiently. In *Proceedings of the 2013 ACM Conference on Computer Communications Security (CCS)*, pages 955–966, 2013.

[26] A. Kapadia and N. Triandopoulos. Halo: High-Assurance Locate for Distributed Hash Tables. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2008.

[27] J. Knockel, G. Saad, and J. Saia. Self-Healing of Byzantine Faults. In *Proceedings of the International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 98–112, 2013.

[28] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(3):382–401, 1982.

[29] D. Li, X. Lu, and J. Wu. FISSIONE: A Scalable Constant Degree and Low Congestion DHT Scheme Based on Kautz Graphs. In *Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.*, volume 3, pages 1677–1688, 2005.

[30] F. Li, P. Mittal, M. Caesar, and N. Borisov. SybilControl: Practical Sybil Defense with Computational Puzzles. In *Proceedings of the Seventh ACM Workshop on Scalable Trusted Computing*, pages 67–78, 2012.

[31] L. Luu, V. Narayanan, C. Zheng, K. Baweja, S. Gilbert, and P. Saxena. A Secure Sharding Protocol For Open Blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 17–30, 2016.

[32] D. Malkhi, M. Naor, and D. Ratajczak. Viceroy: a scalable and dynamic emulation of the butterfly. In *Proceedings of the 21^{st} ACM Symposium on Principles of Distributed Computing (PODC)*, pages 183–192, 2002.

[33] S. Micali. ALGORAND: The Efficient and Democratic Ledger. *CoRR*, abs/1607.01341, 2016.

[34] D. L. Mills. Internet Time Synchronization: The Network Time Protocol. *IEEE Transactions on Communications*, 39(10):1482–1493, 1991.

[35] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge Univ. Press, 1995.

[36] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. http://bitcoin.org/bitcoin.pdf.

[37] A. Nambiar and M. Wright. Salsa: A Structured Approach to Large-Scale Anonymity. In *Proc. of the 13^{th} ACM Conference on Computer and Communications Security*, pages 17–26, 2006.

[38] M. Naor and U. Wieder. A Simple Fault Tolerant Distributed Hash Table. In *Proceedings of the Second International Workshop on Peer-to-Peer Systems (IPTPS)*, pages 88–97, 2003.

[39] M. Naor and U. Wieder. Novel architectures for P2P applications: the continuous-discrete approach. In *Fifteenth ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 50–59, 2003.

[40] N. I. of Standards and Technology. Secure Hash Standard. http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.180-4.pdf, 2015.

[41] B. Parno, D. Wendlandt, E. Shi, A. Perrig, B. Maggs, and Y.-C. Hu. Portcullis: Protecting Connection Setup from Denial-of-Capability Attacks. In *Proceedings of the ACM SIGCOMM 2007 Conference*, pages 289–300, 2007.

[42] R. Rodrigues and B. Liskov. Rosebud: A Scalable Byzantine-Fault-Tolerant Storage Architecture. Technical Report TR/932, MIT LCS, December 2003.

[43] G. Saad and J. Saia. Self-Healing Computation. In *Proceedings of the International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 195–210, 2014.

[44] J. Saia, A. Fiat, S. Gribble, A. Karlin, and S. Saroiu. Dynamically Fault-Tolerant Content Addressable Networks. In *Proceedings of the 1^{st} International Workshop on Peer-to-Peer Systems*, 2002.

[45] J. Saia and M. Young. Reducing Communication Costs in Robust Peer-to-Peer Networks. *Information Processing Letters*, 106(4):152–158, 2008.

[46] C. Scheideler and S. Schmid. A Distributed and Oblivious Heap. In *Proceedings of the 36^{th} International Colloquium on Automata, Languages and Programming (ICALP)*, pages 571–582, 2009.

[47] S. Sen and M. J. Freedman. Commensal Cuckoo: Secure Group Partitioning for Large-Scale Services. *ACM SIGOPS Operating Systems*, 46(1):33–39, Feb. 2012.

[48] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications. In *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*, pages 149–160, 2001.

[49] A. Yao. How to Generate and Exchange Secrets. In *Proceedings of the 27th Annual Symposium on Foundations of Computer Science*, pages 162–167, 1986.

[50] M. Young. Making Chord Robust to Byzantine Attacks. Master's thesis, University of New Mexico, 2006.

[51] M. Young, A. Kate, I. Goldberg, and M. Karsten. Towards Practical Communication in Byzantine-Resistant DHTs. *IEEE/ACM Transactions on Networking*, 21(1):190–203, Feb. 2013.

[52] H. Yu. Sybil Defenses via Social Networks: A Tutorial and Survey. *SIGACT News*, 42(3):80–101, Oct. 2011.

[53] B. Y. Zhao, J. Kubiatowicz, and A. D. Joseph. Tapestry: An Infrastructure for Fault-Resilient Wide-Area Location and Routing. Technical Report UCB//CSD-01-1141, University of California at Berkeley Technical Report, April 2001.

APPENDIX

VI. BACKGROUND MATERIAL

In this section, we provide background material that may provide useful context for our results.

The Search Algorithm. We treat the search algorithm abstractly; it takes as input a value in $[0, 1)$ and returns a ID.

In practice, returning a ID implies returning the information necessary for contacting that corresponding owner of the ID (i.e., the machine) over the Internet (for example, learning the port number and IP address).

A search initiated by ID w is a request that will typically require contacting other IDs on its way to the final recipient y, who holds the resource being sought by w. Typically, searches are recursive, so w's request is forwarded through multiple IDs before reaching y. Alternatively, a search can be iterative, so w contacts a sequence of IDs each of which informs w how to make partial progress towards reaching y. This is not terminology particular to this area, the reader will likely have been exposed to these ideas in a networks course, perhaps in the context of name resolution queries under the Domain Name System. Both w and y, along with any IDs that participate in the search (by forwarding or resolving part of the search) are said to be IDs that are ***traversed*** by the search.

Resources. To flesh out the notions of a ***resource*** and a ***nearest-clockwise ID***, we consider the popular Chord overlay [48]. Each ID maintains links to $O(\log n)$ neighbors, and these neighbors are used to perform recursive searches. Assume that ID w wants to search for, say, a song file; that is, the song file is the ***resource***.

To find this resource, w applies a globally-known hash function to the title of the song to obtain a key value K. Note that there are no tricks here; the hash function is globally known because it comes with the software downloaded onto a user's computer that allows it to participate in the overlay, and it is not changed over the lifetime of the system. Additionally, one could generate key values for resources in other ways, but using the title is a typical example.

For simplicity, keys are assumed to be from a normalized range $[0, 1)$, and this is viewed as a ***unit ring*** where moving from 0 to 1 corresponds to moving ***clockwise*** around this ring.

This key value K is placed in the message sent from w to one of its neighbors, say ID u. If ID u holds the song file corresponding to K, ID u can return the song file directly to w. But, otherwise, u forwards the request onto one of its own neighbors. And so on, until a ID y is located that holds the file corresponding to K. At each forwarding step, the selection of which neighbor to forward the request to is a function of K; this is dependent on the system used (Chord, Viceroy, etc.).

ID y — in our example, the ID who holds the song file on its hard drive — is ***responsible*** for that resource. A simple and often-used rule is that the ID which is ***closest-clockwise*** to K is responsible for the resource; this is also referred to as the ***successor*** of K in the literature. That is, K is a point in the ID space $[0, 1)$ viewed as a unit ring, and if we slide clockwise from K on this unit ring, then the first ID encountered is responsible for the corresponding resource. With modifications, a similar framework applies to situations where the resource is a shared printer or some other service.

IDs and Neighbors. An ID is a value in $[0, 1)$; that is, a point somewhere on the unit ring. We describe how an ID is generated for our particular construction using PoW in Section IV-A.

Note that a participant (i.e., a machine) has a physical location in the real world, while an ID is a virtual location in the overlay. Therefore, a good machine can hold more than 1 ID. Indeed, our construction requires; a good machine u can hold an ID for two old group graphs, and an ID for two new group graphs.

What does this mean? Consider u's ID in an <u>old</u> group graph. This ID is linked to neighboring groups as specified by the input graph topology, and also holds links to members of groups to which it belongs. Similarly, u's ID in a <u>new</u> group graph is also linked to neighboring groups according to the appropriate topology, and to members of group to which it belongs. In other words, each ID is treated as a separate entity in the respective graph/network.

From the perspective of the machine u that generates these IDs, u must must commit resources to maintaining links for each ID it holds. Say u's ID links to w's ID, then in practice this means that u maintains state on, say, a TCP connection between u and w. Therefore, each ID held by u incurs a cost in terms of state maintenance and bandwidth.

In contrast to good machines, the adversary may attempt to create many, many IDs within a single group graph. One danger of this is that the number of bad IDs may now be larger than the number of good IDs; see mention of the Sybil attack under related work. In this case, we cannot hope to create groups with a good majority. Given this, a common assumption is that the adversary holds a minority of the computational power, and therefore computational puzzles allow us to mitigate this attack. The security of Bitcoin, Ethereum, and many blockchain constructions are based on this same assumption.

This is the convention we adopt too, and we prove in Section IV-A that the adversary holds (roughly) at most βn IDs in the system at any given time. This follows from our model section, where we describe the adversary as possessing at most a β-fraction of the computational power.

VII. Proofs for Section III-B

Lemma 5. *Consider a graph H' where the IDs are formed from two sets:*

- $\mathcal{N}_1$ *consists of at least $(1 - \beta)n$ IDs selected u.a.r. from $[0, 1)$.*
- $\mathcal{N}_2$ *is an arbitrary subset of at most βn IDs selected u.a.r. from $[0, 1)$.*

W.h.p., under the same construction as the input graph H, graph H' has properties P1 - P4.

Proof. View the ID space as a unit ring, and place on it the IDs from $\mathcal{N}_1$ and $\mathcal{N}_2$. Let the total number of IDs be m where $m \geq (1 - \beta)n$.

Moving clockwise from any ID, consider a contiguous interval of length $(\lambda \ln m)/m$ where $\lambda > 0$ is any constant. Since IDs in $\mathcal{N}_2$ are selected from a larger set of IDs u.a.r in $[0, 1)$, intuitively the adversary's choice of $\mathcal{N}_2$ cannot significantly change the density/sparseness of IDs on the ring. By Chernoff bounds, regardless of how $\mathcal{N}_2$ is selected, and for any $\lambda > 0$, the following holds:

- with probability at least $1 - (m)^{-\lambda/12}$, every interval contains at least $(\lambda/2) \ln m$ IDs, and
- with probability at least $1 - m^{-\lambda/12}$, every interval contains at most $(3\lambda/2) \ln m$ IDs.

A set of m IDs with locations on the ring that satisfy these two properties is a **λ-well-spread placement**. Observe that no matter how the adversary chooses $\mathcal{N}_2$, w.h.p. the adversary's influence on the distribution of IDs is characterized by some λ-well-spread placement. In other words, we can ignore the adversary since the issue now reduces to: what is the probability of a λ-well-spread placement that degrades a property in H$'$?

We argue by contradiction as follows. Recall the guarantee that the input graph H has some property P with probability at least $1 - 1/m^{c'}$ for a constant $c' > 0$ (Section II). Now assume there one or more λ-well-spread placements that violate this property P for H$'$, and that these occur with aggregate probability $1/m^d > 1/m^{c'}$ for some positive constant d. But this yields a contradiction since, with probability at least $1/m^d > 1/m^{c'}$, placing m IDs u.a.r. on the ring would yield one of these placements for the input graph H and, therefore, violate the guarantee of property P for H. $\qquad\square$

Lemma 6. *W.h.p. a random key value in an old group graph maps to a bad ID with probability at most $(1 + \delta'')\beta$ for an arbitrarily small constant $\delta'' > 0$ depending only on sufficiently large n.*

Proof. By property P2 of the input graph H, w.h.p. a randomly chosen ID in H is responsible for at most a $(1+\delta'')/n$-fraction of the key values for an arbitrarily small $\delta'' > 0$ depending on sufficiently large n (and, by Lemma 5, this holds even if the adversary does not add all of its bad IDs). Since the IDs of the adversary are u.a.r., the βn bad IDs are responsible for at most a $(1 + \delta'')\beta$-fraction of the key values. $\qquad\square$

Lemma 8. *Each group in a new group graph is confused independently with probability at most $O(q_f^2 \log^\gamma n)$*

Proof. Since the bootstrapping group is good, the only way in which a group G$_w$ is confused about a member of L_w is if (1) the two searches for a neighbor in the *old* group graphs both fail or (2) the corresponding group that is asked to be a neighbor erroneously rejects the request.

We analyze these two cases below, but first we elaborate on a subtle point. As described earlier, updates to links occur as the new group graphs are being constructed. For instance, if G$_v$ is the first group to be added to a new group graph, then G$_v$ will need to do searches to update its neighbors correctly. Updates may occur more than once over time as the new group graph is being constructed.

Throughout this construction, there may be times where G$_v$ does not link to the correct neighbors due to a search failing in the old group graphs. Importantly, the probability of failure in the old group graphs is bounded, and is not impacted by any confused groups in the new group graphs being built. So, it is possible for G$_v$ to be temporarily confused in the new group graph. But, importantly, we only care about the final selection

of G$_v$'s neighbors. Thus, when the final group is added to the new group graph, G$_v$'s subsequent update will lead to this final selection, and the probability of confusion at that point is what must be bounded.

We now bound the probability of cases (1) and (2). Applying Lemma 4, case (1) occurs with probability at most q_f^2 per element of L_w. Over $O(\log^\gamma n)$ potential neighbors, a union bound limits the probability of this occurring over all elements of L_w by $O(q_f^2 \log^\gamma n)$. For case (2), the group asked to be a neighbor will perform two searches per request and only reject if both fail; therefore, we get the same bound as in case (1). $\qquad\square$

Lemma 10. *In expectation, each good ID w in a group graph is a member of $O(\log \log n)$ groups and maintains state on $O(|L_w|)$ groups that are either neighbors or have w as a neighbor.*

Proof. We perform our analysis with respect to a good ID w. First, we analyze the state cost incurred by w due to its membership in various groups. Second, we analyze the state cost incurred by w due to (i) w's links to its neighbors, and (ii) those IDs that link to w as a neighbor.

State Cost of Group-Membership. The expected number of groups to which w belongs is at least $d'' \log \log n$, for some constant $d'' > 0$, given that such requests are distributed among all IDs u.a.r. (via the random oracle assumption) and each group requires $O(\log \log n)$ members.

For some ID u, consider a membership request for G$_u$ using the point $h_1(u, i)$. If this returns w (as the successor of this point $h_1(u, i)$), then w accepts the membership request. Given the ε-robustness guarantee in old group graphs, with probability at least $1 - O(1/\log^{k-c} n)$, this acceptance is correct (recall Lemma 4). In other words, the probability of accepting an erroneous member request is at most $1/\text{poly}(\log n)$ where we can tune this polynomial.

How many member requests does w receive? We reiterate a well-known argument: since IDs are u.a.r., the probability that two IDs are separated by more than a $(c'' \ln n)/n$ distance is at most $(1 - (c'' \ln n)/n)^n \leq 1/n^{\Theta(c'')}$, for some constant $c'' > 0$. Therefore, if $h_1(u, i)$ lies outside of the interval $\mathcal{I} = [w - (c'' \ln n)/n, w)$, then w.h.p. w is not the successor of $h_1(u, i)$ and cannot be a valid contender for membership.

Given that $h_1(u, i)$ is u.a.r. (given the random oracle assumption) the probability that $h_1(u, i)$ maps to this interval is $(c'' \ln n)/n$. By property P3, $i = O(\log^\gamma n)$ for some constant $\gamma > 0$, and so over at most n IDs that might own valid IDs in $\mathcal{I}$, there are $O(\log^\gamma n)$ such requests received by w; this holds w.h.p. by a standard Chernoff bound. It follows that w erroneously accepts $O(\log^\gamma n)/\text{poly}(\log n) = O(1)$ malicious requests in expectation so long as k is sufficiently large with respect to c.

State Cost of w's Neighbors. In each group graph, w links to $O(|L_w|)$ groups as neighbors.

State Cost of Other Groups Linking to w as a Neighbor. Via Lemma 5, properties P1 and P3 guarantees that w can

determine via searches whether it should indeed be a neighbor of some ID u, and there are at most $\text{poly}(\log n)$ such IDs u. Using the old group graphs, wherein ε-robustness is guaranteed w.h.p., w initiates a search to check that it should indeed be a neighbor; note that the adversary cannot lie about its IDs since they are verifiable. With a tunable probability at least $1 - O(1/\log^{k-c} n)$ ID w can detect if the request is erroneous. Therefore, in expectation, the number of erroneous acceptances is at most $o(|L_w|)$ so long as our constant k is sufficiently large (recall Lemma 4). $\qquad\square$

VIII. GENERATING AND DISSEMINATING GLOBAL RANDOM STRINGS

Lemma 11. *W.h.p., the adversary generates at most $(1+\epsilon)\beta n$ IDs over $(1\pm\epsilon)(T/2)$ steps and these IDs are u.a.r. in $[0,1)$.*

Proof. Since the adversary has βn computational power to expend over this epoch, w.h.p. it can generate at most $(1+\epsilon)\beta n$ solutions σ_v such that $g(\sigma_v \oplus r_{i-1}) \leq \tau$ within $(1\pm\epsilon)(T/2)$ steps where the constant $\epsilon > 0$ can be made arbitrarily small depending only on sufficiently large n. By the random oracle assumption, w.h.p. applying f to these solutions yields at most $(1+\epsilon)\beta n$ IDs u.a.r. from $[0,1)$. $\qquad\square$

Generating Random Strings. Over epoch i, all good IDs generate random strings and those corresponding to the smallest output under h (a secure hash function) are collected independently by each ID w to create a ***solution set*** R_i^w. To generate a string in epoch i, a ID w uses a string r_{i-1} — the globally-known string from the previous epoch — and an individually generated random string, s_w, to compute the ***output*** $t_w = h(s_w \oplus r_{i-1})$, where $\oplus$ indicates XOR.

Bins and Counters. To facilitate our discussion of how to propagate strings and ease our subsequent analysis, we describe a system of bins and counters maintained by each good ID w. The ***bins*** B_j correspond to intervals in the ID space where $B_j = [1/2^j, 1/2^{j-1})$ for $j = 1, 2, ..., b\ln(nT)$ where $b \geq 1$ is a sufficiently large constant. Since T is known and there are standard techniques for obtaining a constant-factor approximation to $\ln n$, calculating $\ln(nT) = \ln(n) + \ln(T)$ to within a constant factor is possible.[15]

Each bin B_j has an associated ***counter*** C_j. Consider that w receives a string s_u with corresponding output t_u that falls within the interval defined by B_j; we say that B_j ***contains*** t_u. If t_u is smaller than the other values w has seen so far contained in B_j, and $C_j \leq c_0 \ln n$ for some sufficiently large constant $c_0 \geq 1$, then w increments C_j and forwards the corresponding string s_u onto its neighbors. After $C_j = c_0 \ln n$, no value landing within B_j is ever forwarded.

The intuition is that, if $c \ln n$ strings are found with "record-breaking" outputs in B_j, then w.h.p. smaller strings exist

---

[15]A standard technique for estimating $\ln n$ to within a constant factor is as follows. For u.a.r. IDs, the distance $d(u,v)$ between any two IDs u and v satisfies $\frac{\alpha''}{n^2} \leq d(u,v) \leq \frac{\alpha' \ln n}{n}$ w.h.p., depending only on sufficiently large positive constants α', α'' Therefore, w.h.p. $\ln(\frac{1}{d(u,v)}) = \Theta(\ln n)$ and this holds even when an adversary decides to omit some (or all) of its IDs (see Chapter 4 in [50]).

with outputs belonging to B_{j+1}. In other words, those strings corresponding to B_j will not be candidates for a globally-known string, and so they can be ignored.

Protocol for Propagating Strings. The propagation of strings is broken into ***phases*** which make up the first half of an epoch. We describe the protocol for good IDs (although bad IDs can deviate arbitrarily).

Phase 1 executes over steps 1 to $T/2 - 2d' \ln n$ for a constant $d' > 1$ of the current epoch i. Over this time, each ID w generates random strings with associated outputs. After Phase 1 ends, IDs no longer generate new random strings.

Phase 2 begins at step $T/2 - 2d' \ln n + 1$ and runs for $d' \ln n$ steps. Each ID w (using its group G_w) selects the string $s_w^{\min}$ with the smallest output $t_w^{\min}$ that was generated in Phase 1, and then sends $s_w^{\min}$ its neighbors. ID w updates the corresponding bin and counter, as described earlier.

Each neighbor u verifies $s_w^{\min}$. Using $t_w^{\min}$, ID u decides whether to forward $s_w^{\min}$ to its own neighbors (except for w) and, if so, updates the corresponding bin and counter; otherwise, u ignores this value. At the end of Phase 2, each ID w selects the string with the smallest corresponding output it has seen so far; this is denoted by $\boldsymbol{s_w^{i*}}$.

Phase 3 starts at step $T/2 - d' \ln n + 1$ and runs for the final $d' \ln n$ steps. Over these steps, IDs no longer generate new strings, although they will still propagate them according to the above rules.

At the end of the phase, each ID w creates its solution set R_i^w in the following way. ID w finds the largest j for which B_j contains at least one element. Then, w collects the corresponding string, and all the corresponding strings that have outputs contained in subsequent bins for decreasing j, until there are $d_0 \ln n$ elements; the collection of these strings form R_i^w.

This concludes the propagation protocol. We note that immediately (at step $T/2 + 1$) ID w will start generating a new ID signed with the string s_w^{i*} chosen in Phase 2.

Discussion. The adversary may prevent good IDs from agreeing on the same solution set. As mentioned in Subsection IV-A, a $1/\text{poly}(\log n)$-fraction may be unable to partake in the propagation process even with our secure routing, and their loss is already incorporated into our analysis in Subsection III-B. Therefore, we address the ***giant component*** of $(1 - 1/\text{poly}(\log n))n$ good IDs that can reach each other; going forward, we mean this set of IDs when referring to the "good IDs".

The critical source of disagreement between these good IDs is that the adversary may delay releasing a string s' (or multiple strings) with a small output. For example, if this occurs right before the end of Phase 2, then only a subset of good IDs receive s' and their respective solution sets differ from the other good IDs.

We sketch how this disagreement is handled, but first we address the simpler case where there is no adversarial interference.

With No Adversary: Note that the propagation of a string in the giant component requires at most $d' \ln n$ steps. Therefore,

since all IDs send their string at the beginning of Phase 1, then by the end of Phase 2, all IDs accept the same set of strings and agree on the minimum string.

Furthermore, in Phase 3, nothing will occur (since no strings are released late) and so any IDs w and u are guaranteed w.h.p. to have $R_i^w = R_i^u$. What are the outputs corresponding to these solution sets? There are $\Theta(n)$ IDs computing for $\Theta(T)$ steps, so the smallest output in a set R_i^w is $\Theta(\frac{1}{nT})$ and w.h.p. no larger than $O(\frac{\ln n}{nT})$.

With an Adversary: The adversary can propagate a string s' with a small output late in Phase 2.[16] If w receives s' while u does not, then $R_i^w \neq R_i^u$. We argue that w.h.p. for good IDs in the giant component that (1) the size of each solution set remains bounded by $\Theta(\ln n)$, and (2) that the string s_w^{i*} used by each good ID w belongs to every other good IDs' solution set; these two properties enable efficient and correct verification (described below).

How many solutions s' could w receive and add to R_i^w? As noted above, this solution set will hold outputs of value $O(\frac{\ln n}{nT})$. Since the adversary has bounded computational power of βn, w.h.p. there cannot be more than $d'' \ln n$ solutions with output value $\Theta(\frac{1}{nT})$ for some constant $d'' > 0$. This is true even if the adversary computes over the entire epoch. We set the constant c_0 used in the bin counters such that $c_0 \geq d''$ in order to make sure that no smallest values are omitted.

Now, consider two good IDs w and u in the giant component. Assume that w selects s_w^{i*} (recall that this occurs at the end of Phase 2), but that s_w^{i*} is not present in good ID u's solution set R_i^u by the end of Phase 3; we will derive a contradiction. If s_w^{i*} originated from a good ID, then u received s_w^{i*} by the end of Phase 3 since $2d' \ln n$ steps is more then sufficient for the propagation of a string in the giant component (since this is the portion of the network to which we can search). Else, s_w^{i*} originated from the adversary or a ID outside the giant component. Since s_w^{i*} was held by w by the end of Phase 2 (this inclusion could have been delayed by the adversary until the final step of Phase 2), the addition $d' \ln n$ steps in Phase 3 would have allowed s_w^{i*} to reach u and be added to R_i^u. In either case, this yields the contradiction.

Finally, what is the message complexity of the propagation protocol? Recall that for each bin, the associated counter restricts to $O(\ln n)$ the number of times a ID forwards a string to its neighbors. Given that there are $O(\ln(nT))$ bins, the total number of times a ID can forward a string is $O(\ln(n) \ln(nT))$. The number of messages sent between any pair of neighboring groups is $O(|G|^2) = O((\log \log n)^2)$ and the degree in the group graph is $O(\text{poly}(\log n))$. Therefore, the total message complexity over $O(n)$ IDs is $\tilde{O}(n \ln T)$ where $\tilde{O}$ accounts for poly$(\log n)$ terms.

The above discussion supports the following:

Lemma 12. *With high probability, the protocol for propagating strings (i) guarantees that, for each good ID w in the*

[16]The adversary may also delay a string from a good ID outside the giant component, which amounts to the same problem since the adversary controls when this string is released into the giant component.

component, s_w^{i*} is contained within the solution set of every good ID in the component, (ii) $|R_i^w| = O(\ln n)$, and (iii) has message complexity $\tilde{O}(n \ln T)$.

Verifying IDs. For simplicity, our discussion of ID verification in Subsection IV-A assumed that a single r_{i-1} was agreed upon. However, not much changes when using solution sets.

To generate an ID for use in epoch $i + 1$, ID w uses s_w^{i*} to sign its ID. By the above discussion, we are guaranteed w.h.p. that s_w^{i*} belongs to the solution set of each good ID. Therefore, a good ID u that wishes to verify w's new ID checks whether this ID was signed by *any* of the strings in R_i^u; this requires checking only $O(\ln n)$ elements by the above discussion.

IX. BOOTSTRAPPING GROUPS

In prior work, an ID v joins the network by contacting $O(\log n)$ members of a group which, by virtue of its size, has a good majority with high probability. In our construction, we make a similar assumption, except that v contacts $O(\log n / \log \log n)$ groups selected uniformly at random, each of size $O(\log \log n)$. With high probability, the collection of the $O(\log n)$ IDs belonging to these groups will contain a good majority. Therefore, assuming that v's ID is verified, these IDs can fulfill the role of a bootstrapping group G_{boot}, as described in Section III.

During this bootstrapping process, we note that for input graphs with $O(1)$ expected degree, this approach will increase the expected state-maintenance cost to be $O(\log n)$. In the case of input graphs that have $O(\log n)$ degree, our expected state-maintenance cost remains unchanged.

X. SYSTEM INITIALIZATION

How are the group graphs $\mathcal{G}_1^0$ and $\mathcal{G}_2^0$ created? Almost all of the literature treats the problem of building secure overlays in the following manner: (1) prove that an overlay construction yields security guarantees, and then (2) prove these security guarantees can be maintained for each departure/join event. This is typically challenging in its own right, and the issue of how the system is initially placed into a state that satisfies (1) is often not dealt with explicitly.

As motivation, one might consider a system whose development is shepherded by some central authority, or some small group of trustworthy participants, that handles admission control and the assignment of group membership. Once the system reaches some threshold size (sufficient for the w.h.p. arguments), the system becomes fully decentralized.

A notable exception is [21], which gives an explicit solution. The authors specify a protocol that allows all good IDs to learn about the existence of all good IDs; this incurs a communication cost of $O(n \cdot |E|)$ where $|E|$ is the number of edges in the overlay being created. Then, since all good IDs are now aware of each other, a subset of the IDs — called a "representative cluster" — is elected via running a Byzantine Agreement (BA) protocol; this has a communication complexity of soft-$O(n^{3/2})$. The representative cluster, which is shown to have an honest majority, is then responsible for establishing cluster membership, informing each ID about its fellow cluster members, and setting up links between the

clusters. Therefore, the system initialization can be achieved w.h.p by this one-time "heavy-weight" procedure, after which the security guarantees can be maintained.

We believe that a similar approach would work in our setting and allow for the creation of $\mathcal{G}_1^0$ and $\mathcal{G}_2^0$. Since the problem of initialization pertains to much of the literature in this area, a pertinent question is whether one can improve upon the above scheme. We believe this is a problem in its own right and an interesting avenue for future work.

Setting k. We end this section by addressing a subtlety in setting the value of k sufficiently large; this was referred to in Section II.

As described in the main text, we assume that the neighbor sets in $\mathcal{G}_1^0$ and $\mathcal{G}_2^0$ are set up correctly for blue groups. Thus, p_f corresponds to the probability of group being bad, and we can set k to be as large a constant as desired by setting d_1 sufficiently large. By Lemma 9, $k \geq 2c + \gamma$ is sufficient.

Now, consider the next epoch. The construction of the two new group graphs are proven to preserve $p_f \leq 1/\log^k n$, where a red group corresponds to either a group that is bad *or* that has its neighbor set incorrectly established. Therefore, w.h.p., for subsequent epochs, groups graphs will be ε-robustness.