Resource-Competitive Sybil Defenses

Diksha Gupta[1], Jared Saia[2], and Maxwell Young[3]

[1]Dept. of Computer Science, University of New Mexico, NM, USA
dgupta@unm.edu
[2]Dept. of Computer Science, University of New Mexico, NM, USA
saia@cs.unm.edu
[3]Computer Science and Engineering Dept., Mississippi State University, MS, USA
myoung@cse.msstate.edu

Abstract

Proof-of-work (PoW) is an algorithmic tool used to secure networks by imposing a computational cost on participating devices. Unfortunately, traditional PoW schemes require that correct devices perform significant computational work in perpetuity, even when the system is not under attack.

We address this issue by designing general PoW protocols that ensure two properties. First, the fraction of identities in the system that are controlled by an attacker is a minority. Second, the computational cost of our protocol is comparable to the cost of an attacker. In particular, we present an efficient algorithm, GMCOM, which guarantees that the average computational cost to the good ID per unit time is $O(J + \sqrt{T(J+1)})$, where J is the average number of joins by the good IDs and T is the average computational spending of the adversary. Additionally, we discuss a precursor to this algorithm, CCOM, which guarantees an average computational cost to good IDs per unit time of $O(J + T)$.

We prove a lower bound showing that GMCOM's spending rate is asymptotically optimal among a large family of algorithms. Finally, we provide empirical evidence that our algorithms can be significantly more efficient than previous defenses under various attack scenarios.

1 Introduction

Twenty-five years after its introduction by Dwork and Naor [62], ***proof-of-work (PoW)*** is enjoying a research renaissance. Succinctly, PoW is an economic tool to prevent abuse of a system by requiring participants to solve computational puzzles in order to access system resources or participate in group decision making. In recent years, PoW is playing a critical role in cryptocurrencies such as Bitcoin [112], and other blockchain technologies [34, 48, 57, 64, 85, 101, 125]. Yet, despite success with Bitcoin and other cryptocurrencies, PoW has not found a widespread application in mitigating malicious behavior in networks. This is despite numerous proposals for PoW-based defenses [37, 43, 78, 87, 100, 104, 121, 141, 142].

A Barrier to Widespread Use. A major impediment to the widespread use of PoW is "the work". In particular, current PoW approaches have significant computational overhead, given that puzzles must always be solved, *even when the system is not under attack*. This non-stop resource burning translates into a substantial energy — and, ultimately, a monetary – cost [47, 63, 135].[1] Consequently, PoW approaches are currently used primarily in applications where participants have a financial incentive to continually perform

---

[1]For example, in 2018, the Economist calculated that Bitcoin consumes at least 22 terawatt-hours of electricity per year, or enough to power Ireland [63].

work, such as cryptocurrencies. This is a barrier to wide-spread use of a technique that has the potential to fundamentally advance the field of cybersecurity.

In light of this, we seek to reduce the cost of PoW systems and focus on the following question: **Can we design PoW systems where resource costs are low in the absence of attack, and grow slowly with the effort expended by an attacker?**

In this paper, we design and analyze algorithms that answers this question in the affirmative. We begin by discussing a straightforward algorithm, **CCOM** in Section 3, followed by some minor enhancements to this approach in Section 4. Next, we present an improved algorithm, **GMCOM** (Section 5) and analyze it (Section 6). Then we prove a matching lower bound for this class of algorithms in Section 7. We present empirical results which complement our theoretical work in Section 8. Finally, we conclude and offer directions for future research in Section 9.

1.1 Our Model

Our system consists of virtual *identifiers (IDs)*, and an attacker (or *adversary*). Each ID is either good or bad. Each *good* ID follows our algorithm, and all *bad* IDs are controlled by the adversary.

1.1.1 Communication

All communication among good IDs occurs through a broadcast primitive, denoted by **DIFFUSE**, which allows a good ID to send a value to all other good IDs within a known and bounded amount of time, despite the presence of an adversary. We assume that when a message is diffused in the network, it is not possible to determine which ID initiated the diffusion of that message. Such a primitive is a standard assumption in PoW schemes [31, 68, 70, 103]; see [105] for empirical justification. We assume that each message originating at a good ID is signed by the ID's private key.

A *round* is the amount of time it takes to solve our easiest computational puzzle plus the time to communicate the solution to the rest of the network via DIFFUSE. We assume that a constant number of rounds is sufficient to perform Byzantine consensus (see Section 1.3 for definition of Byzantine consensus and an algorithm to solve it).

As is standard, we assume all IDs are synchronized. For simplicity, we initially assume that the time to diffuse a message is small in comparison to the time to solve computational puzzles.[2] However, we later relax this assumption to account for a network latency which is upper bounded by Δ in Section 4.1. We pessimistically assume that the adversary can send messages to any ID at will, and that it can read the messages diffused by good IDs before sending its own.

1.1.2 Puzzles

We assume a source of computational puzzles of varying difficulty, whose solutions cannot be stolen or pre-computed. This is a common assumption in PoW systems [13, 100, 112]. We now provide an overview of the standard way in which this assumption is achieved; more details on the puzzle construction are provided in Section 3.2.

All IDs have access to a hash function, h, about which we make the *random oracle assumption* [26, 95]. Succinctly, this assumption is that when first computed on an input, x, $h(x)$ is selected independently and uniformly at random from the output domain, and that on subsequent computations of $h(x)$ the same output value is always returned. We assume that both the input and output domains are the real numbers between

---

[2] A recent study of the Bitcoin network shows that the communication latency is 12 seconds in contrast to the 10 minutes block interval [50], which motivates our assumption of computational latency dominating the communication latency in our system.

0 and 1. In practice, h may be a cryptographic hash function, such as SHA-2 [120], with inputs and outputs of sufficiently large bit lengths.

Solving a puzzle requires that an ID find an input x such that $h(x)$ is less than some threshold. The input found is the ***puzzle solution***. Decreasing this threshold value will increase the difficulty, since one must compute the hash function on more inputs to find an output that is sufficiently small.

Finally, we must be able to address an adversary who attempts (1) to falsely claim puzzle solutions computed and transmitted by good IDs, and (2) to pre-compute solutions to puzzles. These details are deferred until Section 3.2 when we describe our algorithm.

1.1.3 Adversary

A single adversary controls all bad IDs. This pessimistically represents perfect collusion and coordination by the bad IDs. Bad IDs may arbitrarily deviate from our protocol, including sending incorrect or spurious messages. This type of adversary encapsulates the Sybil attack [60] (see Section 2).

The adversary controls an α-fraction of computational power, where $\alpha > 0$ is a small constant. Thus, in a single round where all IDs are solving puzzles, the adversary can solve an α-fraction of the puzzles. This assumption is standard in past PoW literature [13, 70, 112, 138].

Our algorithms employ digital signatures, but do not require any public key infrastructure. Further, we assume the adversary knows our algorithm, but does not know the private random bits of any good ID.

1.2 System Initialization

We assume an ***initialization phase*** over which initial IDs join the system. At the end of the initialization phase, a heavy-weight protocol (GENID, described in Section 3.1) is run, ensuring that (1) all IDs learn an initial system membership, S_0; and (2) no more than an α-fraction of the membership of S_0 is bad. The initial estimate of the join rate of good IDs is the size of S_0 divided by the time taken for the initialization phase.

1.2.1 Joins and Departures

The system is dynamic with IDs joining and departing over time, subject to the constraint that at most a constant fraction of the good IDs can join or depart in any round. Maintaining performance guarantees amidst many nodes joining and leaving a system is challenging in decentralized systems [15–19].[3] For CCOM we pessimistically assume that all join and departure events are scheduled in a worst-case fashion by the adversary. However, we assume that departing good IDs are selected uniformly at random from the set of all good IDs in the population.

After the initialization phase, we assume that all good IDs announce their departure to the network. In practice, this assumption could be relaxed through the use of ***heartbeat messages*** that are periodically sent out to indicate that an ID is still in the network.

The minimum number of good IDs in the system at any point is assumed to be at least $\boldsymbol{n_0}$. Our goal is to provide security and performance guarantees for $\boldsymbol{O(n_0^\gamma)}$ joins and departures of IDs, for any desired constant $\boldsymbol{\gamma \geq 1}$. In other words, the guarantees on our system hold with high probability (w.h.p.)[4] over this polynomial number of dynamic events. This implies security despite a system size that may vary wildly over time, increasing and decreasing polynomially in n_0 above a minimum number of n_0 good IDs.

---

[3]Informally this process of nodes joining and departing over time is often referred to as "churn".
[4]With probability at least $1 - n_0^{-c}$ for any desired $c \geq 1$.

We assume that at most an ϵ_a-fraction of good IDs may join in a single round, where $\epsilon_a > 0$ is a small constant. Similarly, we assume an ϵ_d-fraction of good IDs may depart in a single round, where $\epsilon_d > 0$ are small constants that depends on the constant number of rounds required to solve Byzantine consensus.

1.3   Problem Definition

We define the ***Defend-ID*** (**DEFID**) problem which consists of maintaining the following two security invariants.

Population Invariant: Ensure that the fraction of bad IDs in the entire system is less than $1/2$.

We use $1/2$ for ease of presentation, but there is nothing special about this fraction. We can adjust this to be any fraction less than 3α by adjusting our algorithm as discussed in Section 4.2. Achieving the Population Invariant will help ensure that system resources – such as jobs executed on a server, or bandwidth obtained from a wireless access point – consumed by the bad IDs are (roughly) proportional to the adversary's computational power. Possible application domains include: content-sharing overlays [66, 133]; and open cloud services [11, 42, 109, 144].

Committee Invariant: Ensure there always exists a committee that (i) is known to all good IDs; (ii) is of *scalable* size i.e., of size $\Theta(\log n_0)$; and (iii) contains less than a $1/2$ fraction of bad IDs.

Again, there is nothing special about $1/2$, and this can be adjusted to a smaller fraction of bad IDs (see Section 9). Achieving the Committee Invariant ensures that the committee can solve the ***Byzantine consensus*** problem [98] in a scalable manner. In this problem, each good ID has an initial input bit. The goal is for (i) all good IDs to decide on the same bit; and (ii) this bit to equal the input bit of at least one good ID. Byzantine consensus enables participants in a distributed network to reach agreement on a decision, even in the presence of a malicious minority. Thus, it is a fundamental building block for many cryptocurrencies [36,65,70,77]; trustworthy computing [38,39,41,44,46,97]; peer-to-peer networks [5,119]; and databases [110, 124, 153].

Establishing Byzantine consensus via the use of committees is a common approach; for examples, see [70, 92, 103, 123]. There are algorithms that enable a set of x nodes to solve Byzantine consensus in constant time and sending $O(x^2)$ messages, when communication is synchronous, nodes have digital signatures, and the fraction of Byzantine nodes is strictly less than $1/2$. We make use of a recent algorithm by Abraham et al. [3].

1.4   Results for CCOM

We measure ***computational cost*** as the effort required to solve computational puzzles (see Section 1.1.2 for details).

We begin by presenting our results for CCOM (see Section 3). The ***system lifetime*** consists of a number of ID joins and departures that is polynomial in n_0. Let T be the ***adversarial spending rate***, which is the cost to the adversary for solving puzzles over the system lifetime. Let J^G denote the ***good ID join rate***, which is the number of good IDs that join during the system lifetime. Finally, the ***algorithmic spending rate*** is the total cost to the good IDs for solving puzzles over the system lifetime.

Theorem 1. *For $\alpha < 1/6$, with error probability polynomially small in n_0 over the system lifetime,* CCOM *solves* DEFID *with an algorithmic spending rate of $O(T + J^G)$.*

Note that the computational cost incurred by the good IDs grows slowly with the cost incurred by the adversary. When there is no attack on the system, the costs are low and solely a function of the number of good IDs; there is no excessive overhead. But as the adversary spends more to attack the system, the costs required to keep the system secure grow commensurately with the adversary's costs.

4

1.5 GMCOM Join and Departure Assumptions for Good IDs

Following are additional assumptions needed to achieve our main result for GMCOM. Time is divided into disjoint intervals called *epochs*, defined as follows.

Definition 1. *Let $\mathcal{G}_i$ be the set of good IDs in the system at the end of epoch i and let epoch 1 begin at system creation. Then, epoch i is defined as the shortest amount of time until $|\mathcal{G}_i - \mathcal{G}_{i-1}| \geq (3/4)|\mathcal{G}_i|$.*

Let ℓ_i denote the number of seconds in epoch i. Let ρ_i be the join rate of good IDs in epoch i; that is, the number of good IDs that join in epoch i divided by ℓ_i.

We make the assumptions A1 and A2 stated below. For ease of exposition, we use notation $c(\mathtt{A}, \mathtt{low})$ to indicate a positive constant at most 1 used to lower bound a quantity in assumption $\mathtt{A}$. Similarly, the constant $c(\mathtt{A}, \mathtt{high}) \geq 1$ is used to upper bound a quantity in assumption $\mathtt{A}$. By parameterizing our analysis later on with these constants, our results generalize to a variety systems where different bounds may be appropriate.

- **A1.** For all $i > 1$, $c(\mathrm{A1}, \mathtt{low})\rho_{i-1} \leq \rho_i \leq c(\mathrm{A1}, \mathtt{high})\rho_{i-1}$.

- **A2.** For any period of time within an epoch that contains at least 2 good join events, the good join rate during that period is between $c(\mathrm{A2}, \mathtt{low})\rho_i$ and $c(\mathrm{A2}, \mathtt{high})\rho_i$.

Informally, assumption A1 implies that the rate at which good IDs join is fairly "smooth", A2 implies that the "rate of change" in the set of good IDs is at least $\Omega(J_i)$ over the course of an epoch, and assumption A2 prohibits good IDs from being too "bunched up" during a small interval of time.

In practice, the performance of our algorithms will depend on these constants, which may differ between various network scenarios. For this reason, we parameterize our analysis of GMCOM (Sections 6.2 and 6.3). Later, in Section 8.1, we provide empirical results using data from real-world systems.

1.6 Results for GMCOM

Next, we present results for GMCOM, presented in Section 5, which ensures that the algorithmic spending rate is asymptotically less than the adversarial spending rate in the presence of an attack, whereas in the absence of an attack it is commensurate with the join rate of good IDs.

Theorem 2. *For $\alpha \leq 1/18$, with error probability polynomially small in n_0 over the system lifetime, GMCOM solves DEFID with an algorithmic spending rate of $O(\sqrt{T(J^G + 1)} + J^G)$.*

This result is complemented by the following lower bound. Define a ***purge-based*** algorithm to be any algorithm where (1) IDs pay a cost of $\Omega(1)$ to join; and (2) after a constant fraction of the population changes, all IDs must pay $\Omega(1)$ to remain in the system; otherwise, they are purged.

The execution of our algorithms, CCOM and GMCOM, consists of contiguous sequences of rounds called ***iterations***. An iteration that starts with round s ends when the number of join and departure events since round s exceeds a constant fraction of the number of IDs present at the start of round s.

Theorem 3. *For any purge-based algorithm, there is an adversarial strategy ensuring the following for any iteration. The algorithmic spending rate is $\Omega(\sqrt{T\, J^G} + J^G)$, where J^G is the good ID join rate and T is the algorithmic spending rate, over the iteration.*

1.6.1 A Corollary for Theorem 2

The following Corollary of Theorem 2 shows that the spending rate for the algorithm remains small, even when focusing on just a subset of iterations. To understand why this is important, consider a long-lived system which suffers a single, significant attack for a small number of iterations, after which there are no more attacks. The cost of any defense may be small when amortized over the lifetime of the system, but this does not give a useful guarantee on performance during the time of attack.

Let $\mathcal{I}$ be any subset of iterations that for integers x and y, $1 \leq x \leq y$, contains every iteration with index between x and y inclusive, and let L be the total length of time of those iterations. Let $\delta(\mathcal{I})$ be $|S_x - S_y|$; and let $\Delta(\mathcal{I})$ be $\delta(\mathcal{I})$ divided by the length of $\mathcal{I}$.

The **adversarial spending rate**, $T_{\mathcal{I}}$, is the cost to the adversary for solving puzzles whose solutions are used in any iteration of $\mathcal{I}$ divided by L. The **good ID join rate**, $J_{\mathcal{I}}^G$, is the number of good IDs that join over the iterations in $\mathcal{I}$ divided by L. Next, the **algorithmic spending rate** is the total cost to the good IDs for solving puzzles whose solutions are used in any iteration of $\mathcal{I}$ divided by L. Then, we state our result formally as:

Corollary 1. *For $\alpha \leq 1/18$, with error probability polynomially small in n_0 over the system lifetime, and for any subset of contiguous iterations $\mathcal{I}$,* GMCOM *has an algorithmic spending rate of* $O\left(\sqrt{T_{\mathcal{I}}\left(J_{\mathcal{I}}^G + 1\right)} + \Delta(\mathcal{I}) + J_{\mathcal{I}}^G\right)$.

1.7 Empirical Results

In Section 8, we empirically evaluate our algorithms and compare their performance to the state of the art. For a variety of networks, our algorithms significantly reduce computational cost. Additionally, we describe heuristics that further improve the performance of our basic algorithms.

2 Related Work

CCOM and GMCOM. Preliminary results on CCOM [80] and GMCOM [81] appeared previously. In this full version of these results, we make several additional contributions. First, improvements to the design of both algorithms have allowed us to provide a more streamlined presentation, and reduce the complexity of our arguments; the accompanying proofs of correctness and performance are now included in full. Second, the specification of GMCOM contains a significant modification for estimating the join rate of good IDs that corrects an error in our preliminary version [81]; the full correctness proofs for our new estimate are now provided in Section 6.2. Third, for GMCOM, we are able to remove an assumption on the behavior of good IDs that was previously required in our preliminary results [81]. Finally, in Section 8, we present additional experimental results that align with our theoretical guarantees.

2.1 Sybil Attack and Defense

The Sybil attack. Our result defends against the Sybil attack [60]. There is large body of literature on mitigating the Sybil attack (for example, see surveys [59, 86, 108, 118]), and additional work documenting real-world Sybil attacks [115, 136, 140, 146]. *Critically, none of these prior results ensure that computational costs for good IDs grow slowly with the cost incurred by the adversary.* We note that our algorithms can be characterized as **resource-competitive** [6, 27, 28, 54, 56, 71, 72, 74, 93, 152], in the sense that the algorithmic cost grows slowly with that of an attacker. However, for simplicity, we omit a discussion of resource-competitive algorithms since it is not critical to understanding our results. We refer interested readers to the survey [28].

We note that obtaining large numbers of machines is expensive: computing power costs money, whether obtained via Amazon AWS [9] or a botnet rental [12]. This motivates our desire to consider the attacker's cost when designing our algorithm's defense.

GENID and System Initialization. In the GENID problem, there is a set of good IDs, and an adversary that wants to impersonate many bad IDs. The adversary has an $\alpha-$fraction of computational power. All good IDs must agree on a set of IDs that contains (1) all good IDs, and (2) at most an $\alpha-$fraction of bad IDs. We make use of a solution to this problem in Section 3.1. Our algorithms call a solution to GENID only once, at system initialization. A number of algorithms have been found to solve GENID [13, 84, 90].

Radio-Resource Testing. In a wireless setting with multiple communication channels, Sybil attacks can be mitigated via *radio-resource testing* which relies on the inability of the adversary to listen to many channels simultaneously [73, 75, 111, 118]. However, this approach may fail if the adversary can monitor most or all of the channels. Furthermore, radio-resource testing suffers from the same issue discussed above, even in the absence of attack, the system must constantly perform tests.

Social Network Properties. There are several results that leverage social networks to yield Sybil resistance [99, 107, 143, 149–151]. However, social-network information may not be available in many settings. Another idea is to use network measurements to verify the uniqueness of IDs [24, 58, 69, 102, 131, 139], but these techniques rely on accurate measurements of latency, signal strength, or round-trip times, for example, and this may not always be possible. Containment strategies are examined in overlays [53, 129], but the extent to which good participants are protected from the malicious actions of Sybil IDs is limited.

Proof of Work and Alternatives. As a choice for PoW, computational puzzles provide certain advantages. First, verifying a solution is much easier than solving the puzzle itself. This places the burden of proof on devices who wish to participate in a protocol rather than on a verifier.

In contrast, bandwidth-oriented schemes, such as [138], require verification that a sufficient number of packets have been received before any service is provided to an ID; this requires effort by the verifier that is proportional to the number of packets.

A recent alternative to PoW is ***proof-of-stake (PoS)*** where security relies on the adversary holding a minority stake in an abstract finite resource [4]. When making a group decision, PoS weights each participant's vote using its share of a limited resource; for example, the amount of cryptocurrency held by the participant. A well-known example is ALGORAND [70], which employs PoS to form a committee. A hybrid approach using both PoW and PoS has been proposed in the Ethereum system [8]. We note that PoS can only be used in systems where the "stake" of each participant is globally known. Thus, PoS is typically used only in cryptocurrency applications.

2.2 Motivation for our Population Invariant

Our algorithms guarantee that the fraction of bad IDs is always less than any fixed constant. This is in contrast to a popular model where the fraction of bad IDs is *assumed* to be bounded below some fixed constant at all times. This latter model features prominently in many areas including Byzantine-resilient replication [38–40, 49, 97], secure peer-to-peer networks [22, 67, 79, 126, 148], reliable broadcast [29, 30, 91, 96], and Byzantine consensus [88, 89, 98]. Thus, our Population Invariant can be viewed as a tool enabling these past results to run correctly even under Sybil attack.

3 Commensurate COMputation (CCOM)

In this section, we describe and analyze our first algorithm CCOM.

3.1 GenID

To initialize our system, we make use of an algorithm created by Andrychowicz and Dziembowski [13], which we call **GenID**. This algorithm creates an initial set of IDs such that no more than an α-fraction are bad. GenID also selects a committee of logarithmic size that has a majority of good IDs. Finally, GenID has significant, but polynomial, computational cost; thus we use it only once during the lifetime of our system.

3.2 Constructing Computational Puzzles

We assume that each good ID can perform μ hash-function evaluations per round for $\mu > 0$. Additionally, we assume that μ is of some size polynomial in n_0 so that $\log \mu = \Theta(\log n_0)$. It is reasonable to assume large μ since, in practice, the number of evaluations that can be performed per second is on the order of millions to low billions [32, 33, 82].

Our technical puzzle generation process follows that of [13]. For any integer $h \geq 1$, we define a ***h-hard puzzle*** to consist of finding $C \log \mu$ solutions using a threshold of $h(1 - \delta)\mu/(C \log \mu)$, where $\delta > 0$ is a small constant and C is a sufficiently large constant depending on δ and μ.

Let X be a random variable (r.v.) giving the expected number of hash evaluations needed to compute $C \log \mu$ solutions. Note that X is a negative binomial r.v. and so, the following concentration bound holds (see, for example, Lemma 2.2 in [22]) for every $0 < \epsilon \leq 1$:

$$Pr(|X - E(X)| \geq \epsilon E(X)) \leq 2e^{-\epsilon^2 (C \log \mu)/(2(1+\epsilon))}$$

Given the above, one can show that every good ID will solve a h-hard puzzle with at most $h\mu$ hash function evaluations, and that the adversary must compute at least $(1 - 2\delta)h\mu$ hash evaluations to solve every h-hard puzzle. This follows from a union bound over $O(n_0^\gamma)$ join and departure events by IDs, where γ is any fixed positive constant, and for C being a sufficiently large constant depending on δ, μ and γ. Note that for small δ, the difference in computational cost is negligible, and that μ is also unnecessary in comparing costs. Thus, for ease of exposition, we assume that each ***h***-hard puzzle requires computational cost ***h*** to solve.

3.3 How Puzzles Are Used

Although, each puzzle is constructed in the same manner, they are used in two distinct ways by our algorithm. First, when a new ID wishes to join the system, it must provide a solution for a 1-round puzzle; this is referred to as ***entrance puzzle***. Here, the input to the puzzle is $K_v||s||\mathcal{T}$, where $\mathcal{T}$ is the timestamp of when the puzzle solution was generated. In order to be verified, the value $\mathcal{T}$ in the solution to an entrance puzzle must be within some small margin of the current time. In practice, this margin would primarily depend on network latency.

We note that, in the case of a bad ID, this solution may have been precomputed by the adversary by using a future timestamp. This is not a problem since the purpose of this puzzle is only to force the adversary to incur a computational cost at some point, and to deter the adversary from reusing puzzle solutions. Importantly, the entrance puzzle is not used to preserve guarantees on the fraction of good IDs in the system.

The second way in which puzzles are used is to limit the fraction of bad IDs in the system; this is referred to as a ***purge puzzle***. An announcement is periodically made that *all* IDs *already in the system* should solve a 1-round puzzle. When this occurs, a ***random string r*** of $\Theta(\log n_0^\gamma)$ bits is generated and included as part of the announcement. The value r must also be appended to the inputs for all requested solutions in this round; that is, the input is $K_v||s||r$. These random bits ensure that the adversary cannot engage in a pre-computation attack — where it solves puzzles and stores the solutions far in advance of launching an

attack — by keeping the puzzles unpredictable. For ease of exposition, we omit further discussion of this issue and consider the use of these random bits as implicit whenever a purge puzzle is issued.

While the same r is used in the puzzle construction for all IDs, we emphasize that a *different* puzzle is assigned to each ID since the public key used in the construction is unique. Again, this is only of importance to the second way in which puzzles are used. Using the public key in the puzzle construction also prevents puzzle solutions from being stolen. That is, ID K_v cannot lay claim to a solution found by ID K_w since the solution is tied to the public key K_w.

Can a message m_v from ID K_v be spoofed? This is prevented in the following manner. ID K_v signs m_v with its private key to get signature $\texttt{sign}_v$, and then sends $(m_v||\texttt{sign}_v||K_v)$ via DIFFUSE. Any other ID can use K_v to check that the message was signed by the ID K_v and thus be assured that ID K_v is the sender.

Algorithm CCOM

Key Variables:
i : iteration number
n_i^a: number of IDs joining since beginning of iteration i.
n_i^d: number of IDs departing since beginning of iteration i.
$\mathcal{S}_i$: set of IDs at the end of iteration i.
$\mathcal{S}_{\text{cur}}$: current set of IDs.
r: current random seed.

Initialization:
$\mathcal{S}_0 \leftarrow$ set of IDs returned by initial call to GENID. The initial committee is also selected by GENID.
$i \leftarrow 1$.

The committee maintains all variables above and makes all decisions using Byzantine Consensus:

1. Each joining ID, v, generates a public key; solves an entrance puzzle using that key and the current timestamp $\mathcal{T}$ and broadcasts the solution via DIFFUSE. Upon verifying the solution, the committee adds v to $\mathcal{S}_{\text{cur}}$.

2. If $n_i^a + n_i^d \geq |\mathcal{S}_{i-1}|/3$, then the current committee does the following:

   Perform Purge:

   (a) Generates a random string r and broadcast via DIFFUSE.

   (b) Sets $\mathcal{S}_i$ and $\mathcal{S}_{\text{cur}}$ to the set of IDs that return valid solutions.

   (c) Selects a new committee of size $\Theta(\log n_0)$ from $\mathcal{S}_i$ and sends out this information via DIFFUSE.

3. $i \leftarrow i + 1$

Figure 1: Pseudocode for COMMENSURATE COMPUTATION (CCOM).

3.4 Overview of CCОМ

We begin by describing *Commensurate Comput- ation* (CCОМ). The execution of CCОМ is conceptually broken into sets of consecutive rounds called *iterations*. Each iteration consists of Steps 1 and 2, and each iteration ends with the execution of Step 2(d).

Notation is introduced as needed, and Table **??** contains our frequently-used notation for easy reference. Throughout, we let log denote the natural logarithm.

3.4.1 Updating and Maintaining System Membership

A key component of our system is a subset of IDs called a committee for which the Committee Invariant must hold; recall Section 1.3. The committee tracks the current membership in the system denoted by the set $\mathcal{S}_{\text{cur}}$. Updates to this set are performed in the following manner. Each ID that wishes to join the system must solve an entrance puzzle. Each good ID in the committee will check the solution to the entrance puzzle and, upon verification, the joining ID is allowed into the system.

Verification of a puzzle solution requires checking that (1) all $C \log \mu$ inputs to h submitted generate an output that is at most $(C \log \mu)/((1 - \delta)\mu)$; and (2) each of these inputs contains the string r and also the individual public key of the ID. Upon verification, the good IDs in the committee solve Byzantine Consensus in order to agree on the contents of $\mathcal{S}_{\text{cur}}$. Those IDs that fail to submit valid puzzle solutions are denied entrance into the system.

Similarly, $\mathcal{S}_{\text{cur}}$ is also updated and agreed upon when a good ID informs the server that it is departing. Of course, bad IDs may not provide such a notification and, therefore, $\mathcal{S}_{\text{cur}}$ is not necessarily accurate at all times. Finally, if we wish to achieve only the Committee Goal, then the last portion of Line 1 is omitted: there is no need for the committee to verify solutions or maintain $\mathcal{S}_{\text{cur}}$.

At the beginning of the system, the committee knows the existing membership denoted by S_0; assume $|S_0| = n_0$ initially. At some point, $n_i^a + n_i^d \geq |S_0|/3$, where n_i^a and n_i^d denote the number of join events and departure events over iteration i; note that if the same ID joins (and departs) more than once, each such event increases n_i^a and n_i^d.

At this point, Step 2 executes, whereby all IDs are issued a purge puzzle and each ID must respond with a valid solution within 1 round. The issuing of these purge puzzles is performed by the committee via the diffusing of r to all IDs; this random string is created via solving Byzantine Consensus (for example, the protocol in [89]) in order to overcome malicious behavior by bad IDs that are committee members.[5]

Recall from Section 1.1 that a round is of sufficient length that the computation time to solve a 1-round puzzle dominates the round trip communication time between the client and server (we remove this assumption later in Section 4.1). Thus, the good IDs can receive solutions to purge puzzles in a round, and then update and agree on $\mathcal{S}_{\text{cur}}$ at the end of the iteration.

Figure 2 illustrates this purging process; good IDs are blue and bad IDs are red. Let x be the number of IDs after the most recent purge; at most $x/3$ of these IDs are bad, since $\alpha = 1/3$. According to Step 2 in CCОМ, the next purge occurs immediately when the number of new IDs entering would become greater than or equal to $x/3$. At the point when a purge is triggered, the number of new IDs that have been allowed to enter is less than $x/3$. Thus, the total number of bad IDs is less than $\frac{2}{3}x$, and the total number of good IDs is at least $\frac{2}{3}x$.

How is r sent such that each good ID trusts it came from the committee? Recall that a participant v that joins the system generates a public key, K_v, and the corresponding private key, k_v; the public key is used as its ID. If ID K_v is a committee member, it will sign the random string r using its private key k_v to obtain sign_v. Then, $(r||\text{sign}_v||K_v)$ is sent using DIFFUSE. Any good ID can verify sign_v via K_v to

---

[5]Agreeing on each random bit of r will suffice. Alternatively, a secure multiparty protocol for generating random values can be used; for example, the result in [132].
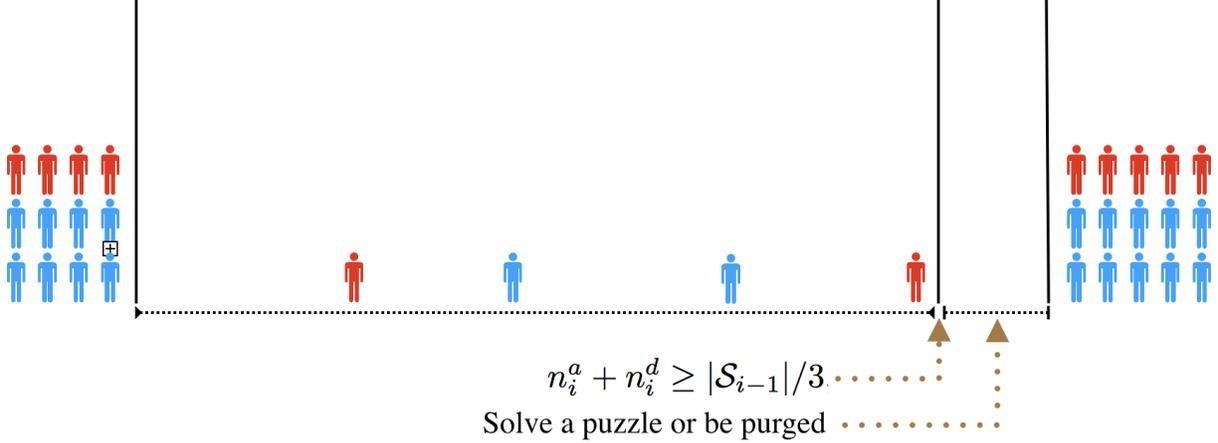
$$n_i^a + n_i^d \geq |\mathcal{S}_{i-1}|/3$$

Solve a puzzle or be purged

Figure 2: An overview of our approach. The system starts with $x = 12$ IDs of which 4 are bad and 8 are good. When the $(x/3)^{\text{th}} = 4^{\text{th}}$ ID joins, a purge is immediately executed and all IDs must solve a purge puzzle or be ejected from the system.

ensure it returns r. This implicitly occurs in Step 2 of Figure 1, but we omit the details in the pseudocode for simplicity.

3.4.2 Updating and Maintaining Committee Membership

Recall from Section 1.1 that good IDs always inform all other IDs of their departure. Thus, a good ID K_v that is a committee member will inform other committee members of its departure. Since bad IDs may not inform of their departure, our estimate of the number of IDs that join and leave will not be accurate but this will not jeopardize the committee goal.

The committee maintenance operations occur in Step 2 of CCOM. In Step 2(a), the committee generates and diffuses a random string r; recall the details discussed in Section 3.3.

In Step 2(b), each ID must respond with a valid solution within 1 round, or else be removed from the system. The committee removes unresponsive or late IDs from its whitelist, which is maintained via Byzantine consensus amongst the committee members. Over the lifetime of the system, the committee should always satisfy the Committee Invariant. As joins and departures occur, this invariant may become endangered. To address this issue, committees are disbanded and rebuilt over time. In particular, a new committee is completed and ready to be used by the start of each successive iteration, and the currently-used committee is disbanded. The construction of a new committee occurs in Step 2(c).

In Step 2(c), the current committee then selects $\Theta(\log n_0)$ IDs uniformly at random from $\mathcal{S}_i$. This selection is done via a committee election protocol, for example [88, 89]. The committee uses DIFFUSE to inform $\mathcal{S}_i$ that the selected IDs are the new committee for iteration $i + 1$. All messages from committee members are verified via public key digital signatures.

Given that the adversary has a minority of the computational power, we expect the newly-formed committee to have a good majority. Byzantine consensus allows for agreement on the members of the new committee, and this is communicated to all good IDs in the network via DIFFUSE. These two properties satisfy the other two criteria of the Committee Invariant; this is argued formally in Section 3.5.1.

Over the round during which this construction is taking place, the existing committee still satisfies the Committee Invariant; this is proved in Lemma 1. Once the new committee members are elected, they are known to all good IDs in the system. Those committee members that were present prior to this construction are no longer recognized by the good IDs. Finally, any messages that are received after this single round are considered late and discarded.

11

Used in Model and Algorithm	
Symbol	**Description**
α	Fraction of total computational power that the adversary controls.
n_0	Lower bound on number of good IDs in the system at any time.
$\mathcal{S}_{\text{cur}}$	Set of all IDs in the system at current time.
$\mathcal{S}_i$	Set of all IDs in the system at the end of iteration i. Equivalently, the set at the beginning of iteration $i + 1$.
ℓ_i	Length of iteration i.

Table 1: A summary of notation for CCOM

3.5 Analysis of CCOM

To simplify our presentation, our claims are proved to hold with probability at least $1 - \tilde{O}(1/n_0^{\gamma+2})$, where $\tilde{O}$ hides a poly($\log n_0$) factor. Of course, we wish the claims of Theorem 1 to hold with probability at least $1 - 1/n_0^{\gamma+1}$ such that a union bound over n_0^γ joins and departures yields a w.h.p. guarantee. By providing this "slack" of an $\tilde{\Omega}(1/n_0)$-factor in each of the guarantees of this section, we demonstrate this is feasible while avoiding an analysis cluttered with specific settings for the constants used in our arguments.

For any iteration i, let G_i and B_i denote the set of good IDs and bad IDs, respectively at the end of iteration i.

3.5.1 Analysis of the Committee Invariant

In this section, we prove that with high probability CCOM preserves the Committee Invariant over a polynomial number of join and leave events.

Lemma 1. *With probability $1 - O(1/n_0)$, over a polynomial number of join and departure events, there is always an honest majority in the committee.*

Proof. For iteration $i = 0$, the committee invariant holds by the correctness of GENID (recall Section 3.1).

Fix an iteration $i > 0$. Recall that a new committee is elected by the existing committee by selecting $c \log |S_i|$ IDs independently and uniformly at random from the set S_i, for a sufficiently large constant $c > 0$ which we define concretely later on in this proof. Let X_G be a random variable which denotes the number of good IDs elected to the new committee in iteration i. Then:

$$E[X_G] = \frac{|G_i|}{|S_i|} c \log |S_i| = (1 - \alpha) c \log |S_i| \tag{1}$$

where the last inequality follows from the fact that the fraction of computational power with the adversary is at most α. Next, we bound the number of good IDs in the committee using Chernoff Bounds [106]:

$$Pr\left(X_G < (1 - \delta)(1 - \alpha) c \log |S_i|\right) \leq \exp\left\{-\frac{\delta^2 (1 - \alpha) c \log |S_i|}{2}\right\} = O\left(n_0^{-(\gamma+1)}\right)$$

where the first step holds for any constant $0 < \delta < 1$, the second step follows from Equation 1, and the last step holds for all $c \geq \frac{28}{13} \frac{(\gamma+1)}{\delta^2}$. For any fixed, constant $\epsilon_0 > 0$, there exists a δ, such that the number of good IDs in the committee is at least $(1 - \alpha - \epsilon_0) c \log |S_i|$ with probability $1 - O(n_0^{-(\gamma+1)})$.

12

Next we bound departures of good IDs from the committee over an iteration. Let Y_g be a random variable which denotes the number of good IDs that depart from the committee during the iteration, i.e. up to the point when the number of departures of IDs from the system is $|S_{i-1}|/3$. Departing good IDs are assumed to be selected uniformly at random from the population of good IDs (Section 1.2.1). Thus, we obtain:

$$E[Y_g] \quad \leq \quad \frac{|S_i|}{3} \left(\frac{c \log |S_i|}{|S_i|} \right) \quad = \quad \frac{c}{3} \log |S_i| \tag{2}$$

Next, we upper bound the number of departures of good IDs from the committee using Chernoff Bounds [106]:

$$Pr\left(Y_g > (1 + \delta') \frac{c \log |S_i|}{3} \right) \quad \leq \quad \exp\left\{ -\frac{\delta'^2 c}{9} \log |S_i| \right\} = O\left(n_0^{-(\gamma+1)} \right)$$

where the first step holds for any constant $0 < \delta' < 1$ and the last step holds for all $c \geq \frac{9(\gamma+1)}{\delta'^2}$.

Thus, for any fixed $\epsilon_1 > 0$, the minimum number of good IDs in the committee at any point during the iteration is greater than $(1 - \alpha - 1/3 - \epsilon_0 - \epsilon_1) c \log |S_i|$. Good IDs can also depart while the committee is running the single-round purge at the end of iteration. But, by assumption, at most an ϵ_d-fraction of good IDs depart per round, where ϵ_d is sufficiently small. Therefore, the number of good IDs exceeds $(1 - \alpha - 1/3 - \epsilon_0 - \epsilon_1 - \epsilon_d) c \log |S_i|$. Choosing $\alpha < 1/6$ ensures that we can choose ϵ_0, ϵ_1, and ϵ_d such that the fraction of good IDs is always strictly greater than $1/2$. Finally, by a union bound over n_0^γ iterations, the committee invariant is maintained over n_0^γ iterations with probability $1 - O(1/n_0)$, which implies the claim. $\qquad\square$

3.5.2 Analysis of the Population Invariant

Lemma 2. *For all iterations $i \geq 0$, $|B_i| \leq \alpha |S_i|$.*

Proof. This is true by the properties of GENID for $i = 0$. Then after Step 2 that ends iteration i for $i > 0$, we have $|B_i| \leq \alpha |S_i|$. $\qquad\square$

Let g_i^a and b_i^a denote the good and bad IDs that join over iteration i, respectively; and let $n_i^a = g_i^a + b_i^a$. Similarly, let g_i^d and b_i^d denote the good and bad IDs that depart over iteration i, respectively; and let $n_i^d = g_i^d + b_i^d$. Note that the committee will have an accurate value for all of these variables except for possibly b_i^d and, consequently, n_i^d, since bad IDs may not give notification when they depart. Therefore, for b_i^d and n_i^d, the committee may hold values which are underestimates of the true values. However, we highlight that, in our argument below, this is not a problem for maintaining the Population Invariant.

Lemma 3. *The fraction of bad IDs is always less than $3/8$.*

Proof. Fix some iteration $i > 0$. Step 1 in iteration i ends when $n_i^a + n_i^d \geq |\mathcal{S}_{i-1}|/3$. Therefore, we have $b_i^a + g_i^d \leq |\mathcal{S}_{i-1}|/3$. We are interested in the maximum value of the ratio of bad IDs to total IDs at any point during the iteration. Thus, we pessimistically assume all additions of bad IDs and removals of good IDs come first in the iteration. We are then interested in the maximum value of the ratio:

$$\frac{|B_{i-1}| + b_i^a}{|\mathcal{S}_{i-1}| + b_i^a - g_i^d}.$$

By Lemma 2, $|B_{i-1}| \leq \alpha |\mathcal{S}_{i-1}|$. Thus, we want to find the maximum of $\frac{\alpha |\mathcal{S}_{i-1}| + b_i^a}{|\mathcal{S}_{i-1}| + b_i^a - g_i^d}$, subject to the constraint that $b_i^a + g_i^d \leq |\mathcal{S}_{i-1}|/3$. This ratio is maximized when the constraint achieves equality, that is when $g_i^d = (|\mathcal{S}_{i-1}|/3) - b_i^a$. Plugging this back into the ratio, we get

$$\frac{\alpha|\mathcal{S}_{i-1}| + b_i^a}{|\mathcal{S}_{i-1}| + b_i^a - g_i^d} \leq \frac{\alpha|\mathcal{S}_{i-1}| + b_i^a}{2|\mathcal{S}_{i-1}|/3 + 2b_i^a}$$

$$\leq \frac{\alpha|\mathcal{S}_{i-1}| + |\mathcal{S}_{i-1}|/3}{2|\mathcal{S}_{i-1}|/3 + 2|\mathcal{S}_{i-1}|/3}$$

$$< 3/8$$

The second line of the inequalities holds since $b_i^a \leq |\mathcal{S}_{i-1}|$, and the last line holds since $\alpha \leq 1/6$.

Finally, we note that this argument is valid even though $\mathcal{S}_{\mathrm{cur}}$ may not account for bad IDs that have departed *without* notifying the committee (recall this is possible as stated in Section 1.1). Intuitively, this is not a problem since such departures can only lower the fraction of bad IDs in the system; formally, the critical equation in the above argument is $b_i^a + g_i^d \leq |S_{i-1}|/3$, and this does not depend on b_i^d. □

3.5.3 Cost Analysis and Proof of Theorem 1

We now examine the cost of running CCOM. Let T_i be the total computational cost incurred by the adversary in iteration i.

Lemma 4. *For any iteration $i \geq 1$, the total computational cost to the good IDs is $O(T_i + g_i^a + g_i^d)$.*

Proof. Fix an iteration $i \geq 1$. Let $\mathcal{A}_i$ denote the total computational cost to the good IDs in this iteration. Note that Note that the total computational cost to the algorithm consists of the total purge cost paid by good IDs in the system at the end of this iteration plus the total entrance cost paid by the new good IDs that join the system in this iteration. Thus, we have:

$$\mathcal{A}_i = |G_i| + g_i^a$$

$$\leq \frac{4}{3}|S_{i-1}| + g_i^a$$

$$\leq 8(g_i^a + g_i^d + 2b_i^a) + g_i^a$$

$$= 9g_i^a + 8g_i^d + 16T_i$$

In the above, the second inequality holds since $|G_i| \leq |S_i|$ and $|S_i| \leq \frac{4}{3}|S_{i-1}|$.

The third inequality holds because $b_i^d \leq \alpha|S_{i-1}| + b_i^a$, since $|B_{i-1}| \leq \alpha|S_{i-1}|$; and also $g_i^a + g_i^d + b_i^a + b_i^d \geq |S_{i-1}|/3$. Thus we have that $g_i^a + g_i^d + 2b_i^a \geq (1/3 - \alpha)|S_{i-1}| > (1/6)|S_{i-1}|$.

Finally, the last inequality holds since $b_i^a \leq T_i$. □

Pulling the pieces together from the previous sections, we can now prove Theorem 1.

Proof of Theorem 1. We first prove that the population and committee goals hold over the lifetime of the system, followed by the analysis of cost to the algorithm.

Population and Committee goals. By Lemma 3.5.1, with probability at least $1 - O(1/n_0^{\gamma+1})$, the committee has a majority of good IDs, has size $\Theta(\log n_0)$, and its members are communicated to all good IDs via DIFFUSE; therefore, the Committee Invariant is met. Given that the Committee Invariant holds, by Lemma 3 the fraction of bad IDs is less than $1/2$ over the $O(n_0^\gamma)$ join and leave events; therefore, the Population Invariant is met.

Spending Rate. Let x be the number of iterations over the lifetime of the system. Then, from Lemma 4, we can obtain the total computational cost to the algorithm to be at most:

$$\sum_{i=1}^{x} O\left(T_i + g_i^a + g_i^d\right) \leq O\left(\sum_{i=1}^{x} T_i + \sum_{i=1}^{x} g_i^a + \sum_{i=1}^{x} g_i^d\right)$$

Note that the number of departing IDs from a system is at most equal to the number of IDs that join the system over the lifetime of the system. Using this, we obtain the total computational cost to the algorithm as:

$$O\left(\sum_{i=1}^{x} T_i + \sum_{i=1}^{x} g_i^a\right)$$

Let ℓ_i denote the length of the i^{th} iteration. Then, we can obtain the average cost to the algorithm over the lifetime of the system as:

$$\frac{O\left(\sum_{i=1}^{x} T_i + \sum_{i=1}^{x} g_i^a\right)}{\sum_{i=1}^{x} \ell_i}$$
$$\leq O\left(\frac{\sum_{i=1}^{x} T_i}{\sum_{i=1}^{x} \ell_i} + \frac{\sum_{i=1}^{x} g_i^a}{\sum_{i=1}^{x} \ell_i}\right)$$
$$= O\left(T + J^G\right)$$

The last step follows since, over all iterations over the lifetime of the system, we have $T = \left(\sum_{i=1}^{x} T_i\right) / \left(\sum_{i=1}^{x} \ell_i\right)$ and $J^G = \left(\sum_{i=1}^{x} g_i^a\right) / \left(\sum_{i=1}^{x} \ell_i\right)$. $\qquad\square$

4 Generalizations of CCOM

In this section, we present two enhancement to CCOM. The first allows for our results to hold when there is a known bounded communication latency. The second is a general rule for when to execute Step 2 in order to maintain a desired strict upper bound (not necessarily $1/2$) on the fraction of bad IDs in the system; that is, this a more flexible Population Invariant.

4.1 Handling Bounded Communication Latency

Previously, all messages were assumed to be delivered with negligible delay when compared with the time to solve computational puzzles. In this section, we relax this assumption to the following: *the adversary may choose to arbitrarily delay messages, subject to the constraint that they are delivered within a bounded number of rounds $\Delta > 0$.* Such a model of communication delay is often referred to as ***partial synchrony*** or ***Δ-bounded delay*** in the literature [61, 122]. We address this issue in the following analysis for each of our algorithms.

In DISTRIBUTED CCOM. In CCOM, messages are exchanged with good IDs every time the committee is constructed. Thus, we increase the wait time for receiving the solutions to the purge puzzle by an additive term of 2Δ. To simplify the analysis of a Δ-bounded delay in the network where the adversary holds an α-fraction of the total computational power, we can instead consider a 0-bounded delay network, but where the adversary holds an increased fraction of the total computational power; this is addressed by Lemma 5.

Lemma 5. *A Δ-bounded network latency with an adversary that holds an α-fraction of total computational power of the network is equivalent to a 0-bounded network latency with an adversary that holds a $\frac{2\alpha\Delta+\alpha}{2\alpha\Delta+1}$-fraction of the total computational power of the network.*

Proof. Recall from Section 1.1, the fraction of computational power with an adversary is α-times the total computational power of the network. Given the Δ-bounded delay in the network, the adversary has $2\Delta + 1$ rounds to perform evaluations. Fix a time step during an iteration. Suppose P is the total computational power of the network at that time step. Then, the number of evaluations of the hash function an adversary can make is $P\alpha(2\Delta + 1)\mu$.

Note that the computational power of the good IDs remains the same, which is $P(1-\alpha)\mu$ since they obey the protocol, and we assume that the bounded delay only benefits the adversary. Thus, the Δ-bounded delay model is equivalent to the 0-bounded delay model with an adversary that holds a $\frac{P\alpha(2\Delta+1)\mu}{P\alpha(2\Delta+1)\mu+P(1-\alpha)\mu} = \frac{2\alpha\Delta+\alpha}{2\alpha\Delta+1}$-fraction of the computational power. $\qquad\square$

Lemma 6. *Assume a Δ-bounded delay and an adversary that holds an α - fraction of computational power, for $\alpha \leq 1/(10\Delta + 6)$. With probability at least $1 - O(1/n_0^{\gamma+1})$, the fraction of good IDs in the committee exceeds $1/2$ for the duration of an iteration.*

Proof. From Lemma 5, the Δ-bounded delay is equivalent to a 0-bounded delay with an adversary that holds a $\frac{2\alpha\Delta+\alpha}{2\alpha\Delta+1}$-fraction of computational power of the network. By Lemma 1, for a 0-bounded delay model with an adversary that holds an α'-fraction of the computational power for $\alpha' \leq 1/6$, the fraction of good IDs exceeds $1/2$. Substituting $\alpha' = \frac{2\alpha\Delta+\alpha}{2\alpha\Delta+1}$, the Δ-bounded delay model with an α-fraction of computational power has at least $1/2$-fraction of good IDs, for $\alpha \leq 1/(10\Delta + 6)$. $\qquad\square$

4.2   Reducing the Fraction of Bad IDs in the Population Invariant

We show how to modify our algorithm in order to maintain the fraction of bad IDs in the system to be less than 3α at all times for any given $\alpha < 1/6$ (i.e., fraction of computational power with the adversary is some value strictly less than $1/6$ of the total computation power in the system). We also show the analogous result for the Committee Invariant.

In order to maintain the fraction of bad IDs to be always some fraction less than 3α, we propose the following modification to Step 2 of Figure 1. Instead of $n_i^a + n_i^d \geq |S_{i-1}|/3$, we have the condition $n_i^a + n_i^d \geq \frac{2\alpha}{1+3\alpha}|S_{i-1}|$.

We prove the more general Population Invariant in the following lemma:

Lemma 7. *For any $\alpha < 1/6$, the fraction of bad IDs in the system is always less than 3α.*

Proof. Fix some iteration $i \geq 1$. Then, similar to Lemma 3, we can calculate the maximum fraction of bad IDs in iteration i as:
$$\frac{|B_{i-1}| + b_i^a}{|S_{i-1}| + b_i^a - g_i^d}$$

Note that $|B_{i-1}| < \alpha|S_{i-1}|$ at the end of iteration $i-1$. Also, in order to handle the variant Population Invariant, we modify our condition to end an iteration to $n_i^a + n_i^d \geq \frac{2\alpha}{1+3\alpha}|S_{i-1}|$ for a given $\beta < 3\alpha$. Thus, the above

fraction is maximized when $b_i^a + g_i^d = \frac{2\alpha}{1+3\alpha}|S_{\text{i-1}}|$. It follows that:

$$
\begin{aligned}
\frac{|B_{\text{i-1}}| + b_i^a}{|S_{\text{i-1}}| + b_i^a - g_i^d} \quad &< \quad \frac{\alpha|S_{\text{i-1}}| + b_i^a}{|S_{\text{i-1}}| + b_i^a - g_i^d} \\[2mm]
&= \quad \frac{\alpha|S_{\text{i-1}}| + b_i^a}{\left(\frac{2\alpha}{1+3\alpha} + \frac{1+\alpha}{1+3\alpha}\right)|S_{\text{i-1}}| + b_i^a - g_i^d} \\[2mm]
&= \quad \frac{\alpha|S_{\text{i-1}}| + b_i^a}{\frac{1+\alpha}{1+3\alpha}|S_{\text{i-1}}| + 2b_i^a} \\[2mm]
&\leq \quad \frac{\alpha|S_{\text{i-1}}| + b_i^a}{\frac{1+\alpha}{1+3\alpha}|S_{\text{i-1}}|} \\[2mm]
&= \quad \frac{1+3\alpha}{1+\alpha}\left(\alpha + \frac{b_i^a}{|S_{\text{i-1}}|}\right) \\[2mm]
&< \quad 3\alpha
\end{aligned}
$$

Where the final inequality holds since $b_i^a < \frac{2\alpha}{1+3\alpha}|S_{\text{i-1}}|$. $\qquad\square$

5 Geometric Mean COMputation (GMCom)

We now describe our algorithm, GMCom. Recall that GMCom makes additional assumptions about the join and departures of good IDs, as detailed in Section 1.1.

In this section, we first discuss a key technical challenge of GMCom: estimating the good join rate (Section 5.1). Next, we give an overview of GMCom (Section 3.4), see Figure 3 for pseudocode. Then we discuss intuition for our new entrance costs (Section 5.3). We collect notation specific to this section in Table 2 for easy reference.

5.1 J_i^G and $\tilde{J}_i^G$

In this section, we describe how we estimate J_i^G, which is intuitively the good join rate in iteration i. More specifically, consider an iteration i that overlaps epochs with indices defined by the set $\mathcal{E}$. For $j \in \mathcal{E}$, let λ_j be the fraction of iteration i that overlaps with epoch j; note that $\sum_{j \in \mathcal{E}} \lambda_j = 1$. Then we define $J_i^G = \sum_{j \in \mathcal{E}} \lambda_j \rho_i$; that is, the fraction of the overlap multiplied by the good join rate in the respective epoch. We refer to this quantity as the ***average good join rate over iteration*** i or simply "the good join rate over iteration i" (omitting "average").

GMCom uses a good join rate estimate, $\tilde{J}_i^G$, in Line 1 of Figure 3 to set the entrance puzzle difficulty. Computing $\tilde{J}_i^G$ is technically challenging, and is done as follows. The committee continually keeps track of sets $\tilde{S}$, which is the most recent system membership, where $|\mathcal{S}_{\text{cur}} - \tilde{S}| \geq \frac{3}{5}|\mathcal{S}_{\text{cur}}|$ (Line (i) in Figure 3) and $\mathcal{S}_{\textbf{\textit{cur}}}$ is the set of good IDs currently in the system. Whenever $\tilde{S}$ is updated — and only when $\tilde{S}$ is updated — the parameter $\tilde{\ell}$ is set to the length of time since the system membership was most recently $\tilde{S}$ (Line (ii) in GMCom), and we refer to this length of time as an ***interval***. Then, $\tilde{J}_i^G = |\mathcal{S}_{\text{cur}}|/\tilde{\ell}$ is used as an estimate for the good join rate over iteration i. The setting of this variable is performed in Line 2(d) of GMCom.

We demonstrate in Section 6.2 that $\tilde{J}_i^G$ is always within a constant-factor of J_i^G. Our approach for estimating the join rate of good IDs may be of independent interest for settings where Assumptions A1 and A2 hold, and it is worthwhile considering why these assumptions are necessary, as discussed below.

Since good and bad IDs cannot be distinguished with certainty, GMCom cannot detect when an epoch begins and ends. Instead, GMCom finds a lower bound on the number of new good IDs that have joined

17

Algorithm GMCOM

Key Variables
i: iteration number
n_i^a: number of IDs joining since beginning of iteration i
n_i^d: number of IDs departing since beginning of iteration i
$\mathcal{S}_i$: set of IDs at end of iteration i
$\mathcal{S}_{\text{cur}}$: current set of IDs in system

Initialization:
$i \leftarrow 1$.
$\mathcal{S}_0 \leftarrow$ set of IDs returned by initial call to GENID. The initial committee is also
    selected by GENID
$\tilde{J}_0^G \leftarrow$ obtained in initialization phase
$\tilde{S} \leftarrow \mathcal{S}_0$

Execution:

The committee maintains all variables above and makes all decisions using Byzantine Consensus. The committee also maintains:

  (i) $\tilde{S} \leftarrow$ most recent system membership ensuring $|\mathcal{S}_{\text{cur}} - \tilde{S}| \geq \frac{3}{5}|\mathcal{S}_{\text{cur}}|$.

  (ii) $\tilde{\ell} \leftarrow$ length of time between last two changes of variable $\tilde{S}$.

For each iteration i, do:

1. Each joining ID solves and diffuses the solution to an entrance puzzle of difficulty equal to the number of IDs that have joined in the last $1/\tilde{J}_i^G$ seconds of the current iteration, including the newly joining ID.

2. When $n_i^a + n_i^d \geq (1/11)|\mathcal{S}_{i-1}|$, do: **Perform Purge:**

  (a) The committee generates and diffuses a random string r to be used in puzzles for this purge and entrance for the next iteration.

  (b) $\mathcal{S}_i \leftarrow$ set of IDs returning difficulty 1 puzzle solutions within 1 round.

  (c) The committee selects a new committee of size $\Theta(\log n_0)$ from $\mathcal{S}_i$ and sends out this information via DIFFUSE.

End Iteration:

  (d) $i \leftarrow i + 1$
  (e) $\tilde{J}_i^G \leftarrow |\mathcal{S}_{\text{cur}}|/\tilde{\ell}$

Figure 3: Pseudocode for GEOMETRICMEAN COMPUTATION (GMCOM).

Table 2: A summary of additional notation used for GMCOM.

Used in Model and Algorithm	
Symbol	**Description**
ℓ_i	Length of epoch i.
ρ_i	The good join rate in epoch i; that is, the number of good IDs joining in epoch i divided by ℓ_i.
$\tilde{J}_i^G$	Estimate of the good join rate in iteration i.
J_i^G	Average good join rate over iteration i; that is, $\sum_{j \in \mathcal{E}} \lambda_j \rho_i$, where $\mathcal{E}$ is the set of indices of epochs overlapped by the iteration, and λ_j is the fraction of the iteration which overlaps epoch $j \in \mathcal{E}$.
J_{cur}	Current number of join events in present iteration divided by current time elapsed in this iteration.
$\tilde{S}$	Most recent system membership where $|\mathcal{S}_{\text{cur}} - \tilde{S}| = (3/5)|\tilde{S}|$
$\tilde{\ell}$	Amount of time since system membership was most recently $\tilde{S}$.
$T_{\mathcal{I}}$	Cost to the adversary for solving puzzles whose solutions are used in any iteration of $\mathcal{I}$ divided by total length of those iterations.
$J_{\mathcal{I}}^G$	Number of good IDs that join over the iterations in $\mathcal{I}$ divided by total length of those iterations.

the system by pessimistically subtracting out the fraction of IDs that could be bad. When this estimate of new good IDs is sufficiently large — that is, when the current interval ends — GMCOM "guesses" that at least one epoch has occurred. However, multiple (but still a constant number of) epochs may have actually occurred during this time. Thus, for this approach to yield a constant-factor estimate of J_i^G, one requirement is that the true good join rate cannot have changed "too much" between these epochs, and Assumption A1 bounds the amount of such change.

Note that Assumption A1 addresses the good join rate measured over entire, consecutive epochs. However, an interval may end *within* some epoch j. If the good join rate over the portion of epoch j overlapped by the interval deviates by more than a constant factor from ρ_j, then the estimate will be inaccurate. Assumption A2 ensures that this cannot happen.

5.2 Overview of GMCOM

We now describe the execution of GMCOM in iteration i while referring to our pseudocode in Figure 3.

As stated earlier in Section 5.1, estimating the good-ID join rate involves tracking $\tilde{S}$ and $\tilde{\ell}$ in Lines (i) and (ii), respectively. The remainder of the iteration executes in two steps.

In Step 1, each joining ID must solve a puzzle of difficulty 1 plus the number of IDs that join within the last $1/\tilde{J}_i^G$ seconds, where $\tilde{J}_i^G$ is our estimate of the join rate for good IDs. We provide intuition for this cost in Section 5.3.

Step 1 lasts until the earliest point in time when the number of IDs that join in iteration i, $\boldsymbol{n_i^a}$, plus the number of IDs that depart in iteration i, $\boldsymbol{n_i^d}$, is at least $(1/11)|\mathcal{S}_{i-1}|$. The quantities n_i^a and n_i^d are tracked by the committee.

When Step 1 ends, a purge is performed by the committee by issuing a 1-hard puzzle via DIFFUSE via Step 2(a). In Step 2(b), each ID must respond with a valid solution within 1 round, or else be removed from the system. The committee removes unresponsive or late IDs from its whitelist, which is maintained via

Byzantine consensus amongst the committee members.

The current committee then selects $\Theta(\log n_0)$ IDs uniformly at random from $\mathcal{S}_i$ in Step 2(c). The committee uses DIFFUSE to inform $\mathcal{S}_i$ that the selected IDs are the new committee for iteration $i + 1$. All messages from committee members are verified via public key digital signatures.

In Step 2 the estimate of the good-ID join rate over this iteration is calculated and iteration i ends.

5.3 Intuition for Entrance-Cost Function

We now give intuition for our entrance-cost function. Consider iteration i. In the absence of an attack, the entrance cost should be proportional to the good join rate. This is indeed the case since the puzzle difficulty is $O(1)$ corresponding to the number of (good) IDs that join within the last $1/\tilde{J}_i^G \approx 1/J_i^G$ seconds.

In contrast, if there is a large attack, then our entrance-cost function imposes a significant cost on the adversary. Consider the case where a batch of many bad IDs is rapidly injected into the system. This drives up the entrance cost since the number of IDs joining within $1/\tilde{J}_i^G$ seconds increases.

More precisely, assume the adversary's spending rate is $T = \xi J_i^{\text{all}}$, where $\boldsymbol{\xi}$ is the entrance cost, and $\boldsymbol{J_i^{\text{all}}}$ is the join rate for all IDs. For the good IDs, the spending rate due to the entrance cost is ξJ_i^G, and the spending rate due to the purge cost is J_i^{all}. Setting these to be equal, and solving for ξ, we get $\xi = J_i^{\text{all}}/J_i^G$; in other words, the number of IDs that have joined over the last $1/J_i^G$ seconds. This is the entrance cost function that best balances entrance and purge costs.

Note that spending rate of good IDs due to the entrance costs and purge costs is:

$$
\begin{aligned}
\xi J_i^G + J_i^{\text{all}} &\leq 2 J_i^{\text{all}} && \text{By our setting of } \xi. \\
&= 2\sqrt{(J_i^{\text{all}})^2} \\
&= 2\sqrt{J_i^{\text{all}} \xi J_i^G} && \text{Since } J_i^{\text{all}} = \xi J_i^G. \\
&= 2\sqrt{J_i^G T} && \text{Since } T = \xi J_i^{\text{all}}.
\end{aligned}
$$

This informal analysis provides intuition for the asymmetric cost guaranteed by GMCOM. Proving that $\tilde{J}_i^G$ is a good estimate for J_i^G, showing that the above analysis holds even only with good estimates of join rates, and incorporating multiple epochs and purge costs, is the subject of our analysis in the next section.

6 Analysis of GMCOM

In this section we prove correctness and analyze computational costs for GMCOM.

6.1 Maintaining the Population and Committee Invariants

GMCOM ensures a modified Population Invariant which guarantees that at most a $1/6$-fraction of IDs are bad.

Lemma 8. *For $\alpha \leq 1/18$,* GMCOM *guarantees that the fraction of bad IDs in the system is less than $1/6$.*

Proof. By Lemma 7, if $\alpha = 1/18$ and a threshold for triggering a purge is $\frac{2\alpha}{1+3\alpha}|\mathcal{S}_{i-1}| = (2/18)/(7/6)|\mathcal{S}_{i-1}| > (1/11)|\mathcal{S}_{i-1}|$, then the fraction of bad IDs in the system is always less than $3\alpha = 1/6$. $\square$

The threshold for the number of join and departure events for triggering a purge also changes for GMCOM, and so we prove that the Committee Invariant still holds.

Lemma 9. *With probability $1 - O(1/n_0)$, over a polynomial number of join and departure events, the fraction of good IDs in the committee is always at least $10/11 - \alpha - 2\epsilon_d$.*

Proof. The argument is almost identical to the proof of Lemma 1, but we provide the details for completeness. For iteration $i = 0$, the committee invariant holds by the use of [84] to initialize the system (recall Section 3.4; for details, see Lemma 6 of [13]).

Fix an iteration $i > 0$. Recall that a new committee is elected by the existing committee, by selecting $c \log |S_i|$ IDs independently and uniformly at random from the set S_i of IDs left after the purge, where $c > 0$ is a sufficiently large constant set later in this proof. Let X_G be a random variable giving the number of good IDs in the committee at formation. Then:

$$E[X_G] \quad = \quad \frac{|G_i|}{|S_i|} c \log |S_i| \quad = \quad (1 - \alpha) c \log |S_i| \tag{3}$$

where the last inequality follows from the fact that the fraction of computational power with the adversary is at most α. By Chernoff bounds [106], for any δ, $0 < \delta < 1$:

$$Pr\left(X_G < (1 - \delta)(1 - \alpha)c \log |S_i|\right) \leq \exp\left\{-\frac{\delta^2(1 - \alpha)c \log |S_i|}{2}\right\}$$

For any constant $\epsilon_0 > 0$, setting $\delta = \frac{\epsilon_0}{1 - \alpha}$, we have that:

$$Pr(X_g < (1 - \alpha - \epsilon_0)c \log |S_i|) = O(n_0^{-(\gamma+1)}).$$

Where the last equality holds for $c \geq \frac{28}{13} \frac{(\gamma+1)}{\delta^2}$.

Next, let Y_g be a random variable denoting the number of good IDs that depart from the committee during iteration i, but before the purge, i.e. up to the point when the number of departures of IDs from the system is $|S_{i-1}|/11$. Departing good IDs are selected uniformly at random from the population of good IDs (See Section 1.2.1). Thus, we obtain:

$$E[Y_g] \quad \leq \quad \frac{|S_i|}{11}\left(\frac{c \log |S_i|}{|S_i|}\right) \quad = \quad \frac{c}{11} \log |S_i| \tag{4}$$

By Chernoff bounds [106]:

$$Pr\left(Y_g > (1 + \delta')\frac{c \log |S_i|}{11}\right) \quad \leq \quad \exp\left\{-\frac{\delta'^2 c}{33} \log |S_i|\right\} = O\left(n_0^{-(\gamma+1)}\right)$$

where the first inequality holds for any δ, $0 < \delta' < 1$ and the equality holds for all $c \geq \frac{33(\gamma+1)}{\delta'^2}$.

Thus, for any fixed $\epsilon_1 > 0$, the minimum number of good IDs in the committee at any point during the iteration is greater than $(1 - \alpha - 1/11 - \epsilon_0 - \epsilon_1)c \log |S_i|$. Good IDs can also depart while the committee is running the single-round purge at the end of iteration. But, by assumption, at most an ϵ_d-fraction of good IDs depart per round, where ϵ_d is sufficiently small. Setting $\epsilon_0 = \epsilon_1 = (1/2)\epsilon_d$, we get that the number of good IDs exceeds $(10/11 - \alpha - 2\epsilon_d)c \log |S_i|$, with probability of error $O\left(n_0^{-(\gamma+1)}\right)$.

Finally, the claim follows by taking a union bound over all n_0^γ iterations. $\square$

6.2 Estimating the Good ID Join Rate

To calculate the entrance cost in iteration i, GMCOM requires an estimate of the good ID join rate, J_i^G. However, good IDs cannot be discerned from bad IDs upon entering the system, and so the adversary may inject bad IDs in an attempt to obscure the true join rate of good IDs.

In this section, we prove that our estimate of the good-ID join rate, $\tilde{J}_i^G$, is within a constant factor of J_i^G. For any time t, let $\mathcal{S}_t$ and $\mathcal{G}_t$ denote the set of all IDs and set of good IDs, respectively, in the system at time t.

We say that interval $\tilde{\ell}$ **touches** an epoch if there is a point in time belonging to both the interval $\tilde{\ell}$ and the interval of time corresponding to the epoch; it does not necessarily mean that ℓ completely contains the epoch, or vice versa.

Lemma 10. *An iteration touches at most 2 epochs.*

Proof. Assume that some iteration i starts at time t_0 and touches at least three epochs; we will derive a contradiction. This assumption implies that there is at least one epoch entirely contained within iteration i. Let this epoch start at time $t_1 \geq t_0$ and end at time $t_2 > t_1$ within iteration i.

By definition, the number of good IDs that join over this epoch is at least $(3/4)|\mathcal{G}_{t_2}|$. Observe that:

$$
\begin{aligned}
(3/4)|\mathcal{G}_{t_2}| &\geq (3/4)(5/6)|\mathcal{S}_{t_2}| \quad \text{by the Population Invariant (Lemma 8)} \\
&= (5/8)|\mathcal{S}_{t_2}| \\
&> (5/8)((10/11)|\mathcal{S}_{t_0}|) \\
&= (25/44)|\mathcal{S}_{t_0}|
\end{aligned}
$$

where the second to last line follows from noting that $|\mathcal{S}_{t_2}| > |\mathcal{S}_{t_0}| - (1/11)|\mathcal{S}_{t_0}| \geq (10/11)|\mathcal{S}_{t_0}|$, since at most $(1/11)|\mathcal{S}_{t_0}|$ departures can happen before a purge occurs.

Finally, observe two bounds. First, during any iteration at most $(1/11)|\mathcal{G}_{t_0}|$ good IDs join before the purge. Second, at most $\epsilon_a(1 + 1/11)|\mathcal{G}_{t_0}| = \epsilon_a(12/11)|\mathcal{G}_{t_0}|$ join during the single round of the purge. Therefore, the total number of good IDs that can join over an iteration is $(1/11)|\mathcal{G}_{t_0}| + \epsilon_a(12/11)|\mathcal{G}_{t_0}|$. The contradiction follows from noting that:

$$
(25/44)|\mathcal{S}_{t_0}| > (1/11)|\mathcal{G}_{t_0}| + \epsilon_a(12/11)|\mathcal{G}_{t_0}|.
$$

for $\epsilon_a < 7/16$. $\qquad\square$

Lemma 11. *An interval touches at most two epochs and cannot completely overlap any single epoch.*

Proof. Assume that some interval starts at time t_0 and touches at least three epochs; we will derive a contradiction. This assumption implies that there is at least one epoch entirely contained within the interval. Consider the first such epoch, and let it start at time $t_1 \geq t_0$ and end at time $t_2 > t_1$.

Observe that:

$$
\begin{aligned}
|\mathcal{S}_{t_2} - \mathcal{S}_{t_1}| &\geq |\mathcal{G}_{t_2} - \mathcal{G}_{t_1}| \\
&\geq (3/4)|\mathcal{G}_{t_2}| \quad \text{by the definition of an epoch} \\
&\geq (3/4)(5/6)|\mathcal{S}_{t_2}| \quad \text{by the Population Invariant (Lemma 8)} \\
&= (5/8)|\mathcal{S}_{t_2}| \\
&> (3/5)|\mathcal{S}_{t_2}|
\end{aligned}
$$

But this is a contradiction since it implies that the interval must end before time t_2. $\qquad\square$

Lemma 12. *Consider any interval of length $\tilde{\ell}$, where there are $|\mathcal{S}_{cur}|$ IDs in the system at the end of the interval, and let epoch $i \geq 1$ be the most recent epoch that the interval touches. Then:*

$$
\left(\frac{5}{6}\right) c(\texttt{A1}, \texttt{low}) c(\texttt{A2}, \texttt{low}) \rho_i \leq \frac{|\mathcal{S}_{cur}|}{\tilde{\ell}} \leq 5 c(\texttt{A1}, \texttt{high}) c(\texttt{A2}, \texttt{high}) \rho_i.
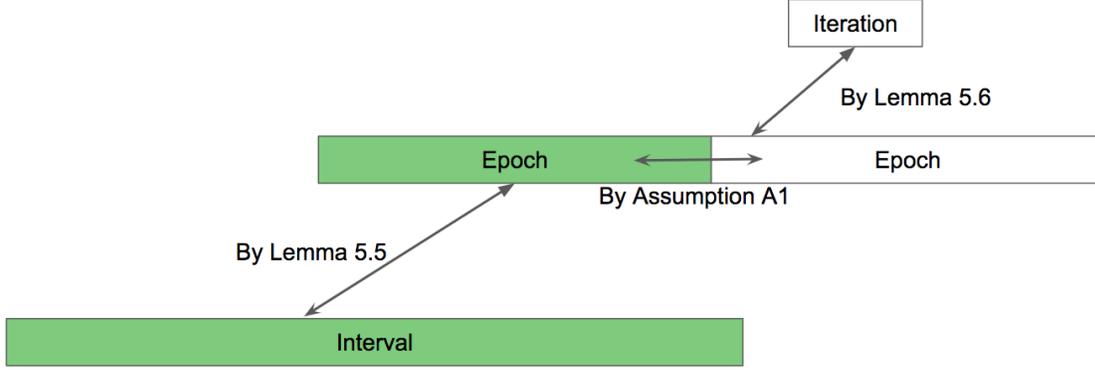$$

Figure 4: A depiction of the relationship between the good join rate in iterations, epochs, and durations. Arrows represent that good join rates of the endpoints are within constant factors of each other. Green indicates the most recently-finished epoch and duration being tracked by the committee to obtain $\tilde{J}_i^G$.

Proof. We break the argument into three cases.

Case 1: The interval touches only a single epoch where the epoch begins at t_0, the interval begins at t_1, and the interval ends at t_2. By Assumption A2, we have:

$$c(\text{A2, low})\rho_i \leq \frac{|\mathcal{G}_{t_2} - \mathcal{G}_{t_1}|}{t_2 - t_1} \leq c(\text{A2, high})\rho_i$$

and by the Population Invariant and the specification of an interval:

$$\left(\frac{2}{5}\right)|\mathcal{S}_{t_2}| \leq \left(\frac{3}{5} - \frac{1}{6}\right)|\mathcal{S}_{t_2}| \leq |\mathcal{G}_{t_2} - \mathcal{G}_{t_1}| \leq \left(\frac{3}{5}\right)|\mathcal{S}_{t_2}|.$$

Therefore, we have:

$$\left(\frac{5}{3}\right)c(\text{A2, low})\rho_i \leq \frac{|\mathcal{S}_{t_2}|}{t_2 - t_1} \leq \left(\frac{5}{2}\right)c(\text{A2, high})\rho_i$$

Case 2: The interval touches epochs $i - 1$ and i and there are at least 2 good ID join events in each epoch. Let epoch $i - 1$ start at time t_0 and end at t_2 (and so epoch i starts at t_2), and let the interval start at time $t_1 \geq t_0$ and end at $t_3 \geq t_2$. By Assumptions A1 and A2, we have:

$$c(\text{A1, low})c(\text{A2, low})\rho_i \leq \frac{|\mathcal{G}_{t_3} - \mathcal{G}_{t_1}|}{t_3 - t_1} \leq c(\text{A1, high})c(\text{A2, high})\rho_i$$

and by the Population Invariant and the specification of an interval:

$$\left(\frac{2}{5}\right)|\mathcal{S}_{t_3}| \leq \left(\frac{3}{5} - \frac{1}{6}\right)|\mathcal{S}_{t_3}| \leq |\mathcal{G}_{t_3} - \mathcal{G}_{t_1}| \leq \left(\frac{3}{5}\right)|\mathcal{S}_{t_3}|.$$

Therefore, we have:

$$\left(\frac{5}{3}\right)c(\text{A1, low})c(\text{A2, low})\rho_i \leq \frac{|\mathcal{S}_{t_3}|}{t_3 - t_1} \leq \left(\frac{5}{2}\right)c(\text{A1, high})c(\text{A2, high})\rho_i \tag{5}$$

Case 3: The interval touches epochs $i - 1$ and i, and w.l.o.g. we have a single good join event in the portion of epoch $i - 1$ that overlaps the interval; denote the length of this overlap by $\lambda' > 0$. As with Case 2, let

epoch $i - 1$ start at time t_0 and end at t_2 (and so epoch i starts at t_2), and let the interval start at time $t_1 \geq t_0$ and end at $t_3 \geq t_2$.

Observe that the single good join event in epoch $i - 1$ increases the numerator of the bounded quantity in Equation 5 by 1, and so twice the upper bound in Case 2 suffices here. The denominator increases by λ', where $\lambda' \leq t_3 - t_1$, so half the lower bound in Case 2 suffices. This implies:

$$\left(\frac{5}{6}\right) c(\texttt{A1}, \texttt{low}) c(\texttt{A2}, \texttt{low}) \rho_i \leq \frac{|\mathcal{S}_{t_3}|}{t_3 - t_1} \leq 5 c(\texttt{A1}, \texttt{high}) c(\texttt{A2}, \texttt{high}) \rho_i$$

$\square$

Lemma 13. *For any epoch $i \geq 1$, and any iteration $i' \geq 1$ that epoch i touches:*

$$c(\texttt{A1}, \texttt{low}) J_{i'}^G \leq \rho_i \leq c(\texttt{A1}, \texttt{high}) J_{i'}^G$$

Proof. By Lemma 10, iteration i' touches at most two epochs, say epochs $i - 1$ and i. Thus, by definition of $J_{i'}^G$, for some $\lambda_1, \lambda_2 \geq 0$, $\lambda_1 + \lambda_2 = 1$:

$$
\begin{aligned}
J_{i'}^G &= \lambda_1 \rho_{i-1} + \lambda_2 \rho_i \\
&\leq \lambda_1 \left(\frac{\rho_i}{c(\texttt{A1}, \texttt{low})} \right) + \lambda_2 \rho_i \\
&\leq \frac{\rho_i}{c(\texttt{A1}, \texttt{low})}
\end{aligned}
$$

The second line in the above holds by Assumption A1. The last line holds since $\lambda_1 + \lambda_2 = 1$ and $c(\texttt{A1}, \texttt{low}) \leq 1$. A similar derivation yields that $J_{i'}^G \geq \frac{\rho_i}{c(\texttt{A1}, \texttt{high})}$. Together this implies that:

$$c(\texttt{A1}, \texttt{low}) J_{i'}^G \leq \rho_i \leq c(\texttt{A1}, \texttt{high}) J_{i'}^G.$$

$\square$

Theorem 4. *For any iteration $i \geq 1$, the following bound holds:*

$$\left(\frac{5}{6}\right) \frac{c(\texttt{A1}, \texttt{low})^2 c(\texttt{A2}, \texttt{low})}{c(\texttt{A1}, \texttt{high})} J_i^G \leq \tilde{J}_i^G \leq \frac{5 \, c(\texttt{A1}, \texttt{high})^2 c(\texttt{A2}, \texttt{high})}{c(\texttt{A1}, \texttt{low})} J_i^G$$

Proof. The estimate $\tilde{J}_i^G$ used in iteration i corresponds to the most recent interval that completed before iteration i started. Let epoch j be the most recent epoch that touches this interval. By Lemma 12:

$$\left(\frac{5}{6}\right) c(\texttt{A1}, \texttt{low}) c(\texttt{A2}, \texttt{low}) \rho_j \leq \tilde{J}_i^G \leq 5 \, c(\texttt{A1}, \texttt{high}) c(\texttt{A2}, \texttt{high}) \rho_j.$$

Epoch j may end prior to the start of iteration i; that is, epoch j may not necessarily touch iteration i. In this case, note that by Lemma 11, the current interval touches epoch $j + 1$ and must end before epoch $j + 1$ ends. This fact, along with the observation that the current interval touches iteration i, implies that epoch $j + 1$ touches iteration i.

By the above, we know that either epoch $j + 1$ or epoch j touches iteration i.. Lemma 13 implies:

$$c(\texttt{A1}, \texttt{low}) J_i^G \leq \rho_j \leq c(\texttt{A1}, \texttt{high}) J_i^G$$

24

or

$$c(\text{A1}, \texttt{low})J_i^G \leq \rho_{j+1} \leq c(\text{A1}, \texttt{high})J_i^G.$$

To employ our top-most equation, we use A1 to derive:

$$\rho_{j+1}/c(\text{A1}, \texttt{high}) \leq \rho_j \leq \rho_{j+1}/c(\text{A1}, \texttt{low})$$

and then plug into our top-most equation, we have:

$$\left(\frac{5}{6}\right)\frac{c(\text{A1}, \texttt{low})^2 c(\text{A2}, \texttt{low})}{c(\text{A1}, \texttt{high})}J_i^G \leq \tilde{J}_i^G \leq \frac{5\,c(\text{A1}, \texttt{high})^2 c(\text{A2}, \texttt{high})}{c(\text{A1}, \texttt{low})}J_i^G$$

$\square$

6.3 Cost Analysis

We make use of the following two algebraic facts that both follow from the Cauchy-Schwartz inequality.

Lemma 14. *Let n be a positive number, and for all $1 \leq i \leq n$, $s_i \geq 0$ and let $S = \sum_{i=1}^n s_i$. Then*

$$\sum_{i=1}^n s_i^2 \geq S^2/n$$

Proof. Let u be a vector of length n with for all $1 \leq i \leq n$, $u[i] = s_i$, and let v be a vector of length n with, for all $1 \leq i \leq n$, $v[i] = 1$. Then by Cauchy-Schwartz:

$$\begin{aligned}
|\langle u, v \rangle|^2 &\leq \langle u, u \rangle \cdot \langle v, v \rangle \\
S^2 &\leq \left(\sum_{i=1}^n s_i^2\right) \cdot n
\end{aligned}$$

Rearranging completes the proof. $\square$

Let $\mathcal{T}_i$ denote the cost rate paid by the adversary during iteration i. Let J_i^B be the join rate of bad IDs during iteration i. For simplicity, we provide notation for the following constants:

$$c(JE, \texttt{high}) = \frac{5\,c(\text{A1}, \texttt{high})^2 c(\text{A2}, \texttt{high})}{c(\text{A1}, \texttt{low})}$$

and

$$c(JE, \texttt{low}) = \left(\frac{5}{6}\right)\frac{c(\text{A1}, \texttt{low})^2 c(\text{A2}, \texttt{low})}{c(\text{A1}, \texttt{high})}$$

where JE stands for "join estimate".

Lemma 15. *For any iteration $i > 1$, $J_i^B \leq c(L8)\sqrt{T_i(J_i^G + 1)}$, where $c(L8) = \sqrt{2c(JE, \texttt{high})}$.*

Proof. For simplicity, we normalize time units so that $\ell_i = 1$. Partition iteration i from left to right into sub-iterations, all of length $1/\tilde{J}_i^G$, except the last, which is of length at most $1/\tilde{J}_i^G$. We lower bound the cost paid by the adversary for joins by pessimistically assuming that only bad IDs are counted when computing entrance costs. For $1 \leq x \leq \lceil \tilde{J}_i^G \rceil$, let j_x be the total number of bad IDs that join in sub-iteration x. Since $\sum_{y=1}^{j_x} y = (j_x + 1)j_x/2 \geq (j_x)^2/2$, the total entrance cost paid by bad IDs is at least $(1/2)\sum_{x=1}^{\lceil \tilde{J}_i^G \rceil}(j_x)^2$.

Since $\sum_{x=1}^{\lceil \tilde{J}_i^G \rceil} j_x = J_i^B$, by applying Lemma 14, we have:

25

$$T_i \geq \frac{1}{2} \sum_{x=1}^{\lceil \tilde{J}_i^G \rceil} (j_x)^2 \geq \frac{(J_i^B)^2}{2(\lceil \tilde{J}_i^G \rceil)}$$

Cross-multiplying and taking the square root, we get:

$$
\begin{aligned}
J_i^B &\leq \sqrt{2T_i \lceil \tilde{J}_i^G \rceil} \\
&\leq \sqrt{2T_i (\tilde{J}_i^G + 1)} \\
&\leq \sqrt{2T_i \left(c(JE, \texttt{high}) J_i^G + 1 \right)}
\end{aligned}
$$

where the second step follows from noting that $\lceil x \rceil \leq x + 1$ for all x, and the final step follows from Lemma 4 which states that:

$$\tilde{J}_i^G \leq \frac{5c(\text{A1}, \texttt{high})^2 c(\text{A2}, \texttt{high})}{c(\text{A1}, \texttt{low})} J_i^G = c(JE, \texttt{high}) J_i^G$$

which yields the lemma statement. $\qquad\square$

Lemma 16. *Let $\mathcal{A}_i$ be the average spend rate for the algorithm in any iteration $i > 1$. Then:*

$$\mathcal{A}_i \leq \frac{c(L9)|S_{i-1}|}{\ell_i}$$

where $c(L9) = \left(\frac{12}{11} + \frac{c(\text{A1, high})c(\text{A2, high})}{11c(JE, low)} \right)$.

Proof. For simplicity, we first normalize time units so that $\ell_i = 1$. Partition iteration i from left to right into sub-iterations, all of length $1/J_i^G$, except the last, which is of length at most $1/J_i^G$.

The spend rate for the algorithm due to purge costs is $|S_{i-1}|$.

For entrance costs, note the following two facts. First, by Assumptions A1 and A2, there are at most $c(\text{A1}, \texttt{high})c(\text{A2}, \texttt{high})$ good IDs in any sub-iteration (note that a sub-iteration might span two epochs). Second, the entrance cost for any good ID is 1 plus the number of join events over the past $1/\tilde{J}_i^G$ seconds. By Theorem 4, $1/\tilde{J}_i^G \leq 1/(c(JE, \texttt{low})J_i^G)$, and so the entrance cost is at most 1 plus the number of join events over the past $1/c(JE, \texttt{low})$ sub-iterations; that is $1 + c(\text{A1}, \texttt{high})c(\text{A2}, \texttt{high})/c(JE, \texttt{low})$.

By these two facts, and by the fact that there are at most $|S_{i-1}|/11$ join events in an iteration, the total entrance costs paid by good IDs in iteration i is at most:

$$\left(1 + \frac{c(\text{A1}, \texttt{high})c(\text{A2}, \texttt{high})}{c(JE, \texttt{low})} \right) \left(\frac{|S_{i-1}|}{11} \right).$$

Adding the bounds for both entrance and purge costs and dividing by the value of ℓ_i yields a total cost of at most:

$$
\begin{aligned}
&|S_{i-1}| + \left(1 + \frac{c(\text{A1}, \texttt{high})c(\text{A2}, \texttt{high})}{c(JE, \texttt{low})} \right) \left(\frac{|S_{i-1}|}{11} \right) \\
&= \frac{|S_{i-1}|}{\ell_i} \left(\frac{12}{11} + \frac{c(\text{A1}, \texttt{high})c(\text{A2}, \texttt{high})}{11c(JE, \texttt{low})} \right)
\end{aligned}
$$

which completes the proof. $\qquad\square$

Lemma 17. *Let u and v be n-dimensional vectors. Then*

$$\sum_{j=1}^{n} \sqrt{u_j v_j} \leq \sqrt{\left(\sum_{j=1}^{n} u_j\right)\left(\sum_{j=1}^{n} v_j\right)}.$$

Proof. Let x and y be vectors such that for all $1 \leq i \leq n$, $x_i = \sqrt{u_i}$ and $y_i = \sqrt{v_i}$. Then applying the Cauchy-Schwarz inequality to x and y, we have that $x \cdot y \leq |x||y|$, from which the lemma follows directly. $\qquad\square$

Let $\mathcal{I}$ be any subset of iterations that for integers x and y, $1 \leq x \leq y$, contains every iteration with index between x and y inclusive. Let $\delta(\mathcal{I})$ be $|S_x - S_y|$; and let $\Delta(\mathcal{I})$ be $\delta(\mathcal{I})$ divided by the length of $\mathcal{I}$. Let $\mathcal{A}_\mathcal{I}$ and $T_\mathcal{I}$ be the algorithmic and adversarial spend rates over $\mathcal{I}$; and let $J_\mathcal{I}^G$ be the good join rate over all of $\mathcal{I}$. Then we have the following lemma.

Lemma 18. *For any subset of contiguous iterations, $\mathcal{I}$, which starts after iteration 1, it holds that:*

$$11c(L9)\left(2\Delta(\mathcal{I}) + c(L8)\sqrt{2T_\mathcal{I}(c(JE,\texttt{high})J_\mathcal{I}^G + 1)} + J_\mathcal{I}^G\right).$$

Proof. By Lemma 16, we have:

$$\sum_{i \in \mathcal{I}} \mathcal{A}_i \ell_i \leq c(L9) \sum_{i \in \mathcal{I}} |S_{i-1}| \tag{6}$$

Let D_i be the departure rate of both good and bad IDs during iteration i. By the conditions for when a purge occurs, we have that:

$$
\begin{aligned}
\sum_{i \in \mathcal{I}} \frac{|S_{i-1}|}{11} &\leq \sum_{i \in \mathcal{I}} (D_i + J_i^B + J_i^G)\ell_i \\
&\leq 2\delta(\mathcal{I}) + \sum_{i \in \mathcal{I}} J_i^B \ell_i + \sum_{i \in \mathcal{I}} J_i^G \ell_i \\
&\leq 2\delta(\mathcal{I}) + \sum_{i \in \mathcal{I}} \sqrt{2T_i \ell_i (c(JE,\texttt{high})J_i^G + 1)\ell_i} + \sum_{i \in \mathcal{I}} J_i^G \ell_i \\
&\leq 2\delta(\mathcal{I}) + \sqrt{\left(\sum_{i \in \mathcal{I}} 2T_i \ell_i\right)\left(\sum_{i \in \mathcal{I}} (c(JE,\texttt{high})J_i^G + 1)\ell_i\right)} + \sum_{i \in \mathcal{I}} J_i^G \ell_i
\end{aligned}
$$

The second line in the above follows from the fact that every ID that departs must have departed from the set of IDs in the system at the start of $\mathcal{I}$ or else must have been an ID that joined during $\mathcal{I}$. The third line follows from the Lemma 15 bound on J_i^B, and noting that $\ell_i = \sqrt{\ell_i^2}$. Finally, the last line follows from Lemma 17, by letting $u_i = 2T_i \ell_i$ and $v_i = (c(JE,\texttt{high})J_i^G + 1)\ell_i$ for all $i \in \mathcal{I}$.

Now combining Equation 6 with the above inequality, we have:

$$\sum_{i \in \mathcal{I}} \mathcal{A}_i \ell_i \leq 11c(L9)\left(2\delta(\mathcal{I}) + \sqrt{\left(\sum_{i \in \mathcal{I}} 2T_i \ell_i\right)\left(\sum_{i \in \mathcal{I}} (c(JE,\texttt{high})J_i^G + 1)\ell_i\right)} + \sum_{i \in \mathcal{I}} J_i^G \ell_i\right)$$

Dividing both sides of the above inequality by $\ell_{tot} = \sum_{i \in \mathcal{I}} \ell_i$ and recalling that $c(L8) = \sqrt{2c(JE, \texttt{high})}$, we get

$$
\begin{aligned}
\mathcal{A}_{\mathcal{I}} &\leq 11c(L9) \left(\frac{2\delta(\mathcal{I})}{\ell_{tot}} + c(L8) \sqrt{\frac{\left(\sum_{i \in \mathcal{I}} T_i \ell_i\right)}{\ell_{tot}} \frac{\left(\sum_{i \in \mathcal{I}} (J_i^G + 1)\ell_i\right)}{\ell_{tot}}} + \frac{\sum_{i \in \mathcal{I}} J_i^G \ell_i}{\ell_{tot}} \right) \\
&= 11c(L9) \left(2\Delta(\mathcal{I}) + c(L8)\sqrt{2T_{\mathcal{I}}(c(JE, \texttt{high})J_{\mathcal{I}}^G + 1)} + J_{\mathcal{I}}^G \right)
\end{aligned}
$$

which completes the proof. $\square$

Maintenance of the Population Invariant is guaranteed by Lemma 8, and the Committee Invariant is guaranteed by Lemma 9. Corollary 1 follows directly from Lemma 18. Theorem 2 also follows from Lemma 18 by noting that $\Delta(\mathcal{I}) = 0$ when $\mathcal{I}$ is all iterations, since the system is initially empty.

Finally, in concluding this section, we note that the same enhancements to CCOM presented in Section 4 can be applied to GMCOM.

7 Lower Bounds

In this section, we provide a lower bound that applies to the class of algorithms which have the attributes $B1 - B3$ described below.

- **B1.** Each new ID must pay an entrance fee in order to join the system and this is defined by a ***cost function f***, which takes as inputs the good join rate and the adversarial join rate.

- **B2.** The algorithm executes over iterations, delineated by when the condition $n_i^a + n_i^d \geq \delta |\mathcal{S}_i|$ holds, for any fixed positive δ.

- **B3.** At the end of each iteration, each ID must pay $\Omega(1)$ to remain in the system.

We emphasize that B1 captures any cost function where the cost during an iteration depends only on the good join rate and the bad join rate. Our analysis of GMCOM makes this assumption since its cost function depends on estimates of the good join rate and the total, i.e. the good-ID join rate plus the bad-ID join rate. With regard to B2 and B3, recall that we wish to preserve the population invariant (i.e., a majority of good IDs). It is hard to imagine an algorithm that preserves this invariant without a computational test being imposed on all IDs whenever the system membership changes significantly.

7.1 Lower-Bound Analysis

Restating in terms of the conditions above, we have the following result.

Theorem 2. *Suppose an algorithm satisfies conditions B1-B3, then there exists an adversarial strategy that forces the algorithm to spend at a rate of $\Omega(\sqrt{T\,J^G} + J^G)$, where J^G is the good ID join rate, and T is the algorithmic spending rate, both taken over the iteration.*

Proof. Fix an iteration i. Let n be the number of IDs in the system at the start of iteration i. The adversarial will have bad IDs join uniformly at the maximum rate possible, and then have the bad IDs drop out during the purge. In particular, let J^B be the rate at which bad IDs join, and let $f(J^B, J^G)$ be the algorithm's entrance cost function based on J^B and J^G; we pessimistically assume that the algorithm knows both J^B and J^G exactly. Then $J^B = T/f(J^B, J^G)$.

We first calculate the algorithmic spending rate due to purge puzzle costs in iteration i. Since the iteration ends after $\Theta(n)$ join events (B2), and since each purge puzzle has a cost of $\Omega(1)$ (B3), the rate of spending on purge puzzles is asymptotically at least equal to the rate at which good and bad IDs join the system. Thus, the spending rate due to purge puzzles is $\Omega(J^G + J^B)$.

We now have two cases:

Case 1: $f(J^B, J^G) \leq J^B/J^G$. In this case, we have:

$$J^B \geq J^G f(J^B, J^G) = J^G T/J^B$$

where the above equality holds since $f(J^B, J^G) = T/J^B$. Solving for J^B we get:

$$J^B \geq \sqrt{T J^G}$$

Since the algorithmic spending rate due to purge costs is $\Omega(J^B + J^G)$, the total algorithmic spending rate is:

$$\Omega(J^B + J^G) = \Omega\left(\sqrt{T J^G} + J^G\right)$$

Case 2: $f(J^B, J^G) > J^B/J^G$. In this case, we have:

$$J^B < J^G f(J^B, J^G) = J^G T/J^B$$

where the above equality follows since $f(J^B, J^G) = T/J^B$. Solving for J^B, we get:

$$J^B < \sqrt{T J^G}.$$

The spending rate for the algorithm due to entrance costs is $\Omega(J^G f(J^B, J^G))$. Adding in the spending rate for purge costs of $\Omega(J^G + J^B)$, we get that the total spend rate of the algorithm is:

$$
\begin{aligned}
\Omega(J^G f(J^B, J^G) + (J^G + J^B)) &= \Omega(J^G T/J^B + (J^G + J^B)) \\
&= \Omega\left(\sqrt{T J^G} + J^G\right)
\end{aligned}
$$

$\square$

8 Experiments

In this section, we report our simulation results for CCOM and GMCOM. Our goal is to examine the performance predicted by our theoretical results with respect to the computational cost of solving puzzles. Given this goal, we do not model Byzantine consensus or committee formation. Throughout all of our experiments, we assume a computational cost of k for solving a puzzle of difficulty k.

We present our findings as follows. In Section 8.1, we validate our assumptions from Section 1.5 over real-world peer-to-peer networks. In Section 8.2, we evaluate the computational cost for the proposed algorithms, CCOM and GMCOM, as a function of the adversarial cost, and we compare our algorithms against another PoW-based Sybil Defense algorithm. Finally, in Section 8.3, we propose and implement several heuristics to optimize the performance of GMCOM.

Our experiments make use of data from four real-world networks:

- *Bitcoin:* This dataset consists of roughly 7 days of join/departure-event timestamps; this data was obtained by personal correspondence with Till Neudecker [116]. The dataset records the join and departure events of IDs in the network, timestamped to the second.

- *BitTorrent Debian:* Peers connecting to the BitTorrent network to obtain a Debian ISO image.

- *BitTorrent RedHat:* Peers connecting to the BitTorrent network to obtain a RedHat ISO image.

- *BitTorrent FlatOut:* Peers connecting to the BitTorrent network to obtain aa demo of the game Flatout.

For the BitTorrent data, our simulations assume the Weibull distribution to model session time, with shape parameter values 0.38, 0.59 and 0.34, and scale parameter values 42.2, 41.0 and 21.3; these values are chosen based on a well-cited study by Stuzbach and Rejaie [134].

Finally, we note that our simulation code is written in MATLAB and can be obtained from [1]. All of our experiments were executed on a MacBook Air with macOS Mojave (version 10.14.3) using a 1.6 GHz Intel Core i5 processor and 8 GB of 2133 MHz LPDDR3.

8.1 Testing GMCOM Join and Departure Assumptions

To validate our assumptions from Section 1.5, we performed empirical studies on the join rate of IDs. Recall from Section 1.5, we define ρ_i to be the rate of join of good IDs in epoch i. Since we are interested in only the join rate of good IDs, we assume all IDs that ever join the system during our simulations in this section are good.

Experimental Setup. For the Bitcoin network, the system initially consists of 9212 IDs, and the join and departure events are based off the dataset from [117]. For the BitTorrent networks, we initialize the three networks with 1000 IDs each, and simulate the join and departure events over 1000 epochs. We verified our observations over 20 independent runs for the four networks.

In order to test if Assumption A1 holds, for each epoch $i \in \{2, 3, ..., 1000\}$, the join rate of epoch i, ρ_i, was compared to that join rate of the previous epoch, ρ_{i-1}. This was done for each of the four networks, the results are plotted in Figure 5(a) for the Bitcoin network, and Figure 6(a), (c) and (e) for BitTorrent Debian, BitTorrent RedHat and BitTorrent FlatOut networks, respectively.
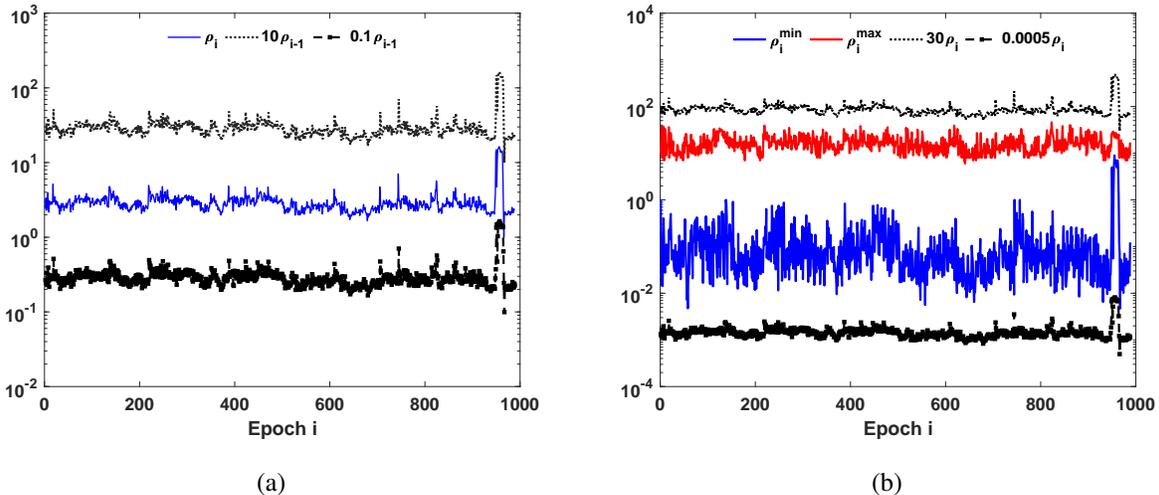

(a)   (b)

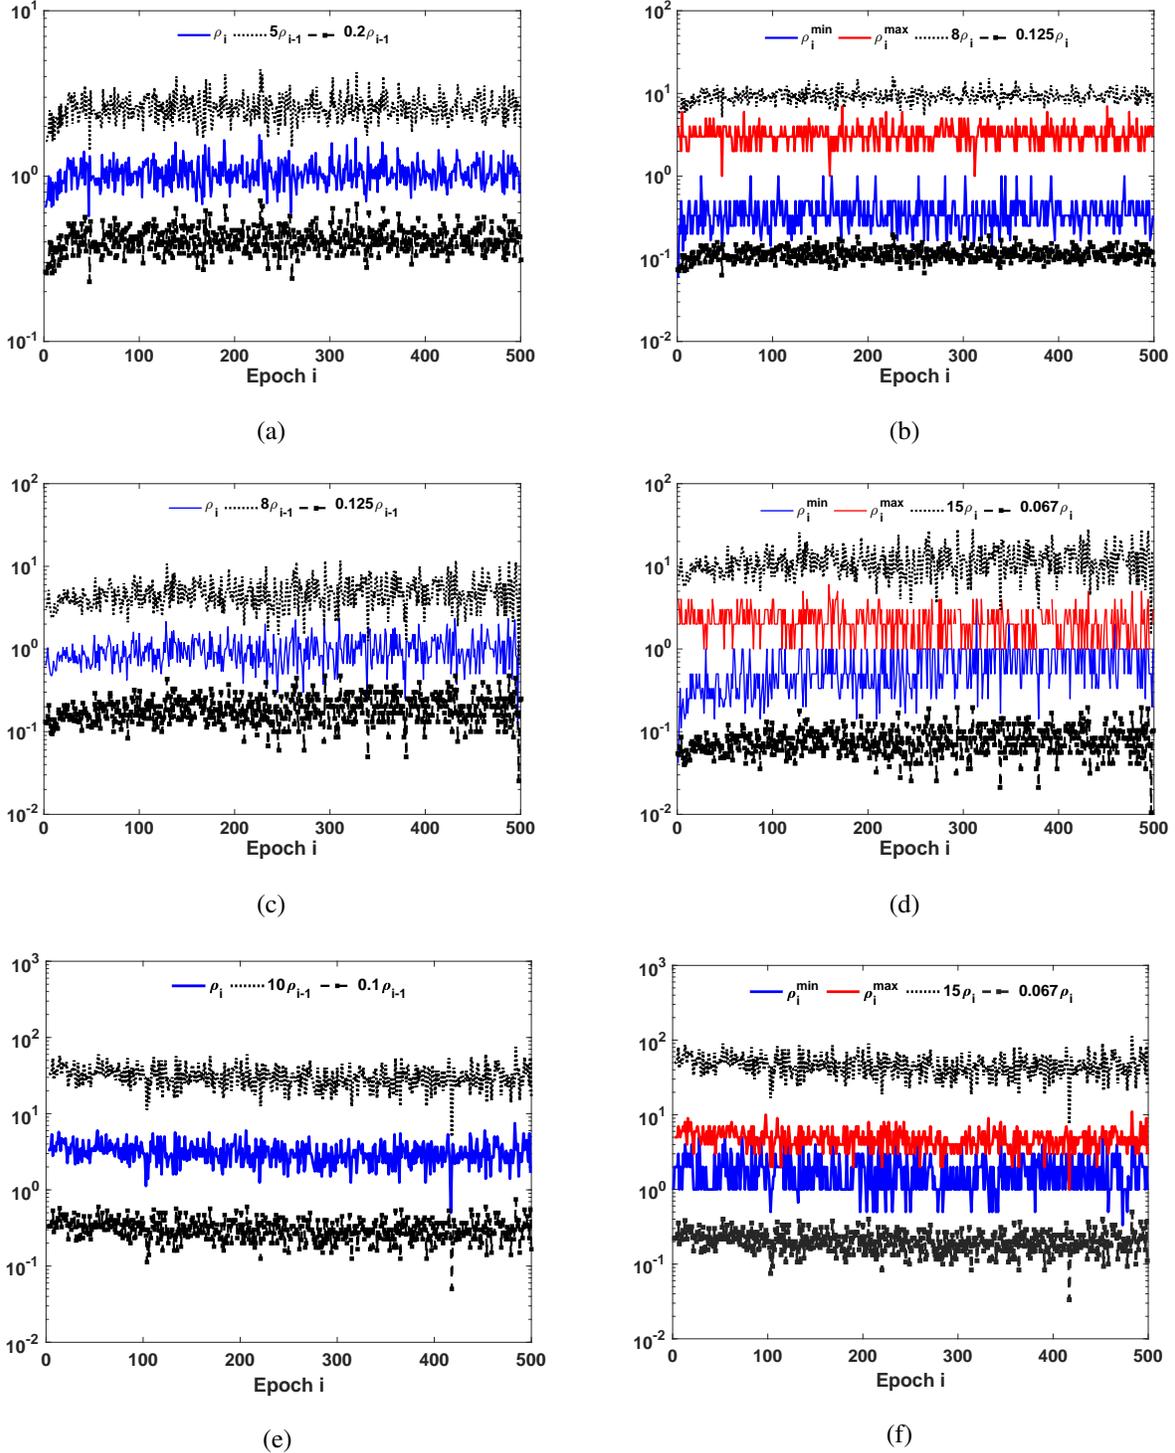Figure 5: Testing Assumptions A1 and A2 for Bitcoin.

Figure 6: Testing Assumptions A1 and A2 for: (a) and(b) Debian, (c) and (d) RedHat, (e) and (f) FlatOut.

Next, we test Assumption A2 for each of the four networks. In particular, we test for each epoch $i \in \{1, 2, ..., 1000\}$, that the join rate between two successive joins by good IDs is within some constant of

ρ_i. To do so, we measured the minimum and maximum join rate for epoch i, denoted by ρ_i^{min} and ρ_i^{max}, respectively, and compared these values against ρ_i.

The result are plotted in Figure 5 (b) for the Bitcoin network, and Figure 6(b), (d) and (f) for BitTorrent Debian, BitTorrent RedHat and BitTorrent FlatOut Networks, respectively. Table 3 lists the values of constants attached to our assumptions.

Network	$c(\text{A1}, \texttt{low})$	$c(\text{A1}, \texttt{high})$	$c(\text{A2}, \texttt{low})$	$c(\text{A2}, \texttt{high})$
Bitcoin Network	0.1	10	0.0005	30
BitTorrent Debian	0.2	5	0.125	8
BitTorrent RedHat	0.125	8	0.067	15
BitTorrent FlatOut	0.1	10	0.067	15

Table 3: Constants Assumptions A1 and A2 for Section 8.1

8.2 Testing Computational Cost

We simulate CCOM and GMCOM to validate that they exhibit the predicted spending rate from Theorem 2. Specifically, we validate the asymmetric spending rate $O(\sqrt{T(J^G + 1)} + J^G)$, where T is the adversary's computational cost for solving puzzles divided by the duration of an attack, and J^G is the average join rate of good IDs over the duration of the experiment for the four networks.

We compare the performance of our algorithms against another PoW-based Sybil defense, SYBILCONTROL [100]. Under this algorithm, each ID must solve a puzzle to join the system. Additionally, each ID periodically tests its neighbors with a puzzle, removing from its list of neighbors those IDs that fail to provide a solution within a time limit; these tests are not coordinated between IDs. An ID may be a neighbor to more than one ID and so receive multiple puzzles; these are combined into a single puzzle whose solution satisfies all the received puzzles.

Similar to Section 8.1, the Bitcoin network initially consists of 9212 IDs and the join and departure events are scheduled based on the events provided in the dataset. Also, for the three BitTorrent network, we initialize the system with 1000 IDs and the join/ departure events are scheduled based on the parameters and distribution from [117].

We assume $\alpha = 1/14$, and T ranges over $[2^0, 2^{100}]$, where for each value of T, the system is simulated for $10,000$ seconds. The adversary solves entrance puzzles to add bad IDs to the system. We pessimistically ignore the cost paid by the adversary for solving purge puzzles. We average our results over 20 independent runs for each of the four networks for the three algorithms. To avoid cluttering the plots, we omit error bars as they are negligible.

Figure 7 illustrates our results. The blue line depicts the average computational cost to good IDs per second obtained by executing GMCOM when the adversary spends T per second. The green and red lines are the plots for CCOM and SYBILCONTROL, respectively.

Note that the x-axis and y-axis are both $\log$ scaled. Initially, the blue line increases slowly due to T not being substantially larger than J^G, and hidden constants. However, as T grows, we observe behavior very close to $G = \sqrt{T}$, which validates the asymptotic behavior of the asymmetric cost. We cut off the red line at the point at which SYBILCONTROL is no longer able to ensure that the fraction of bad IDs is less than $1/2$.
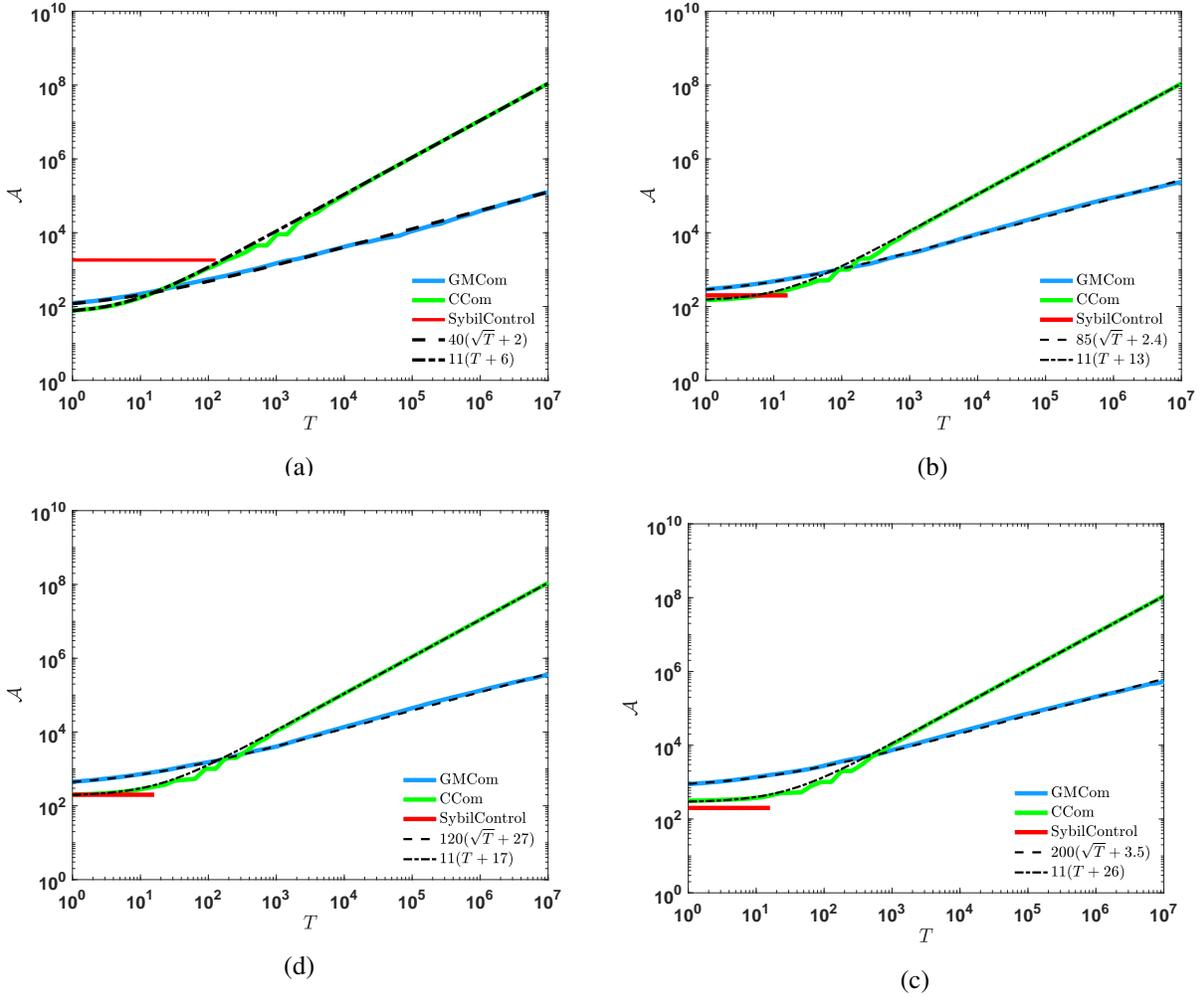
Figure 7: Comparison of algorithmic cost to adversarial cost for GMCOM, CCOM and SybilControl.

8.3 Proposed Heuristics

In this section, we explore heuristics to improve performance of our algorithm. To determine effective heuristics, we first study two separate costs: the purge cost and the entrance cost to the good IDs. We studied these two costs for the Bitcoin Network in the absence of an attack.

In Figure 8, we plot the costs as a function of time. In order to better visualize the frequency of purges, the x-axis is a log plot. It is clear that our algorithmic cost is dominated by the cost of purges. Hence, we focus on reducing the frequency of purges, without compromising correctness. We describe three heuristics below, followed by an evaluation of their performance.
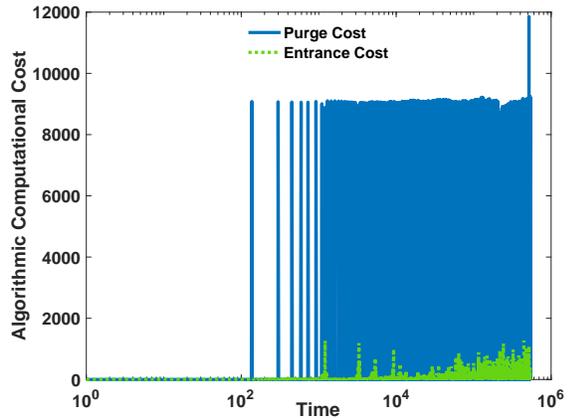

Figure 8: Comparison of purge cost to entrance cost for Bitcoin Network.

Heuristic 1: We use the symmetric difference to determine when to do a purge. Specifically, for iteration i, if $|(\mathcal{S}_{\text{cur}} \cup \mathcal{S}_{i-1}) - (\mathcal{S}_{\text{cur}} \cap \mathcal{S}_{i-1})| \geq |\mathcal{S}_{i-1}|/11$, then a purge is executed. This ensures that the fraction of bad IDs can increase by no more than in our original specification. But it potentially lengthens the amount of time between purges, for example, in the case when some ID joins and departs repeatedly.

Heuristic 2: We use the estimated good join rate to bound the maximum number of bad IDs that might have joined during the current iteration. This allows us to upper-bound the fraction of bad IDs in the system and purge only when the system's correctness may be at risk. We use this fact in conjunction with Heuristic 1 to safely delay the purge.

Heuristic 3: We propose the use of existing Sybil-detection techniques to remove some Sybil IDs from the system. For example, recent works have explored the possibility of identifying Sybil IDs based on the network topology [**?**, **?**, **?**]. In our experiments, we focus on SybilInfer from [**?**], and use it as a black box to possibly detect and discard Sybil IDs as they enter the system. The accuracy of the blackbox in detecting honest and Sybil IDs is a parameter in our experiments. We focus on two specific accuracy values, 0.98 and 0.92, based on emperical results in [**?**].

To evaluate the performance of these heuristics, we use GMCOM as a baseline. The experimental setup is the same as Section 8.2, with the same four data sets. We define ***GMCH1*** to be GMCOM using both
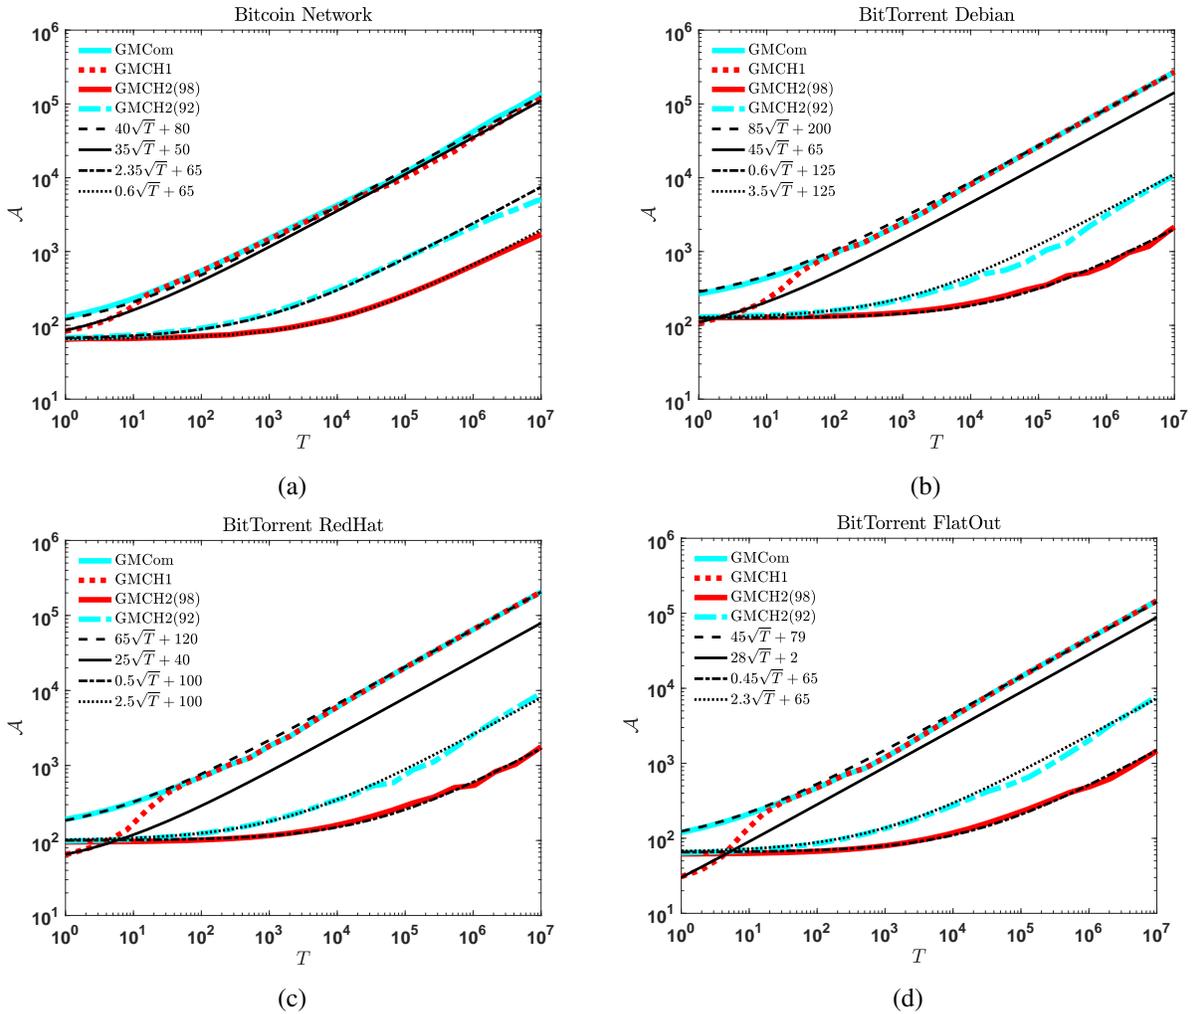

Figure 9: Comparison of algorithmic cost to adversarial cost for GMCOM and proposed heuristics.

Heuristic 1 and Heuristic 2. We define **GMCH2(98)** and **GMCH2(92)** to be GMCOM using Heuristics 1 and 2, and also Heuristic 3, with the accuracy parameter of Heuristic 3 as 0.98 and 0.92.

Figure 9 illustrates our results. These indicate that the GMCH1 is effective in reducing costs when adversarial effort is limited. On the other hand, GMCH2 reduces costs significantly during adversarial attack, with improvements of up to two orders of magnitude during the most significant attacks tested.

9 Conclusion and Future Work

We have presented and analyzed two PoW-based algorithms for defending against the Sybil attack. Our empirical work suggests that these algorithms are competitive with the contemporary PoW-based Sybil defense, SYBILCONTROL, and that they can be significantly more efficient when the system is under attack. Finally, we have proved a lower bound showing that our algorithm's computational cost is asymptotically optimal among a large class of Sybil-defense algorithms.

Many open problems remain including the following. One of particular interest is: Can we adapt our technique to secure multi-party computation? The problem of *secure multi-party computation (MPC)* involves designing an algorithm for the purpose of computing the value of an n-ary function f over private inputs from n IDs $x_1, x_2, ..., x_n$, such that the IDs learn the value of $f(x_1, x_2, ..., x_n)$, but learn nothing more about the inputs than what can be inferred from this output of f. The problem is generally complicated by the assumption that an adversary controls a hidden subset of the IDs. In recent years, a number of attempts have been made to solve this problem for very large n manner [14, 25, 35, 51, 52, 55, 76]. We believe that our technique for forming committees in a dynamic network could be helpful to solve a dynamic version of this problem, while ensuring that the resource costs to the good IDs is commensurate with the resource costs to an adversary. Secure MPC may have applications in the area of smart contracts for cryptocurrencies [**?**, **?**, **?**].

Another avenue of future work considers scenarios where "good" IDs may not always blindly follow our algorithm. Instead, these IDs may be rational but selfish, in that they seek to optimize some known utility function. The adversary still behaves in a worst-case manner, capturing the fact that the utility function of the adversary may be completely unknown. This approach is similar to BAR (Byzantine, Altruistic, Rational) games [45] in distributed computing, see also [2, 7, 137, 145]. Extending our result to accommodate this model is of interest.

Finally, there is a substantial body of literature on attack-resistant overlays; for example [10, 20–23, 67, 79, 83, 94, 113, 114, 127, 128, 130, 147, 148]. These results critically depend on the fraction of bad IDs always being upper-bounded by a constant less than $1/2$. However, there are only a handful of results that propose a method for guaranteeing this bound. A natural idea is to examine whether our algorithms can be applied to this setting in order to guarantee this bound, via the Population Invariant.

References

[1] Source code is available at `https://dikshagupta.blog/research/`.

[2] ABRAHAM, I., ALVISI, L., AND HALPERN, J. Y. Distributed Computing Meets Game Theory: Combining Insights from Two Fields. *SIGACT News 42*, 2 (June 2011), 69–76.

[3] ABRAHAM, I., DEVADAS, S., DOLEV, D., NAYAK, K., AND REN, L. Synchronous byzantine agreement with expected o (1) rounds, expected o (n2) communication, and optimal resilience. *IACR Cryptology ePrint Archive 2018* (2018), 1028.

[4] ABRAHAM, I., AND MALKHI, D. The Blockchain Consensus Layer and BFT. *Bulletin of the EATCS: Distributed Computing Column* (2017).

[5] ADYA, A., BOLOSKY, W. J., CASTRO, M., CERMAK, G., CHAIKEN, R., DOUCEUR, J. R., HOWELL, J., LORCH, J. R., THEIMER, M., AND WATTENHOFER, R. P. FARSITE: Federated, Available, and Reliable Storage for Incompletely Trusted Environment. In 5^{th} USENIX *Symposium on Operating Systems Design and Implementation (OSDI)* (2002), pp. 1–14.

[6] AGGARWAL, A., DANI, V., HAYES, T., AND SAIA, J. Secure One-Way Interactive Communication. In *Proceedings of the 15th International Conference on Distributed Computing and Networking (ICDCN)* (2017).

[7] AIYER, A. S., ALVISI, L., CLEMENT, A., DAHLIN, M., MARTIN, J.-P., AND PORTH, C. BAR Fault Tolerance for Cooperative Services. In *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles (SOSP)* (2005), pp. 45–58.

[8] ALYSSA HERTIG. Ethereum's Big Switch: The New Roadmap to Proof-of-Stake. www.coindesk.com/ethereums-big-switch-the-new-roadmap-to-proof-of-stake/.

[9] AMAZON. Amazon Web Service (AWS). `https://aws.amazon.com/`.

[10] ANCEAUME, E., LUDINARD, R., RAVOAJA, A., AND BRASILEIRO, F. V. PeerCube: A Hypercube-Based P2P Overlay Robust against Collusion and Churn. In *Proceedings of the 2^{nd} IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)* (2008), pp. 15–24.

[11] ANDERSON, D. P. BOINC: A System for Public-Resource Computing and Storage. In *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing (GRID)* (2004), pp. 4–10.

[12] ANDERSON, R., BARTON, C., BÖHME, R., CLAYTON, R., VAN EETEN, M. J., LEVI, M., MOORE, T., AND SAVAGE, S. Measuring the Cost of Cybercrime. In *The Economics of Information Security and Privacy*. Springer, 2013, pp. 265–300.

[13] ANDRYCHOWICZ, M., AND DZIEMBOWSKI, S. PoW-Based Distributed Cryptography with No Trusted Setup. In *Annual Cryptology Conference* (2015), Springer, pp. 379–399.

[14] APPLEBAUM, B., ISHAI, Y., AND KUSHILEVITZ, E. From secrecy to soundness: Efficient verification via secure computation. *Automata, languages and Programming* (2010), 152–163.

[15] AUGUSTINE, J., MOLLA, A. R., MORSY, E., PANDURANGAN, G., ROBINSON, P., AND UPFAL, E. 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)* (2013), pp. 53–62.

[16] AUGUSTINE, J., PANDURANGAN, G., AND ROBINSON, P. Fast Byzantine Agreement in Dynamic Networks. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)* (2013), pp. 74–83.

[17] AUGUSTINE, J., PANDURANGAN, G., AND ROBINSON, P. Fast Byzantine Leader Election in Dynamic Networks. In *International Symposium on Distributed Computing (DISC)* (2015), Springer, pp. 276–291.

[18] AUGUSTINE, J., PANDURANGAN, G., ROBINSON, P., ROCHE, S., AND UPFAL, E. 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)* (2015), pp. 350–369.

[19] AUGUSTINE, J., PANDURANGAN, G., ROBINSON, P., AND UPFAL, E. 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)* (2012), pp. 551–569.

[20] AWERBUCH, B., AND SCHEIDELER, C. Group Spreading: A Protocol for Provably Secure Distributed Name Service. In 31^{st} *International Colloquium on Automata, Languages, and Programming (ICALP)* (2004), pp. 183–195.

[21] AWERBUCH, B., AND SCHEIDELER, C. Robust Random Number Generation for Peer-to-peer Systems. In *10th International Conference On Principles of Distributed Systems (OPODIS)* (2006), pp. 275–289.

[22] AWERBUCH, B., AND SCHEIDELER, C. Towards a Scalable and Robust DHT. In *Proceedings of the 18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)* (2006), pp. 318–327.

[23] AWERBUCH, B., AND SCHEIDELER, C. Towards Scalable and Robust Overlay Networks. In *Proceedings of the 6^{th} International Workshop on Peer-to-Peer Systems (IPTPS)* (2007), p. n. pag.

[24] BAZZI, R. A., AND KONJEVOD, G. On the Establishment of Distinct Identities in Overlay Networks. In *Proceedings 24^{th} Annual ACM Symposium on Principles of Distributed Computing (PODC)* (2005), pp. 312–320.

[25] BEERLIOVÁ-TRUBINIOVÁ, Z., AND HIRT, M. Efficient multi-party computation with dispute control. In *TCC* (2006), vol. 3876, Springer, pp. 305–328.

[26] BELLARE, M., AND ROGAWAY, P. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security (CCS)* (1993), pp. 62–73.

[27] BENDER, M. A., FINEMAN, J. T., GILBERT, S., AND YOUNG, M. How to Scale Exponential Backoff: Constant Throughput, Polylog Access Attempts, and Robustness. In *Proceedings of the 27^{th} Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2016), pp. 636–654.

[28] BENDER, M. A., FINEMAN, J. T., MOVAHEDI, M., SAIA, J., DANI, V., GILBERT, S., PETTIE, S., AND YOUNG, M. Resource-Competitive Algorithms. *SIGACT News 46*, 3 (Sept. 2015), 57–71.

[29] BHANDARI, V., AND VAIDYA, N. H. On Reliable Broadcast in a Radio Network. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)* (2005), pp. 138–147.

[30] BHANDHARI, V., KATZ, J., KOO, C.-Y., AND VAIDYA, N. Reliable Broadcast in Radio Networks: The Bounded Collision Case. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)* (2006), pp. 258 – 264.

[31] BITCOINWIKI. BitcoinWiki Network. `https://en.bitcoin.it/wiki/Network\#Standard_relaying`.

[32] BITCOINWIKI. Mining Hardware Comparison, 2016. `https://en.bitcoin.it/wiki/Mining_hardware_comparison`.

[33] BITCOINWIKI. Non-Specialized Hardware Comparison, 2016. `https://en.bitcoin.it/wiki/Non-specialized_hardware_comparison`.

[34] BLOCKSTACK. Blockstack: The New Decentralized Internet, 2017. `https://blockstack.org/`.

[35] BOGETOFT, P., CHRISTENSEN, D. L., DAMGÅRD, I., GEISLER, M., JAKOBSEN, T. P., KRØIGAARD, M., NIELSEN, J. D., NIELSEN, J. B., NIELSEN, K., PAGTER, J., ET AL. Secure multiparty computation goes live. In *Financial Cryptography* (2009), vol. 5628, Springer, pp. 325–343.

[36] BONNEAU, J., MILLER, A., CLARK, J., NARAYANAN, A., KROLL, J. A., AND FELTEN, E. W. SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *Proceedings of the IEEE Symposium on Security and Privacy (SP)* (2015), pp. 104–121.

[37] BORISOV, N. Computational Puzzles As Sybil Defenses. In *Proceedings of the Sixth IEEE International Conference on Peer-to-Peer Computing (P2P)* (2006), pp. 171–176.

[38] CACHIN, C., AND PORITZ, J. A. Secure Intrusion-Tolerant Replication on the Internet. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN)* (2002), pp. 167–176.

[39] CASTRO, M., AND LISKOV, B. Practical Byzantine Fault Tolerance. *Operating Systems Review 33* (1998), 173–186.

[40] CASTRO, M., AND LISKOV, B. Byzantine Fault Tolerance Can Be Fast. In *International Conference on Dependable Systems and Networks* (2001), pp. 513–518.

[41] CASTRO, M., AND LISKOV, B. Practical Byzantine Fault Tolerance and Proactive Recovery. *ACM Transactions on Computer Systems (TOCS) 20*, 4 (2002), 398–461.

[42] CHANDRA, A., AND WEISSMAN, J. Nebulas: Using Distributed Voluntary Resources to Build Clouds. In *Proceedings of the 2009 Conference on Hot Topics in Cloud Computing (HotCloud)* (2009).

[43] CHANG FENG, W., KAISER, E., CHI FENG, W., AND LUU, A. The Design and Implementation of Network Puzzles. In *Proceedings of the Annual Joint Conference of IEEE Computer and Communications Societies (INFOCOM)* (2005), pp. 2372–2382.

[44] CLEMENT, A., MARCHETTI, M., WONG, E., ALVISI, L., AND DAHLIN, M. Byzantine Fault Tolerance: The Time is Now. In *Proceedings of the Second Workshop on Large-Scale Distributed Systems and Middleware (LADIS)* (2008), pp. 1–4.

[45] CLEMENT, A., NAPPER, J., LI, H., MARTIN, J.-P., ALVISI, L., AND DAHLIN, M. Theory of BAR Games. In *Proceedings of the 26th Annual ACM Symposium on Principles of Distributed Computing (PODC)* (2007), pp. 358–359.

[46] CLEMENT, A., WONG, E., ALVISI, L., DAHLIN, M., AND MARCHETTI, M. Making Byzantine Fault Tolerant Systems Tolerate Byzantine Faults. In *Proceedings of the Sixth USENIX Symposium on Networked Systems Design and Implementation (NSDI)* (2009), pp. 153–168.

[47] COINDESK. Think Tank Reignites Debate Over Bitcoin Mining's Environmental Effects. `http://www.coindesk.com/think-tank-debate-bitcoin-mining-environment/`.

[48] CORE, C. Chain Core, 2017. `https://chain.com/`.

[49] COWLING, J., MYERS, D., LISKOV, B., RODRIGUES, R., AND SHRIRA, L. HQ Replication: A Hybrid Quorum Protocol for Byzantine Fault Tolerance. In *Proceedings of the Third USENIX Symposium on Operating Systems Design and Implementation (OSDI)* (1999), pp. 177–190.

[50] CROMAN, K., DECKER, C., EYAL, I., GENCER, A. E., JUELS, A., KOSBA, A., MILLER, A., SAXENA, P., SHI, E., SIRER, E. G., ET AL. On scaling decentralized blockchains. In *International Conference on Financial Cryptography and Data Security* (2016), Springer, pp. 106–125.

[51] DAMGARD, I., AND ISHAI, Y. Scalable secure multiparty computation. In *Crypto* (2006), vol. 4117, Springer, pp. 501–520.

[52] DAMGÅRD, I., ISHAI, Y., KRØIGAARD, M., NIELSEN, J., AND SMITH, A. Scalable multiparty computation with nearly optimal work and resilience. *Advances in Cryptology–CRYPTO 2008* (2008), 241–261.

[53] DANEZIS, G., LESNIEWSKI-LAAS, C., KAASHOEK, M. F., AND ANDERSON, R. Sybil-Resistant DHT Routing. In *Proceedings of the 10^{th} European Symposium On Research In Computer Security (ESORICS)* (2005), pp. 305–318.

[54] DANI, V., HAYES, T., MOVAHEDI, M., SAIA, J., AND YOUNG, M. Interactive Communication with Unknown Noise Rate. *Information and computation* (2017).

[55] DANI, V., KING, V., MOVAHEDI, M., AND SAIA, J. Quorums quicken queries: Efficient asynchronous secure multiparty computation. In *International Conference on Distributed Computing and Networking* (2014), Springer, pp. 242–256.

[56] DANI, V., MOVAHEDI, M., SAIA, J., AND YOUNG, M. Interactive Communication with Unknown Noise Rate. In *Proceedings of the Colloquium on Automata, Languages, and Programming (ICALP)* (2015).

[57] DASH. Dash is Digital Cash, 2017. (`https://www.dash.org/`.

[58] DEMIRBAS, M., AND SONG, Y. An RSSI-based Scheme for Sybil Attack Detection in Wireless Sensor Networks. In *Proceedings of the 2006 International Symposium on on World of Wireless, Mobile and Multimedia Networks (WOWMOM)* (2006), pp. 564–570.

[59] DINGER, J., AND HARTENSTEIN, H. Defending the Sybil Attack in P2P Networks: Taxonomy, Challenges, and a Proposal for Self-Registration. In *Proceedings of the First International Conference on Availability, Reliability and Security (ARES)* (2006), pp. 756–763.

[60] DOUCEUR, J. The Sybil Attack. In *Proceedings of the Second International Peer-to-Peer Symposium (IPTPS)* (2002), pp. 251–260.

[61] DWORK, C., LYNCH, N., AND STOCKMEYER, L. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM) 35*, 2 (1988), 288–323.

[62] DWORK, C., AND NAOR, M. Pricing via Processing or Combatting Junk Mail. In *Proceedings of the 12^{th} Annual International Cryptology Conference on Advances in Cryptology* (1993), pp. 139–147.

[63] ECONOMIST, T. Why bitcoin uses so much energy. https://www.economist.com/the-economist-explains/2018/07/09/why-bitcoin-uses-so-much-energy.

[64] ETHEREUM. Ethereum: Blockchain App. Program, 2016. `https://www.ethereum.org/`.

[65] EYAL, I., GENCER, A. E., SIRER, E. G., AND VAN RENESSE, R. Bitcoin-NG: A Scalable Blockchain Protocol. In *Proceedings of the 13^{th} USENIX Symposium on Networked Systems Design and Implementation (NSDI)* (2016), pp. 45–59.

[66] FALKNER, J., PIATEK, M., JOHN, J. P., KRISHNAMURTHY, A., AND ANDERSON, T. Profiling a Million User DHT. In 7^{th} *ACM SIGCOMM Conference on Internet Measurement* (2007), pp. 129–134.

[67] FIAT, A., SAIA, J., AND YOUNG, M. Making Chord Robust to Byzantine Attacks. In *Proceedings of the 13^{th} European Symposium on Algorithms (ESA)* (2005), pp. 803–814.

[68] GARAY, J., KIAYIAS, A., AND LEONARDOS, N. The Bitcoin Backbone Protocol: Analysis and Applications. In *Proceedings of 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)* (2015), pp. 281–310.

[69] GIL, S., KUMAR, S., MAZUMDER, M., KATABI, D., AND RUS, D. Guaranteeing Spoof-Resilient Multi-Robot Networks. In *Proceedings of Robotics: Science and Systems* (Rome, Italy, July 2015).

[70] GILAD, Y., HEMO, R., MICALI, S., VLACHOS, G., AND ZELDOVICH, N. Algorand: Scaling Byzantine Agreements for Cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP)* (2017), pp. 51–68.

[71] GILBERT, S., KING, V., PETTIE, S., PORAT, E., SAIA, J., AND YOUNG, M. (Near) Optimal Resource-Competitive Broadcast with Jamming. In *Proceedings of the 26^{th} ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)* (2014), pp. 257–266.

[72] GILBERT, S., KING, V., SAIA, J., AND YOUNG, M. Resource-Competitive Analysis: A New Perspective on Attack-Resistant Distributed Computing. In *Proceedings of the 8^{th} ACM International Workshop on Foundations of Mobile Computing* (2012).

[73] GILBERT, S., NEWPORT, C., AND ZHENG, C. Who Are You? Secure Identities in Ad Hoc Networks. In *Proceedings of the 28^{th} International Symposium on Distributed Computing (DISC)* (2014), pp. 227–242.

[74] GILBERT, S., AND YOUNG, M. Making Evildoers Pay: Resource-Competitive Broadcast in Sensor Networks. In *Proceedings of the 31^{th} Symposium on Principles of Distributed Computing (PODC)* (2012), pp. 145–154.

[75] GILBERT, S., AND ZHENG, C. SybilCast: Broadcast on the Open Airwaves. In *Proceedings of the 25^{th} Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)* (2013), pp. 130–139.

[76] GOLDREICH, O. Secure multi-party computation. *Manuscript. Preliminary version* (1998), 86–97.

[77] GORBUNOV, S., AND MICALI, S. Democoin: A Publicly Verifiable and Jointly Serviced Cryptocurrency. Cryptology ePrint Archive, Report 2015/521, 2015. `http://eprint.iacr.org/2015/521`.

[78] GREEN, J., JUEN, J., FATEMIEH, O., SHANKESI, R., JIN, D., AND GUNTER, C. A. Reconstructing Hash Reversal Based Proof of Work Schemes. In *Proceedings of the 4th USENIX Conference on Large-scale Exploits and Emergent Threats (LEET)* (2011), pp. 10–10.

[79] GUERRAOUI, R., HUC, F., AND KERMARREC, A.-M. Highly Dynamic Distributed Computing with Byzantine Failures. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing (PODC)* (2013), pp. 176–183.

[80] GUPTA, D., SAIA, J., AND YOUNG, M. Proof of Work Without All the Work. In *Proceedings of the* 19th *International Conference on Distributed Computing and Networking (ICDCN)* (2018).

[81] GUPTA, D., SAIA, J., AND YOUNG, M. Peace through superior puzzling: An asymmetric sybil defense. In *33rd IEEE International Parallel and Distributed Processing Symposium (IPDPS)* (2019).

[82] HASHCAT. Benchmark Results, 2016. `http://thepasswordproject.com/oclhashcat_benchmarking`.

[83] HILDRUM, K., AND KUBIATOWICZ, J. Asymptotically Efficient Approaches to Fault-Tolerance in Peer-to-peer Networks. In *Proceedings of the 17th International Symposium on Distributed Computing* (2004), pp. 321–336.

[84] HOU, R., JAHJA, I., LUU, L., SAXENA, P., AND YU, H. Randomized view reconciliation in permissionless distributed systems.

[85] IOTA. IOTA: The Economy of Things, 2017. `http://iota.org/`.

[86] JOHN, R., CHERIAN, J. P., AND KIZHAKKETHOTTAM, J. J. A Survey of Techniques to Prevent Sybil Attacks. In *Proc. of the Intl. Conference on Soft-Computing and Networks Security (ICSNS)* (2015), pp. 1–6.

[87] KAISER, E., AND FENG, W. C. Mod_kaPoW: Mitigating DoS with Transparent Proof-of-Work. In *Proceedings of the 2007 ACM CoNEXT Conference* (2007), pp. 74:1–74:2.

[88] KAPRON, B., KEMPE, D., KING, V., SAIA, J., AND SANWALANI, V. Fast asynchronous byzantine agreement and leader election with full information. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (2008), SODA '08, pp. 1038–1047.

[89] KATZ, J., AND KOO, C.-Y. On expected constant-round protocols for byzantine agreement. *J. Comput. Syst. Sci. 75*, 2 (Feb. 2009), 91–112.

[90] KATZ, J., MILLER, A., AND SHI, E. Pseudonymous secure computation from time-lock puzzles.

[91] KING, V., PHILLIPS, C. A., SAIA, J., AND YOUNG, M. Sleeping on the Job: Energy-Efficient and Robust Broadcast for Radio Networks. *Algorithmica 61*, 3 (2011), 518–554.

[92] KING, V., SAIA, J., SANWALANI, V., AND VEE, E. Scalable Leader Election. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithm (SODA)* (2006), pp. 990–999.

[93] KING, V., SAIA, J., AND YOUNG, M. Conflict on a Communication Channel. In *Proceedings of the 30th Symposium on Principles of Distributed Computing (PODC)* (2011), pp. 277–286.

[94] KNOCKEL, J., SAAD, G., AND SAIA, J. Self-Healing of Byzantine Faults. In *Proceedings of the International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)* (2013), pp. 98–112.

[95] KOBLITZ, N., AND MENEZES, A. J. The Random Oracle Model: A Twenty-Year Retrospective. *Designs, Codes and Cryptography 77*, 2-3 (2015), 587–610.

[96] KOO, C.-Y. Broadcast in Radio Networks Tolerating Byzantine Adversarial Behavior. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)* (2004), pp. 275–282.

[97] KOTLA, R., ALVISI, L., DAHLIN, M., CLEMENT, A., AND WONG, E. Zyzzyva: Speculative Byzantine Fault Tolerance. In *Proceedings of* 21^{st} *Symposium on Operating Systems Principles* (2007), pp. 45–58.

[98] LAMPORT, L., SHOSTAK, R., AND PEASE, M. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems (TOPLAS) 4*, 3 (1982), 382–401.

[99] LESNIEWSKI-LAAS, C., AND KAASHOEK, M. F. Whanau: A Sybil-Proof Distributed Hash Table. In *Proceedings of the* 7^{th} *USENIX Conference on Networked Systems Design and Implementation* (2010), NSDI'10, pp. 8–8.

[100] LI, F., MITTAL, P., CAESAR, M., AND BORISOV, N. SybilControl: Practical Sybil Defense with Computational Puzzles. In *Proceedings of the Seventh ACM Workshop on Scalable Trusted Computing* (2012), pp. 67–78.

[101] LITECOIN. litecoin: Global Decentralized Currency, 2017. `https://litecoin.org/`.

[102] LIU, Y., BILD, D. R., DICK, R. P., MAO, Z. M., AND WALLACH, D. S. The Mason Test: A Defense Against Sybil Attacks in Wireless Networks Without Trusted Authorities. *IEEE Transactions on Mobile Computing 14*, 11 (2015), 2376–2391.

[103] LUU, L., NARAYANAN, V., ZHENG, C., BAWEJA, K., GILBERT, S., AND SAXENA, P. A Secure Sharding Protocol For Open Blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)* (2016), pp. 17–30.

[104] MARTINOVIC, I., ZDARSKY, F. A., WILHELM, M., WEGMANN, C., AND SCHMITT, J. B. Wireless Client Puzzles in IEEE 802.11 Networks: Security by Wireless. In *Proceedings of the First ACM Conference on Wireless Network Security* (2008), WiSec '08, pp. 36–45.

[105] MILLER, A., LITTON, J., PACHULSKI, A., SPRING, N., GUPTA, N., LEVIN, D., AND BHATTACHARJEE, B. Discovering Bitcoin's Public Topology and Influential Nodes, 2015. `http://cs.umd.edu/projects/coinscope/coinscope.pdf`.

[106] MITZENMACHER, M., AND UPFAL, E. *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, 2017.

[107] MOHAISEN, A., AND HOLLENBECK, S. Improving Social Network-Based Sybil Defenses by Rewiring and Augmenting Social Graphs. In *Proceedings of the* 14^{th} *International Workshop on Information Security Applications (WISA)* (2014), pp. 65–80.

[108] MOHAISEN, A., AND KIM, J. The Sybil Attacks and Defenses: A Survey. *Smart Computing Review 3*, 6 (2013), 480–489.

[109] MOHAISEN, A., TRAN, H., CHANDRA, A., AND KIM, Y. Trustworthy distributed computing on social networks. In *Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security* (2013), ASIA CCS '13, pp. 155–160.

[110] MOLINA, H. G., PITTELLI, F., AND DAVIDSON, S. Applications of Byzantine Agreement in Database Systems. *ACM Transactions on Database Systems (TODS) 11* (1986), 27–47.

[111] MÓNICA, D., LEITAO, L., RODRIGUES, L., AND RIBEIRO, C. On the Use of Radio Resource Tests in Wireless Ad-Hoc Networks. In *Proceedings of the 3rd Workshop on Recent Advances on Intrusion-Tolerant Systems* (2009), pp. F21–F26.

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

[113] NAOR, M., AND WIEDER, U. A Simple Fault Tolerant Distributed Hash Table. In *Proceedings of the Second International Workshop on Peer-to-Peer Systems (IPTPS)* (2003), pp. 88–97.

[114] NAOR, M., AND WIEDER, U. Novel Architectures for P2P Applications: The Continuous-Discrete Approach. In *Fifteenth ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)* (2003).

[115] NEUDECKER, T. Bitcoin Cash (BCH) Sybil Nodes on the Bitcoin Peer-to-Peer Network, 2017. `http://dsn.tm.kit.edu/publications/files/332/bch_sybil.pdf`.

[116] NEUDECKER, T., ANDELFINGER, P., AND HARTENSTEIN, H. A Simulation Model for Analysis of Attacks on the Bitcoin Peer-to-Peer Network. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management (IM)* (2015), pp. 1327–1332.

[117] NEUDECKER, T., ANDELFINGER, P., AND HARTENSTEIN, H. Timing analysis for inferring the topology of the bitcoin peer-to-peer network. In *Proceedings of the 13th IEEE International Conference on Advanced and Trusted Computing (ATC)* (July 2016).

[118] NEWSOME, J., SHI, E., SONG, D., AND PERRIG, A. The Sybil Attack in Sensor Networks: Analysis & Defenses. In *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks (IPSN)* (2004), pp. 259–268.

[119] OCEANSTORE. The Oceanstore Project. http://oceanstore.cs.berkeley.edu.

[120] OF STANDARDS, N. I., AND TECHNOLOGY. Secure Hash Standard (SHS). `http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.180-4.pdf`, 2015.

[121] PARNO, B., WENDLANDT, D., SHI, E., PERRIG, A., MAGGS, B., AND HU, Y.-C. Portcullis: Protecting Connection Setup from Denial-of-Capability Attacks. *ACM SIGCOMM Computer Communication Review 37*, 4 (2007), 289–300.

[122] PASS, R., SEEMAN, L., AND SHELAT, A. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques* (2017), Springer, pp. 643–673.

[123] PASS, R., AND SHI, E. Hybrid consensus: Efficient consensus in the permissionless model. In *31st International Symposium on Distributed Computing (DISC 2017)* (2017), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

[124] PREGUIÇA, N., RODRIGUES, R., HONORATO, C., AND LOURENÇO, J. Byzantium: Byzantine-Fault-Tolerant Database Replication Providing Snapshot Isolation. In *Proceedings of the Fourth Conference on Hot Topics in System Dependability* (2008), USENIX Association, p. 9.

[125] PRIMECOIN. Primecoin: A New Type of Cryptocurrency Which is Proof-of-work Based on Searching for Prime Numbers, 2017. (`http://primecoin.io/`.

[126] RODRIGUES, R., AND LISKOV, B. Rosebud: A Scalable Byzantine-Fault-Tolerant Storage Architecture. Tech. Rep. TR/932, MIT LCS, December 2003.

[127] SAAD, G., AND SAIA, J. Self-Healing Computation. In *Proceedings of the International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)* (2014), pp. 195–210.

[128] SAIA, J., AND YOUNG, M. Reducing Communication Costs in Robust Peer-to-Peer Networks. *Information Processing Letters 106(4)* (2008), 152–158.

[129] SCHEIDELER, C., AND SCHMID, S. A Distributed and Oblivious Heap. In *Proceedings of the 36th Internatilonal Collogquium on Automata, Languages and Programming: Part II* (2009), ICALP '09, pp. 571–582.

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

[131] SHERR, M., BLAZE, M., AND LOO, B. T. Veracity: Practical Secure Network Coordinates via Vote-based Agreements. In *Proceedings of the USENIX Annual Technical Conference* (2009), pp. 13 –13.

[132] SRINATHAN, K., AND RANGAN, C. P. Efficient asynchronous secure multiparty distributed computation. In *INDOCRYPT 2000, Lecture Notes in Computer Science*, vol. 1977. Springer-Verlag, 2000, pp. 117–129.

[133] STEINER, M., EN-NAJJARY, T., AND BIERSACK, E. W. A Global View of KAD. In 7^{th} *ACM SIGCOMM Conference on Internet Measurement* (2007), pp. 117 – 122.

[134] STUTZBACH, D., AND REJAIE, R. Understanding Churn in Peer-to-peer Networks. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement (IMC)* (New York, NY, USA, 2006), ACM, pp. 189–202.

[135] TECHNICA, A. Mining Bitcoins Takes Power, But is it an Environmental Disaster? http://arstechnica.com/business/2013/04/mining-bitcoins-takes-power-but-is-it-an-environmental-disaster/.

[136] TRAN, N., MIN, B., LI, J., AND SUBRAMANIAN, L. Sybil-Resilient Online Content Voting. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation (NSDI)* (2009), pp. 15–28.

[137] VILAÇA, X., LEITÃO, J. A., CORREIA, M., AND RODRIGUES, L. N-Party BAR Transfer. In *Proceedings of the 15th International Conference on Principles of Distributed Systems (OPODIS)* (2011), pp. 392–408.

[138] WALFISH, M., VUTUKURU, M., BALAKRISHNAN, H., KARGER, D., AND SHENKER, S. DDoS Defense by Offense. *ACM Transactions on Computer Systems (TOCS) 28*, 1 (2010), 3.

[139] WANG, H., ZHU, Y., AND HU, Y. An Efficient and Secure Peer-to-Peer Overlay Network. In *Proceedings of the IEEE Conference on Local Computer Networks* (2005), pp. 764–771.

[140] WANG, L., AND KANGASHARJU, J. Real-World Sybil Attacks in BitTorrent Mainline DHT. In *2012 IEEE Global Communications Conference (GLOBECOM)* (2012), pp. 826–832.

[141] WANG, X., AND REITER, M. K. Defending Against Denial-of-Service Attacks with Puzzle Auctions. In *Proceedings of the 2003 IEEE Symposium on Security and Privacy* (2003), p. 78.

[142] WATERS, B., JUELS, A., HALDERMAN, A., AND FELTEN, E. New Client Puzzle Outsourcing Techniques for DoS Resistance. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS)* (2004), pp. 246–256.

[143] WEI, W., XU, F., TAN, C. C., AND LI, Q. SybilDefender: A Defense Mechanism for Sybil Attacks in Large Social Networks. *IEEE Transactions on Parallel & Distributed Systems 24*, 12 (2013), 2492–2502.

[144] WEISSMAN, J. B., SUNDARRAJAN, P., GUPTA, A., RYDEN, M., NAIR, R., AND CHANDRA, A. Early Experience with the Distributed Nebula Cloud. In *Proceedings of the Fourth International Workshop on Data-intensive Distributed Computing (DIDC)* (2011), pp. 17–26.

[145] WONG, E. L., LENERS, J. B., AND ALVISI, L. It's on Me! The Benefit of Altruism in BAR Environment. In *Proceedings of the 24th International Conference on Distributed Computing (DISC)* (2010), pp. 406–420.

[146] YANG, Z., WILSON, C., WANG, X., GAO, T., ZHAO, B. Y., AND DAI, Y. Uncovering Social Network Sybils in the Wild. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference (IMC)* (2011), pp. 259–268.

[147] YOUNG, M., KATE, A., GOLDBERG, I., AND KARSTEN, M. Practical Robust Communication in DHTs Tolerating a Byzantine Adversary. In *Proceedings of the 30^{th} IEEE International Conference on Distributed Computing Systems (ICDCS)* (2010), pp. 263–272.

[148] YOUNG, M., KATE, A., GOLDBERG, I., AND KARSTEN, M. Towards Practical Communication in Byzantine-Resistant DHTs. *IEEE/ACM Transactions on Networking 21*, 1 (Feb. 2013), 190–203.

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

[150] YU, H., GIBBONS, P. B., KAMINSKY, M., AND XIAO, F. SybilLimit: A Near-Optimal Social Network Defense Against Sybil Attacks. *IEEE/ACM Trans. Netw. 18*, 3 (June 2010), 885—898.

[151] YU, H., KAMINSKY, M., GIBBONS, P. B., AND FLAXMAN, A. SybilGuard: Defending Against Sybil Attacks via Social Networks. *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM) 36* (Aug. 2006), 267–278.

[152] ZAMANI, M., SAIA, J., AND CRANDALL, J. TorBricks: Blocking-Resistant Tor Bridge Distribution. In *International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)* (2017), Springer, pp. 426–440.

[153] ZHAO, W. A Byzantine Fault Tolerant Distributed Commit Protocol. In *Proceedings of the 3^{rd} IEEE International Symposium on Dependable, Autonomic and Secure Computing* (2007), pp. 37–46.