

IRS: An Incentive-compatible Reward Scheme for Algorand

Maizi Liao
University of Waterloo
Waterloo, Canada
m7liao@uwaterloo.ca

Wojciech Golab
University of Waterloo
Waterloo, Canada
wgolab@uwaterloo.ca

Seyed Majid Zahedi
University of Waterloo
Waterloo, Canada
smzahedi@uwaterloo.ca

ABSTRACT

Founded in 2017, Algorand is one of the world’s first carbon-negative, public blockchains inspired by proof of stake. Algorand uses a Byzantine agreement protocol to add new blocks to the blockchain. The protocol can tolerate malicious users as long as a supermajority of the stake is controlled by non-malicious users. The protocol achieves about 100x more throughput compared to Bitcoin and can be easily scaled to millions of nodes. Despite its impressive features, Algorand lacks a reward-distribution scheme that can effectively incentivize nodes to participate in the protocol. In this work, we study the incentive issue in Algorand through the lens of game theory. We model the Algorand protocol as a Bayesian game and propose a novel reward scheme to address the incentive issue in Algorand. We derive necessary conditions to ensure that participation in the protocol is a Bayesian Nash equilibrium under our proposed reward scheme even in the presence of a malicious adversary. We also present quantitative analysis of our proposed reward scheme by applying it to two real-world deployment scenarios. We estimate the costs of running an Algorand node and simulate the protocol to measure the overheads in terms of computation, storage, and networking.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic mechanism design; Algorithmic game theory.**

KEYWORDS

Incentive-compatible reward schemes; Algorand; Bayesian games

ACM Reference Format:

Maizi Liao, Wojciech Golab, and Seyed Majid Zahedi. 2023. IRS: An Incentive-compatible Reward Scheme for Algorand. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

The concept of blockchain is first popularized by Bitcoin [44] as a tamper-resistant distributed transaction ledger. Blockchains could be classified into two categories: permissioned and permissionless. Permissioned blockchains, also known as private blockchains, implement an access-control mechanism to restrict unauthorized users from accessing the ledger [14, 37, 46]. Examples include HyperLedger Fabric [8] and Libra (now called Diem) [18]. In contrast, permissionless blockchains do not impose any access restrictions [32, 42]. Examples include Bitcoin [44] and Ethereum [56].

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

When obtaining permission is not required, the system could become prone to Sybil attacks¹. To mitigate the Sybil attack threat, permissionless consensus protocols often use additional mechanisms. For example, Bitcoin uses proof of work (PoW), which requires nodes to solve a computationally intensive puzzle. Winners earn the right to add blocks to the blockchain and collect rewards for their computational effort. PoW suffers from high energy and computational costs [24]. Proof of stake (PoS) has been proposed to mitigate these costs [35, 49]. In most PoS consensus protocols, nodes stake their cryptocurrency assets to gain rights to add blocks and earn rewards.

Inspired by proof of stake, Algorand is one of the world’s first carbon negative blockchain protocols [27]. The Algorand blockchain runs a randomized, committee-based consensus protocol [15, 29]. The core of the protocol is a Byzantine agreement protocol that allows nodes to reach consensus on a new block in the presence of Byzantine faults². Nodes are selected randomly to participate in the Byzantine agreement protocol as committee members. The original reward scheme of Algorand³ rewards all nodes proportionally to their account balance. Although simple, this reward scheme suffers from the free-rider problem: nodes have no incentive to participate in the protocol as doing so imposes computational and communication costs. This can be seen by tracking the number of nodes that actively participate in Algorand. According to [3], in May 2022, only about 1.6 billion units of Algorand’s cryptocurrency were registered to participate while a total of about 8.4 billion units of Algorand’s cryptocurrency were available. Moreover, while there were more than 1.7 million active accounts, only 361 unique accounts had recently participated in Algorand’s consensus protocol. Since the safety and liveness of Algorand depend mainly on high participation of nodes, the lack of participation poses a serious threat to Algorand by making it prone to attacks from malicious participants.

There exist other proof of stake based consensus protocols which have different properties comparing to Algorand. Avalanche’s consensus protocol [50] is inspired by gossip protocols and does not need a committee. Tendermint [10] is committee-based consensus protocol but the membership of the committee is deterministic and public. The consensus protocols of Ethereum 2, Polkadot and Cardano are similar. They use proof of stake to randomly select block proposers and then use a chain selection rule to resolve forks [11–13, 17, 34, 54, 55].

¹In a Sybil attack, the attacker creates a large number of pseudonymous identities to gain disproportionate control and/or influence over the system.

²Byzantine faults cause a node to inconsistently appear both failed and functioning to others. The term “Byzantine” is taken from the “Byzantine general problem” [38].

³Algorand is moving from its original reward scheme to a new reward scheme called the Governance Rewards [28]. Under the new reward scheme, only agents who commit to participate in the governance of the Algorand ecosystem will be rewarded. Agents have to prove their commitment by locking their cryptocurrency assets for a potentially long term. This new scheme is in line with proof-of-stake protocols.

The incentive problem in Bitcoin-like blockchains has been studied extensively in recent years [16, 21, 39, 47, 52]. However, the results do not apply to Algorand due to its unique consensus protocol. In this paper, we propose IRS, an incentive-compatible reward scheme for Algorand. Under IRS, nodes’ participation is monitored through committee votes. These votes serve as evidence of the committee members’ participation in the protocol. In addition, IRS requires committee members to include with their vote the identities of the nodes from which they have received a valid message. This facilitates the monitoring of the collaboration of nodes that are not chosen as committee members. We model Algorand as a Bayesian game and study nodes’ strategies under our proposed reward scheme. Our analysis considers an adversary that can corrupt nodes in a probabilistic manner. We show that if certain conditions are met, all nodes are incentivized to participate in the protocol regardless of being selected as committee members.

In summary, we make the following contributions. In §3.3, we present a detailed cost model for nodes’ participation in Algorand. In §4, we model the Algorand protocol as a Bayesian game. In §5, we propose IRS, a novel incentive-compatible reward scheme to address the free-rider problem. In §5, we study equilibrium strategies under IRS and derive necessary conditions to ensure nodes participation. In §6, we present detailed implementation requirements for real-world deployment of IRS. Finally, in §7, we quantitatively analyze our proposed reward scheme and simulate the protocol to measure its overheads in terms of computation, storage, and networking.

2 ALGORAND PROTOCOL

The Algorand protocol maintains a permissionless blockchain. In Algorand, adding a new block to the blockchain requires multiple steps. Algorithm (1) provides high-level pseudocode of the Algorand protocol (please refer to Appendix C in [41] for an in-depth overview of the protocol). At each step, all nodes wait for the messages from the previous step for a fixed amount of time. Each node then validates and propagates received messages to its neighbors. The protocol can terminate at specific steps (i.e., $k > 4$ where $k \not\equiv 1 \pmod{3}$) if a termination condition is met (i.e., enough votes are received or the final step, K_{\max} , is reached). At each non-terminal step, a random committee of self-selected nodes is formed. Committee members generate and propagate a message according to the protocol. The message is a block proposal if $k = 1$, and its a vote on a proposed block if $k > 1$.

Cryptographic sortition. In Algorand, nodes are assumed to have access to a unique signature scheme (e.g., [43]). As shown in Algorithm (1), at step k , given a publicly known random seed, Q , each node privately computes a hashlen-bit-long random string $x_k = H(\text{SIG}(k, Q))$ by digitally signing (k, Q) and then hashing it using a random oracle H [30]. The string x_k is interpreted as a binary expansion of a number between 0 and 1, denoted by $.x_k = x_k / 2^{\text{hashlen}}$. If this number is less than a known threshold, p_c , then the node is a member of the committee at step k . The threshold is set such that the expected size of the committee is τ (i.e., $p_c = \tau/n$, where n is the number of nodes in the system). $.x_k$ is also used to represent the priority of the node; the smaller $.x_k$ the higher the priority of the node will be. For nodes in the committee,

Algorithm 1: High-level pseudocode for Algorand

```

for  $k = 1, \dots, K_{\max}$  do
  if  $k > 1$  then Validate and gossip step- $(k - 1)$  messages;
  if  $k > 4$  and  $k \not\equiv 1 \pmod{3}$  then Exit if termination
    condition for step  $k$  is met;
   $(x_k, \sigma_k) \leftarrow \text{Sortition}(k)$ ;
  if  $.x_k \leq p_c$  then
    message  $\leftarrow$  Generate a message;
    Propagate message and  $\sigma_k$ ;

procedure  $\text{Sortition}(k)$  begin
   $\sigma \leftarrow \text{SIG}(k, Q)$ ;
   $x \leftarrow H(\sigma)$ ;
  return  $(x, \sigma)$ ;

```

$\sigma_k = \text{SIG}(k, Q)$ is the committee credential. Committee members propagate their credential alongside their generated message.

Gossip network and protocol. In Algorand, each node is provided with an *address-book* file containing the IP address and the port number of other nodes. Nodes form a gossip network by selecting a subset of n_{rp} random peers to gossip messages to. The parameter n_{rp} depends on the number of nodes, and it is set such that the gossip network is strongly connected. Messages are disseminated on the gossip network using a gossip protocol. The message dissemination is initiated by committee members at each step. Each committee member propagates their generated message to their randomly selected peers. Those peers then forward the message to their own peers. And this process continues until the message is received by all the nodes in the network. To avoid forwarding loops, nodes do not propagate the same message twice.

3 PRELIMINARIES

3.1 Adversary Model

The committee is guaranteed to reach Byzantine agreement [19, 22, 48] in the presence of an adversary that can corrupt nodes and control their actions. The Algorand protocol is resilient to such adversary as long as it cannot corrupt more than $1/3$ of the nodes. This is achieved by setting the expected size of the committee, τ , such that, with high probability, at least $2/3$ of the committee members are non-Byzantine nodes. In this paper, however, we consider a slightly different adversary model. In particular, we assume that the adversary corrupts each node with a fixed probability $p_b < 1/3$. Under our probabilistic adversary model, it is possible for the adversary to corrupt more than $1/3$ of nodes ex post. However, we show in §5 that for large systems, under our adversary model, non-Byzantine nodes still constitute more than $2/3$ of the committee with high probability. Consequently, Algorand protocol is guaranteed to reach Byzantine agreement with high probability.

Non-Byzantine supermajority. Since Algorand is a permissionless blockchain, the adversary can easily introduce as many new nodes as it wishes. Therefore, instead of assuming that the system has at least a $2/3$ majority of non-Byzantine nodes, it is often more meaningful to assume that at least $2/3$ of the cryptocurrency assets are controlled by non-Byzantine nodes. In other words, instead of

assuming that the adversary can corrupt up to 1/3 of the nodes, it is often assumed that the adversary can control up to 1/3 of the assets in the blockchain. Algorand achieves this by assigning *sub-nodes* to each node in proportion to the balance of its account. The cryptographic sortition algorithm then randomly selects each sub-node as a committee member. In this paper, we present our analysis under the simpler assumption that each node has a single sub-node. We then show how to modify the Algorand protocol and our analysis to consider the more realistic assumption that each node controls multiple sub-nodes.

3.2 Network Model

In this paper, we assume that the gossip network is strongly synchronous. This is a widely adopted network assumption [6, 7, 10, 29] which states that all messages propagated initially by non-Byzantine nodes are received by all other non-Byzantine nodes within a known time period. We further assume that the network remains strongly synchronous if a majority of nodes run the gossip protocol. This means that for large systems, the adversary cannot launch an Eclipse attack [31] with high probability.

3.3 Cost Model

Nodes running the Algorand protocol incur processing and communication costs at each step of the protocol. These costs are measurable in quantitative terms (e.g., energy consumption) and can be expressed in monetary values (e.g., cryptocurrency or Dollar).

We denote the total cost incurred by any node i at step k by $C_i(k)$. This cost has two components: (a) baseline costs, $C_i^b(k)$, and (b) committee costs, $C^c(k)$. We model $C_i(k)$ as follows.

$$C_i(k) = C_i^b(k) + C^c(k) \times \mathbb{1}(x_{i,k} < p_c)^4. \quad (1)$$

$C_i^b(k)$ represents the baseline costs that do not depend on whether node i is selected as a committee member at step k . Examples include costs of running cryptographic sortition and propagating messages. $C^c(k)$ represents the committee costs incurred by a node when it is selected as a committee member at step k . Examples include costs of generating blocks and votes.

4 THE ALGORAND GAME

To study nodes' incentives, we model the participation of nodes in the Algorand protocol as a Bayesian game. We formally define the Algorand game and describe nodes' strategies and utilities. We then discuss solution concepts for our proposed game.

4.1 Game Model

We model the Algorand protocol as a Bayesian game. A Bayesian game consists of a set of agents. Each agent has a type and a set of available actions. Agents do not know their types before the start of the game. They, however, know a common prior probability distribution over types. At the beginning of the game, each agent privately observes its own type. Agents then simultaneously take their actions without knowing each others' types. Finally, agents receive a real-valued utility (a.k.a. payoff) given their joint types and actions. A Bayesian game is formally defined as follows.

DEFINITION 1 (BAYESIAN GAME [53]). A Bayesian game is represented by a tuple (N, A, Θ, p, u) where:

- $N = \{1, \dots, n\}$ is a set of agents;
- $A = A_1 \times \dots \times A_n$, where A_i is a set of actions available to agent i ;
- $\Theta = \Theta_1 \times \dots \times \Theta_n$, where Θ_i is the type space of agent i ;
- $P : \Theta \mapsto [0, 1]$ is a common prior over types; and
- $u = (u_1, \dots, u_n)$, where $u_i : A \times \Theta \mapsto \mathbb{R}$ is the utility for agent i .

Agents and actions. In our setting, agents represent Algorand nodes. Agents are assumed to be rational in the sense that they selfishly choose an action to maximize their utility function. We consider three actions: (i) cooperate, C , (ii) defect, D , (iii) misbehave, M . A cooperative agent fully runs the Algorand protocol's code and consequently incurs all the processing and communication costs associated with it. A defective agent does not run any code (e.g., logs off from the system) and incurs no costs. A misbehaving agent runs a malicious code to sabotage the system. Note that the malicious code can imitate the behaviour of cooperating or defecting. Non-Byzantine agents that are not corrupted by the adversary do not misbehave. They only choose between cooperating and defecting. They cooperate if and only if their expected rewards exceed their expected costs. Byzantine agents, however, always misbehave (i.e., they run the adversary's malicious code).

Formally, we denote the action of agent i by $a_i \in A_i = \{C, D, M\}$. A vector of actions $a = (a_1, \dots, a_n) \in A$ is called an action profile. An action profile a can be written as (a_i, a_{-i}) , where a_{-i} is an action profile without agent i 's action⁵.

Types. Type of agent i is defined to be $\theta_i = (\theta_{i,0}, \theta_{i,1}, \dots, \theta_{i,K_{\max}})$ where $\theta_{i,0} \in \{0, 1\}$ indicates if agent i is corrupted by the adversary (1 if i is Byzantine and zero otherwise), and $\theta_{i,k}$ is a number between 0 and 1 represented in binary by the hash result of the sortition algorithm run by agent i at step $k = 1, \dots, K_{\max}$ ($\theta_{i,k} = x_{i,k}$, where $x_{i,k}$ is returned by $\text{Sortition}_i(k)$). The type space of agent i is denoted as Θ_i . At the beginning of the Algorand game, each agent observes whether it is corrupted by the adversary. Agents also receive a random seed for the cryptographic sortition algorithm. Given the random seed, agents can run the sortition algorithm for all steps to know what exactly their type vector is.

Prior probabilities. The probability that agent i is corrupted by the adversary is $\mathbb{P}(\theta_{i,0} = 1) = p_b$. For $k = 1, \dots, K_{\max}$, $\theta_{i,k}$'s are drawn independently from a uniform distribution between 0 and 1. Therefore, the probability that agent i is selected as a committee member is $\mathbb{P}(\theta_{i,k} \leq p_c) = p_c$. The adversary corrupts agents independently. Agents also are selected as committee members at any given step independently.

Utility functions. For agent i , the utility function, u_i , maps action profiles, $a = (a_1, \dots, a_n)$, and type vectors, $\theta = (\theta_1, \dots, \theta_n)$, to real-valued payoffs. If $\theta_{i,0} = 1$, then we assume that $u_i(a, \theta)$ is $-B$ if $a_i \in \{C, D\}$ and 0 if $a_i = M$ where B is a large real number. Under this assumption, regardless of their type, Byzantine agents always prefer M to C and D . To model non-Byzantine agents preferences, we assume that if $\theta_{i,0} = 0$, then $u_i(a, \theta) = -B$ if $a_i = M$. For $a_i \in \{C, D\}$, the utility of non-Byzantine agents is equal to the rewards they receive minus the costs they incur. We analyze utility functions in more details in §5.

⁴ $\mathbb{1}(\cdot)$ is the indicator function which returns 1 if the condition is true, and 0 otherwise.

⁵Throughout the paper, we use $-i$ to denote all agents except agent i .

4.2 Strategies and Equilibria

A strategy defines a description of how a game would be played in every contingency. In a Bayesian game, a strategy prescribes a distribution over actions for every type that an agent could have. Let $\Delta(A_i)$ be the set of all probability distributions over A_i . For agent i , a strategy $s_i : \Theta_i \mapsto \Delta(A_i)$ is a mapping from agent i 's types to distributions over agent i 's actions. The set of all strategies for agent i is denoted by S_i . By $s_i(a_i | \theta_i)$, we indicate the probability that agent i takes action a_i under s_i given that agent i 's type is θ_i . Similar to action profiles, a strategy profile $s = (s_1, \dots, s_n) \in S$ is a vector of strategies, where $S = S_1 \times \dots \times S_n$ is a set of all possible strategy profiles.

Expected utilities. In Bayesian games, there are two main sources of uncertainty: (i) types and (ii) actions. Types are drawn from the prior probability distribution, P , and actions are taken based on agents' strategies. To capture both sources of uncertainty, the ex ante expected utility of agent i is modeled as follows.

$$EU_i(s) = \mathbb{E}_{a,\theta}[u_i(a, \theta)] = p_b \mathbb{E}_{a,\theta}[u_i(a, \theta) | \theta_{i,0} = 1] + (1 - p_b) \mathbb{E}_{a,\theta}[u_i(a, \theta) | \theta_{i,0} = 0].$$

The expectation is taken with respect to θ and $a \sim s(\cdot | \theta)$. This formula models the expected utility of agent i before the start of the game and before the agent observes its type.

Given the defined expected utility model, we can define the set of agent i 's best responses to strategy profile s_{-i} as:

$$BR_i(s_{-i}) = \operatorname{argmax}_{s_i \in S_i} EU_i((s_i, s_{-i})).$$

Intuitively, a best response is a strategy which provides the highest expected utility given the strategy of others. Note that there may be more than one strategy that maximizes agent i 's expected utility for a given s_{-i} .

Bayesian Nash equilibrium. As discussed before, a strategy is a full contingency plan. Agents **simultaneously** choose their strategies **before** the start of the game and do not change their adopted strategies during the game. Once the game starts, each agent observes its type and acts as prescribed by its strategy. Agents strategies form a Bayesian Nash equilibrium (BNE) when the strategy of each agent is a best response to the strategies adopted by other agents. Formally, a strategy profile s^* is a BNE if and only if $s_i^* \in BR_i(s_{-i}^*)$, for all i . Informally, in a BNE, agent i does not have any incentive to unilaterally change its strategy from s_i^* if it knows that other agents have fixed their strategies to s_{-i}^* .

5 INCENTIVE ANALYSIS IN ALGORAND

In this section, we formulate agents' utilities and study their BNE strategies. We first consider Algorand's original reward scheme. We show that under this reward scheme, cooperation is not a BNE strategy. We then propose a novel reward scheme and show that under certain conditions, our proposed reward scheme incentivizes all non-Byzantine agents to cooperate regardless of their type.

5.1 Algorand's Original Reward Scheme

Algorand's original reward scheme is called Participation Rewards [28]. Under this reward scheme, *The Algorand Foundation* distributes a fixed amount of cryptocurrency assets as a reward among all agents. Agents are assigned sub-nodes in proportion to the balance

of their accounts. Under Algorand's original reward scheme, the fixed reward, R , is distributed equally among all sub-nodes. If agent i is assigned w_i sub-nodes, then agent i 's reward, R_i , is equal to $R w_i / W$, where W is the total number of sub-nodes in the system. The first advantage of this reward scheme is its simplicity: it is easy to implement, and it is easy to explain to agents how they are rewarded. The most important advantage of Algorand's original reward scheme is that it provides proportional rewards.

DEFINITION 2. *Proportional rewards.* Let R_i denote the expected reward of agent i . A reward scheme provides proportional rewards if for any agent i and $j \in N$, $R_i / R_j = w_i / w_j$.

The proportional rewards property ensures that the expected fraction of assets controlled by the adversary does not increase by the action of the reward scheme. Although Algorand's original reward scheme is simple and provides proportional rewards, it suffers from a key drawback: it fails to prevent free riding. Agents do not have any incentive to cooperate as they receive their rewards irrespective of their cooperation.

THEOREM 3. *Let s^* be a strategy profile where for each agent i , if $\theta_{i,0} = 0$, then $s_i^*(C | \theta_i) = 1$, and otherwise, $s_i^*(M | \theta_i) = 1$. Under Algorand's original reward scheme, s^* is not a BNE.*

PROOF. Suppose agent i is a non-Byzantine agent (i.e., $\theta_{i,0} = 0$). If agent i defects, it receives its rewards without incurring any costs. Formally, $u_i(a, \theta) = R/n$ for all $\theta \in \Theta$ if $a_i = D^b$. Let s'_i be a strategy that chooses D regardless of the type (i.e., $s'_i(D | \theta_i) = 1$ for all $\theta_i \in \Theta_i$). It can be easily shown that $EU_i((s'_i, s_{-i}^*)) = (1 - p_b) R/n$. If a non-Byzantine agent i cooperates, it receives its reward but incurs some strictly positive costs. This means that $u_i(a, \theta) < R/n$ if $a_i = C$. Let $a^*(\theta)$ be an action profile where a_i^* is C if $\theta_{i,0} = 0$ and M otherwise. The expected utility of agent i for s^* is:

$$EU_i(s^*) = (1 - p_b) \mathbb{E}_{\theta}[u_i(a^*(\theta), \theta) | \theta_{i,0} = 0] < (1 - p_b) R/n = EU_i((s'_i, s_{-i}^*)).$$

This implies that s_i^* is not a best response to s_{-i}^* . Therefore, s^* is not a BNE under Algorand's original reward scheme. \square

5.2 Incentive-compatible Reward Scheme (IRS)

To address the free-rider problem, we propose IRS, a novel incentive-compatible reward scheme for Algorand. Under IRS, the Algorand Foundation distributes rewards among agents based on their cooperation. The cooperation of committee members can be easily tracked as their messages are guaranteed to reach all other agents. However, tracking the cooperation of the agents that are not selected as a committee member is challenging as they do not initiate any messages. To address this challenge, our proposed reward scheme requires committee members at step $k = 2, \dots, K_{\max}$ to include with their vote the identities of the agents from which they have received a valid message at step $k - 1$.

Let $R^c(k)$ and $R^b(k)$ denote a fixed baseline reward and a fixed committee reward at step k , respectively. Under IRS, a cooperating agent i receives a committee reward of $R^c(k)$ if it is selected as

⁶We present the proof for the case where all agents are assigned a single sub-node. Our proof easily extends to the case where agents are assigned different number of sub-nodes.

a committee member at step k . Additionally, if agent i 's identity is included in the vote generated by agent $j \in N_i$ at step k , then agent i receives a baseline reward of $R^b(k)/n_{rs}$. We assume that Byzantine agents do not include the identity of non-Byzantine agents in their vote when they are selected as committee members. In other words, non-Byzantine agents do not receive any baseline reward for propagating messages to their Byzantine peers. We make this assumption to calculate a lower bound on the expected utility of non-Byzantine agents. If Byzantine agents include the identity of non-Byzantine agents, the expected utility of non-Byzantine agents can only increase. Given this assumption, the total reward of a cooperating agent i at step k given an action profile a and a type vector θ can be formulated as:

$$R_i(a, \theta, k) = (R^b(k)/n_{rs}) \sum_{j \in N_i} \mathbb{1}(\theta_{j,k} \leq p_c, a_j = C) \quad (2)$$

$$+ R^c(k) \mathbb{1}(\theta_{i,k} \leq p_c). \quad (3)$$

To show that our reward scheme prevents the free-rider problem, we first prove that any non-Byzantine agent cannot unilaterally change the outcome of each step by defecting. Given that, we formulate the expected utility of each non-Byzantine agent assuming that all other non-Byzantine agents cooperate regardless of their type. We then prove that if certain conditions are met, cooperation is a best response for a non-Byzantine agent when other non-Byzantine agents cooperate. This then shows that cooperation is a BNE strategy for non-Byzantine agents under IRS.

LEMMA 4. Let NR_k and NB_k be random variables indicating the number of non-Byzantine and the number of Byzantine agents selected as committee members at any step k , respectively. Let $\mu_r = \mathbb{E}(NR_k)$, $\mu_b = \mathbb{E}(NB_k)$, $\mu = \mu_b + \mu_r/2$, $\delta_r = 1 - T/(1 - p_b)$, and $\delta = 2 \times T/(1 + p_b) - 1$. Then we have:

- $\mathbb{P}(NR_k \leq T \tau) \leq e^{-\mu_r \delta_r^2/2}$, and
- $\mathbb{P}(NB_k + NR_k/2 \geq T \tau) \leq e^{-\mu \delta^2/(2+\delta)}$.

PROOF. Let $X_{i,k}$ be a random variable that takes value 1 if agent i is selected as a committee member and agent i is not corrupted by the adversary, and takes value 0 otherwise. Similarly, let $Y_{i,k}$ be a random variable that takes value 1 if agent i is selected as a committee member and agent i is corrupted by the adversary, and takes 0 otherwise. We can write $NR_k = \sum_i X_{i,k}$, and $NB_k = \sum_i Y_{i,k}$. We have $\mathbb{E}(X_{i,k}) = (1 - p_b) p_c$, and $\mathbb{E}(Y_{i,k}) = p_b p_c$. Therefore, $\mu_r = \mathbb{E}(NR_k) = \sum_i \mathbb{E}(X_{i,k}) = (1 - p_b) \tau$, and $\mu_b = \mathbb{E}(NB_k) = \sum_i \mathbb{E}(Y_{i,k}) = p_b \tau$. Since $\delta_r \geq 0$, according to Chernoff bound, the following inequality holds for any k .

$$\mathbb{P}(NR_k \leq T \tau) = \mathbb{P}(NR_k \leq (1 - \delta_r) (1 - p_b) \tau)$$

$$= \mathbb{P}(NR_k \leq (1 - \delta_r) \mu_r) \leq e^{-\mu_r \delta_r^2/2}.$$

Similarly, since $\delta \geq 0$, we have the following inequality for any k .

$$\mathbb{P}(NB_k + NR_k/2 \geq T \tau) = \mathbb{P}(NB_k + NR_k/2 \geq (1 + \delta) (1 + p_b) \tau/2)$$

$$= \mathbb{P}(NB_k + NR_k/2 \geq (1 + \delta) \mu)$$

$$\leq e^{-\mu \delta^2/(2+\delta)}.$$

□

We assume that p_b , τ , and T are set such that $NR_k > T \tau$ and $NB_k + NR_k/2 < T \tau$ are true with overwhelming probability. For

example, using Lemma (4), with $p_b = 0.2$, $\tau = 4000$ and $T = 0.7$, the probability that $NR_k \leq T \tau$ is less than 10^{-10} and the probability that $NB_k + NR_k/2 \geq T \tau$ is less than 10^{-13} . The two inequalities imply that more than 2/3 of the selected committee members at each step are non-Byzantine agents with overwhelming probability.

PROPOSITION 5. Consider any agent $i \in N$. Suppose that the strategy profile of all agents except agent i is s_{-i}^* where for each agent $j \in N \setminus i$, if $\theta_{j,0} = 0$, then $s_j^*(C | \theta_j) = 1$, and $s_j^*(M | \theta_j) = 1$ otherwise. In a large system (i.e., $n \rightarrow \infty$), the safety and liveness guarantees of the Algorand protocol are met with high probability regardless of the strategy of agent i .

PROOF. Assume that p_b , τ , and T are set such that $NR_k > T \tau$ and $NB_k + NR_k/2 < T \tau$ with overwhelming probability. Consider a new system consisting of all agents except agent i . For large systems, it is easy to modify Lemma (4) to show that the two inequalities still hold for the new system with the same p_b , τ , and T . Moreover, the Chernoff bound can be applied to show that with $p_b \leq 1/3$, the network is synchronous with high probability regardless of agent i 's cooperation. Given that the network is synchronous, and the two inequalities hold with high probability, Theorem 1 from [15] can be applied to the new system to guarantee safety and liveness. □

We next formulate the expected utility of each agent. The Algorand protocol implements a randomized algorithm. As discussed in §2, the algorithm runs in multiple steps. We use $K(a, \theta)$ to denote the total number of steps it takes the algorithm to complete as a function of agents' types and their actions. Given a and θ , the total utility of agent i with $\theta_{i,0} = 0$ can be formulated as:

$$u_i(a, \theta) = \sum_{k=1}^{K(a, \theta)} u_i(a, \theta, k),$$

where $u_i(a, \theta, k)$ is $R_i(a, \theta, k) - C_i(a, \theta, k)^7$ if $a_i = C$ and 0 otherwise.

LEMMA 6. Consider any agent $i \in N$. Suppose that the strategy profile of all agents except agent i is s_{-i}^* where for each agent $j \in N \setminus i$, if $\theta_{j,0} = 0$, then $s_j^*(C | \theta_j) = 1$, and $s_j^*(M | \theta_j) = 1$ otherwise. Suppose further that agent i adopts strategy s_i which plays M if $\theta_{i,0} = 1$ and plays C with probability s_c and D with probability $1 - s_c$ if $\theta_{i,0} = 0$. Let $a^*(\theta)$ be an action profile where a_i^* is C if $\theta_{i,0} = 0$ and M otherwise. Define $K(\theta) = K(a^*(\theta), \theta)$. The expected utility of agent i is:

$$EU_i((s_i, s_{-i}^*)) = s_c (1 - p_b) \sum_{\ell=1}^{K_{\max}} \mathbb{P}(K(\theta) = \ell) \sum_{k=1}^{\ell} u(k),$$

where $u(k) = R^b(k) p_c (1 - p_b) - C^b(k) + (R^c(k) - C^c(k)) p_c$.

The proof is presented in Appendix B of [41]. Next, we prove that s^* is a BNE under certain conditions. These conditions are derived simply by requiring that the expected rewards should outweigh the expected costs.

THEOREM 7. Let s^* be a strategy profile where for each agent i , if $\theta_{i,0} = 0$, then $s_i^*(C | \theta_i) = 1$, and otherwise, $s_i^*(M | \theta_i) = 1$. Given p_c , p_b , $C^b(k)$, and $C^c(k)$ for $k = 1, \dots, K_{\max}$, s^* is a BNE under IRS if:

$$R^b(k)(1 - p_b) + R^c(k) \geq C^b(k)/p_c + C^c(k), \quad \forall k = 1, \dots, K_{\max}.$$

⁷In Equation (1), C_i depends on θ through $.x_{i,k}$, and it is formulated assuming that agent i cooperates.

PROOF. To prove that s^* is a BNE, it suffices to show that s_i^* is a best response to s_{-i}^* for all agent $i \in N$. Recall that s_i^* is a best response to s_{-i}^* when $EU_i((s_i^*, s_{-i}^*)) \geq EU_i((s_i, s_{-i}^*))$ for all $s_i \in S_i$. If $R^b(k)(1-p_b) + R^c(k) \geq C^b(k)/p_c + C^c(k)$ for all $k = 1, \dots, K_{\max}$, then $u(k)$ in Lemma (6) is greater than or equal to zero for all $k = 1, \dots, K_{\max}$. Consequently, $EU_i((s_i, s_{-i}^*))$ is maximized when $s_c = 1$. In this case, we have $s = s^*$ which means s^* is a BNE. \square

6 IMPLEMENTATION DETAILS

6.1 Gossip Protocol in IRS

To implement IRS, we make three main modifications to Algorand’s default gossip protocol. First, we require the randomness of the peer-selection mechanism to be verifiable (e.g., through verifiable pseudo-random peer selection [40]). This requirement prevents the adversary from gaining unauthorized awards. Byzantine nodes cannot be rewarded for propagating messages to nodes that are not among their randomly selected peers. Byzantine committee members also cannot refer other Byzantine nodes that are not randomly connected to them.

Second, we require nodes to disable selective propagation. Selective propagation is an optimization technique that prevents nodes from propagating low-priority block proposals [15, 29]. Although this technique reduces network congestion, it prevents low-priority block proposals from reaching all nodes in the gossip network. This in turn prevents the Algorand Foundation from tracking cooperation of some committee members at step 1. One optimization that could be implemented to replace selective propagation is to only send the committee member’s credential for the low-priority block proposals without sending the entire block proposal.

Third, we require nodes to track the identity of all nodes from which they have received a valid message, even if the message is a duplicate message. For example, suppose that agent i receives message m at step k first from agent j and later from agent j' . Agent i propagates message m to its peers only the first time it receives it from agent j . However, it saves the identity of both agents j and j' as propagators for m at step k . If agent i becomes a committee member at step $k + 1$, it includes the identity of both agents j and j' in its generated vote. This allows the Algorand Foundation to not only track the cooperation of the committee members but also the cooperation of their peers. Since committee members are selected randomly, their cooperating peers can be considered random samples of non-voting agents that cooperate.

6.2 Consideration of Assets in IRS

Algorand assigns sub-nodes to each node in proportion to the balance of its account. For simplicity, in our analysis so far, we have considered a single sub-node per node. We now discuss the necessary changes needed to allow multiple sub-nodes per node.

Cryptographic sortition. The sortition algorithm can be easily extended to consider node with more than one sub-node. Suppose that node i has w_i sub-nodes. A simple way to extend sortition is for node i to run the sortition algorithm (see Algorithm (1)) on each of its w_i sub-nodes. Although simple, this method is computationally expensive. An alternative method, proposed in [15], is to use the inverse transform sampling. In this method, the interval of $[0, 1)$ is

partitioned into consecutive intervals of the following form.

$$I_{w_i, p}^M = [B(M; w_i, p), B(M + 1; w_i, p)), \forall M \in \{0, 1, \dots, w_i - 1\}.$$

If x_i falls in the $I_{w_i, p}^M$ interval, node i has M selected sub-nodes. **Adversary model.** So far, we have assumed that the adversary corrupts each node with probability $p_b < 1/3$. We can extend this assumption by allowing the adversary to corrupt each sub-node with probability p_b . Suppose that the total number of sub-nodes is $W = \sum_{i \in N} w_i$. We define $W_{-i} = \sum_{j \in N \setminus i} w_j$ for all agents i . For our results to hold under the new adversary model, we require W_{-i} to go to infinity as n goes to infinity. We further require that $w_i/W < (1 - 3 \times p_b)$ for all nodes i . This requirement ensures that if any single node is removed from the set of nodes, in expectation, the adversary does not control more than one-third of the remaining sub-nodes. For example, if $p_b = 0.3$, then no single agent should have greater than or equal to one-tenth of all sub-nodes. Under these assumptions, it can be easily verified that our results in §5 hold for the new adversary model because non-Byzantine sub-nodes still constitute more than $2/3$ of the committee with high probability.

Gossip network and protocol. When each agent has more than one sub-node, instead of forming the random gossip network among nodes, we require the gossip network to be constructed among sub-nodes. This means that nodes need to select a subset of random sub-nodes as peers for each of their sub-nodes. This can be achieved by including indexes of sub-nodes as an extra input for the verifiable pseudo-random peer selection procedure [40]. When a sub-node propagates a message to another sub-node, it is required to include the indexes of both the sender and the receiver sub-nodes in the message. When a node receives a message, it will verify if the sender sub-node and the receiver sub-node are indeed connected.

IRS rewards. The simplest way to extend IRS to consider multiple sub-nodes per node is to distribute rewards on a per-sub-node basis. In other words, IRS rewards each sub-node independently. In this way, the expected total reward of a node is the sum of the expected rewards of its sub-nodes. The expected total cost depends on the implementation of the sortition process and the gossip protocol. With the simple extensions presented above, the total expected cost is the sum of the expected costs of its sub-nodes.

7 EXPERIMENTS

In this section, we estimate the costs of running an Algorand node for two real-world deployment scenarios. Based on the costs, we identify the baseline reward, R^b , and committee reward, R^c , for the two scenarios to guarantee participation as a BNE strategy. We also simulate the extended gossip protocol outlined in §6 for IRS and measure its overheads in terms of computation, storage and networking.

7.1 Estimated Costs and Rewards

Compared to PoW blockchains, Algorand nodes require much less computation power to participate in the consensus protocol. Any tentative participant can run an Algorand node with commodity machines. We consider two deployment scenarios. First, the participant runs an Algorand node on a virtual machine from a cloud service provider such as Amazon Web Services (AWS) [5]. Second,

Comp. & Storage Cost	72.54 \$/mon	8.2×10^{-5} Algo/s
Networking Cost	0.09 \$/GB	0.2647 Algo/GB

Table 1: Costs of an AWS EC2 instance

Step	C^b	R^b	C^c	R^c
1	1.77×10^{-3}	18603	8.47×10^{-5}	8.5×10^{-5}
2	6.4×10^{-3}	133891	2.12×10^{-6}	2.2×10^{-6}
3	3.26×10^{-3}	20408	2.12×10^{-6}	2.2×10^{-6}
4 to K_{\max}	0.011	66823	2.12×10^{-6}	2.2×10^{-6}

Table 2: Estimated costs and rewards (Algo/block) on AWS.

the participant runs the Algorand protocol on a personal computer. We estimate the costs in Algo⁸ (1 Algo = 0.34 USD).

Cloud-hosted Algorand node. Consider an AWS EC2 instance with 4 vCPU, 8GB memory, and 256GB SSD storage. The costs associated with renting such instance are listed in Table 1. The average size of each block in Algorand ledger is about 40KB [4]. We estimate the size of a vote message to be about 1KB based on the official implementation of Algorand in Go programming language [26]. On average, it takes about 5 seconds for the Algorand protocol to commit a block [3]. We assume that each round takes 5 steps, which means that each step takes about 1 second. We further assume that each Algorand node has 8 peers in the gossip network. This is the default value in Algorand [29]. According to current specification of Algorand [25], $p_b = 1/5$ and $\tau = 20, 2990, 1500, 5000, 5000, \dots, 5000$ for step 1 to K_{\max} . As of July 24th, 2022, according to the AlgoExplorer [36], there are 550 nodes with a total of more than 25 billion sub-nodes actively participating in running the Algorand protocol.

We estimate the baseline cost at each step as computation and storage costs plus communication costs. Computation and storage costs are listed in Table 1 and are fixed across steps. The communication cost per step is calculated as a linear function of the number of neighbours per node and the per GB networking cost listed in Table 1. Similarly, the committee cost at each step is estimated as the sum of computation and storage costs and communication costs. The computation and storage costs are covered in our estimated baseline cost (the same rented AWS EC2 instance executes both baseline and committee related portions of the Algorand protocol). The communication cost per step is calculated similarly as before. Given the costs, we derive the baseline rewards and committee rewards at each step by ensuring that $R^b(k)(1 - p_b) \geq C^b(k)/p_c$ and $R^c(k) \geq C^c(k)$ according to Theorem (7). The estimated costs and derived IRS rewards per block are listed in Table 2.

Self-hosted Algorand node. Next, we consider the costs and associated IRS rewards when an Algorand node runs on a personal computer (PC). We consider the energy consumption of a typical PC to be about 200W [1]. We consider an average energy price of 0.0944 USD/kWh [2]. We further consider an average price of an Internet plan with unlimited traffic to be around 100 USD/month. Given these costs, the overall cost of running an Algorand node on a personal computer is about 0.00013 Algo/s. Similar to our cost estimation for the AWS scenario, the baseline cost covers the total computation and storage costs. Since there is no limit on the traffic,

⁸Algo is the unit of Algorand’s cryptocurrency.

	Step 1	Step 2	Step 3	Steps 4 to K_{\max}
R^b	1362	2627	815	815

Table 3: Estimated baseline rewards (Algo/block) on PC.

we can simply set $C^b = 0.00013$ Algo and $C^c = 0$ Algo. The baseline IRS rewards are listed in Table 3.

Comparison against Algorand’s original reward scheme. Algorand’s Participation Rewards only distributes about 20 to 30 Algos for each block. In comparison, our calculated rewards might seem prohibitively large for a practical implementation and deployment of IRS. We note that most blocks are generated within the first 5 steps. This means that for most blocks, no reward is given for steps 6 to K_{\max} . Second, our calculations are based on the assumption that each agent has a single sub-node, which leads to a worst-case analysis (unlike expected rewards, expected costs do not necessarily increase linearly as a function of sub-nodes). In the real-world deployment of Algorand, each agent has on average 40 million sub-nodes. Restricting the analysis to such agents could lead to more feasible rewards. Finally, in real-world settings, some agents are altruistic (i.e., they participate regardless of the rewards). Inclusion of such agents in the analysis would lead to lower rewards. Extending IRS to include altruistic agents would be an interesting future work, both theoretically and in practice.

7.2 Overhead of Gossip Protocol in IRS

The overhead of the extended gossip protocol in IRS consists of three main parts: (i) computational overhead of running verifiable random peer selection, (ii) networking overhead of disabling selective propagation, and (iii) storage overhead of storing identities of all valid message propagators.

Computation overhead. The verifiable random peer selection procedure randomly connects sub-nodes to each other. To implement this procedure, the inverse transform sampling technique in Algorand’s sortition procedure can be used (see §6.2). To measure the overhead of this implementation, we run the benchmark of the sortition function in the official implementation of Algorand [26] on a machine powered by AMD EPYC 7H12 processors. The average execution time to run the sortition once is about 0.2 ms. Algorand reconstructs the gossip network for each new block [29], so the overall computation overhead is then $0.0002 \times (N - 1) \approx 0.11$ seconds per block for $N = 550$. Given that each block is committed to the ledger every 5 seconds, this is an overhead of 2.2%. One way to reduce this overhead is to reduce the network reconstruction frequency (e.g., reconstruct the gossip network every 4 or 5 blocks). **Network and storage overhead.** To measure the network and the storage overheads, we implement a simulator of the gossip protocol in IRS. The simulator has 132 lines of Python code and is available at <https://anonymous.4open.science/r/IRS-Simulator/>. We run our simulator on real-world data extracted from AlgoExplorer [36]. We randomly connect each sub-node to 8 other sub-nodes. According to our simulation results, on average, each Algorand node is connected to 252 peers (each node has on average more than 45 million sub-nodes). Moreover, our results show that each node has to gossip the hashes of low-priority blocks 18 times (on average 20 selected sub-nodes as block proposers). Assuming that

the size of each block hash is 1KB, the network bandwidth overhead for each node is about $252 \times 18 \times 1\text{KB} = 4536\text{KB}$ per committed block. For storage overhead, each node needs to store the public keys of all valid message propagators. The size of the public key in Algorand is 32 bytes [26], so the storage overhead is about $252 \times 32\text{B} = 8.064\text{KB}$ per committed block. Regarding networking and storage, IRS has limited carbon footprint because networking and storage are energy disproportional (i.e., the energy consumption does not proportionally change as a function of network/storage utilization).

8 RELATED WORKS

Incentives in blockchains. Incentive compatibility of reward schemes is crucial to the security of blockchains. Eyal and Sirer [21] prove negative results with regard to incentive compatibility of Bitcoin. They model the competitive mining among the miners as a strategic game and propose a novel mining strategy called selfish mining. The authors examine the profitable threshold of selfish mining, which is the computation power needed to gain more revenue. The authors then propose a modification to the Bitcoin chain-selection protocol to increase the profitable threshold of selfish mining to 1/4. Selfish mining has been studied extensively ever since [33, 45, 47, 51, 57].

Incentives for information propagation. Most permissionless blockchains [29, 44, 56] rely on a peer-to-peer gossip protocol to propagate transactions. However, these protocols do not provide an incentive for the nodes in the blockchains to participate in propagating the transactions. In fact, Babaiou et al. [9] argue that these protocols incentivize nodes not to propagate transactions.

The authors in [9] propose a reward scheme to incentivize transaction propagation. They prove that under the proposed reward scheme, the strategy where all nodes propagate transactions and do not create fake identities (Sybil attacks) is a Nash equilibrium. In addition, the proposed scheme guarantees that most of the nodes in the network will be aware of the transaction. The additional rewards required to implement the scheme are a constant in expectation and the user only needs to send the transaction to a small number of nodes in the beginning. A main drawback of the model in [9] is that it is highly restricted. The model only considers networks in the form of a forest of d -ary directed trees with height H and assumes that each node has the same hashing power.

Ersoy et al. [20] propose another incentive mechanism for transaction propagation under a network model with minimal restrictions. The resulting reward scheme encourages propagating without creating fake identities.

Incentives in committee-based protocols. Amoussou-Guenou et al. [6, 7] analyze a simplified committee-based consensus protocol. They propose to make some committee members pivotal such that they have incentives to participate. In this case, a rational node can unilaterally determine the result of the consensus protocol. If a pivotal node does not follow the protocol, consensus is not reached and a penalty is applied. As a result, all pivotal nodes participating in the protocol as required by the protocol becomes a Nash equilibrium. Nevertheless, the solution is not practical in Algorand because the assumptions in [6, 7] do not hold. For example, the solution proposed by [7] assumes that all nodes are treated equally and have the same voting power, which is not the case in Algorand.

Also, the solution assumes that all the committee members are ordered by publicly known indexes, while the membership is private in Algorand until nodes publish it. In addition, each node communicates with all other nodes directly in the simplified protocol while Algorand adopts a gossip protocol to disseminate messages.

Fooladgar et al. [23] analyze the Algorand protocol as a static non-cooperative game. They demonstrate that all nodes participating the protocol is not a Nash equilibrium under Algorand’s original reward scheme. If all other nodes participate in the protocol, a node can free-ride to get its reward without paying the costs of running an Algorand node. To address the free-rider problem, the authors propose a new reward scheme where only participating nodes are rewarded. In this work, the analysis is limited to the rational-agent-only case. In addition, the proposed reward scheme only incentivizes committee members and fails to incentivize message propagation by all nodes. Under the proposed method, only nodes whose participation is necessary (for the network to remain synchronous) are incentivized to propagate messages.

We extend the prior work in two main aspects. First, our model preserves important features of the Algorand protocol, such as randomized membership selection and private membership information. Second, we propose a new reward scheme and prove that all nodes participating faithfully is a Bayesian Nash equilibrium under certain conditions even in the presence of an adversary.

9 CONCLUSION & FUTURE WORKS

In this paper, we model the Algorand protocol as a Bayesian game. We propose a reward-distribution scheme and derive necessary conditions such that all nodes participating in the protocol is a Bayesian Nash equilibrium. Our work highlights some open problems in designing incentive-compatible reward schemes for randomized, committee-based consensus protocols that depend on nodes for message propagation. In particular, like Algorand’s original reward scheme, our proposed reward scheme relies on a central authority to collect messages to identify the role of nodes and to verify the participation of nodes. A more decentralized reward scheme is preferred where nodes in Algorand network determine how the rewards are distributed without a central authority. Moreover, although our model captures some important features of Algorand, modeling the protocol as a sequential game is an interesting direction for future work. Finally, our analysis is based on the assumption that the network is synchronous. Extending the analysis to a setting where the network is not synchronous would be a possible future direction.

ACKNOWLEDGMENTS

This work was partially supported by the NSERC-RGPIN-2019-04936, CFI-JELF-38850, ORF-RI-38850, Ripple, and NSERC Discovery grants.

REFERENCES

- [1] 2018. Power Management Statistics. <https://www.it.northwestern.edu/hardware/eco/stats.html>.
- [2] 2021. State Electricity Profiles. <https://www.eia.gov/electricity/state/>.
- [3] Algorand. 2022. Algorand Developer Portal. <https://metrics.algorand.org/>.
- [4] Algoscan. 2022. Developer | Algoscan. <https://developer.algoscan.app/>.
- [5] Amazon. 2006. Cloud Computing Services - Amazon Web Services (AWS). <https://aws.amazon.com/>.

- [6] Yackolley Amoussou-Guenou, Bruno Biais, Maria Potop-Butucaru, and Sara Tucci-Piergiorgianni. 2020. Rational Behaviors in Committee-Based Blockchains. In *Proceedings of the 24th International Conference on Principles of Distributed Systems (OPODIS)*, 12:1–12:16.
- [7] Yackolley Amoussou-Guenou, Bruno Biais, Maria Potop-Butucaru, and Sara Tucci-Piergiorgianni. 2020. Rational vs Byzantine Players in Consensus-Based Blockchains. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 43–51.
- [8] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolić, Sharon Weed Cocco, and Jason Yellick. 2018. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. In *Proceedings of the 13th EuroSys Conference*, 1–15.
- [9] Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar. 2012. On Bitcoin and Red Balloons. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 56.
- [10] Ethan Buchman, Jae Kwon, and Zarko Milosevic. 2019. The Latest Gossip on BFT Consensus. *arXiv:1807.04938 [cs]* (2019).
- [11] Vitalik Buterin and Virgil Griffith. 2017. Casper the Friendly Finality Gadget. <https://doi.org/10.48550/ARXIV.1710.09437>
- [12] Vitalik Buterin, Diego Hernandez, Thor Kamphofner, Khiem Pham, Zhi Qiao, Danny Ryan, Julhyeok Sin, Ying Wang, and Yan X Zhang. 2020. Combining GHOST and Casper. <https://doi.org/10.48550/ARXIV.2003.03052>
- [13] Vitalik Buterin, Daniël Reijnders, Stefanos Leonardos, and Georgios Piliouras. 2019. Incentives in Ethereum’s Hybrid Casper Protocol. In *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 236–244. <https://doi.org/10.1109/BLOC.2019.8751241>
- [14] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine Fault Tolerance. In *Proceedings of the Symposium on Operating Systems Design and Implementation (OSDI)*, 173–186.
- [15] Jing Chen and Silvio Micali. 2019. Algorand: A Secure and Efficient Distributed Ledger. *Theoretical Computer Science* 777 (2019), 155–183.
- [16] Lin William Cong, Zhiguo He, and Jiasun Li. 2021. Decentralized Mining in Centralized Pools. *The Review of Financial Studies* 34, 3 (2021), 1191–1235.
- [17] Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. 2018. Ouroboros Praos: An Adaptively-Secure, Semi-synchronous Proof-of-Stake Blockchain. In *Advances in Cryptology – EUROCRYPT 2018*. Springer International Publishing, 66–98.
- [18] Diem. 2020. The Diem Association. <https://www.diem.com/en-us/>.
- [19] Danny Dolev and H. Raymond Strong. 1983. Authenticated algorithms for Byzantine agreement. *SIAM J. Comput.* 12, 4 (1983), 656–666.
- [20] Oguzhan Ersoy, Zhijie Ren, Zekeriya Erkin, and Reginald L. Legendijk. 2018. Transaction Propagation on Permissionless Blockchains: Incentive and Routing Mechanisms. In *Proceedings of the Crypto Valley Conference on Blockchain Technology (CVCBT)*, 20–30.
- [21] Ittay Eyal and Emin Gün Sirer. 2014. Majority Is Not Enough: Bitcoin Mining Is Vulnerable. In *Proceedings of the International Conference on Financial Cryptography and Data Security (FC)*, 436–454.
- [22] Michael J Fischer. 1983. The consensus problem in unreliable distributed systems (a brief survey). In *Proceedings of the International Conference on Fundamentals of Computation Theory (FCT)*, 127–140.
- [23] Mehdi Foadlgar, Mohammad Hossein Manshaei, Murtuza Jaidliwala, and Mohammad Ashiqur Rahman. 2020. On Incentive Compatible Role-Based Reward Distribution in Algorand. In *Proceedings of the 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 452–463.
- [24] Cambridge Centre for Alternative Finance. 2022. Cambridge Bitcoin Electricity Consumption Index (CBECI). <https://ccaf.io/cbeci/index/>.
- [25] Algorand Foundation. 2019. Algorand Blockchain Features Specification Version 1.0. https://github.com/algorandfoundation/specs/blob/5615adc36bad610c7f165fa2967f4ecfa75125f0/overview/Algorand_v1_spec-2.pdf.
- [26] Algorand Foundation. 2019. Algorand’s official implementation in Go. - GitHub. <https://github.com/algorand/go-algorand>.
- [27] Algorand Foundation. 2021. Algorand Pledges to be the Greenest Blockchain with a Carbon-Negative Network Now and in the Future. https://www.algorand.com/resources/algorand-announcements/carbon_negative_announcement.
- [28] Algorand Foundation. 2021. Introducing Governance: Earn rewards for your participation in the decision making. <https://algorand.foundation/governance>.
- [29] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. 2017. Algorand: Scaling Byzantine Agreements for Cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP)*, 51–68.
- [30] Oded Goldreich. 2007. *Foundations of cryptography: Volume 1, basic tools*. Cambridge university press.
- [31] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. 2015. Eclipse Attacks on Bitcoin’s Peer-to-Peer Network. In *Proceedings of the 24th USENIX Security Symposium (USENIX Security)*, 129–144.
- [32] Christine V Helliar, Louise Crawford, Laura Rocca, Claudio Teodori, and Monica Veneziani. 2020. Permissionless and permissioned blockchain diffusion. *International Journal of Information Management* 54 (2020), 102–136.
- [33] Charlie Hou, Mingxun Zhou, Yan Ji, Phil Daian, Florian Tramèr, Giulia Fanti, and Ari Juels. 2021. SquirRL: Automating Attack Analysis on Blockchain Incentive Mechanisms with Deep Reinforcement Learning. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*.
- [34] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. 2017. Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol. In *Advances in Cryptology – CRYPTO 2017*. Springer International Publishing, 357–388.
- [35] Sunny King and Scott Nadal. 2012. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. (2012).
- [36] Rand Labs. 2020. Algorand (ALGO) Blockchain Explorer. <https://algoexplorer.io/>.
- [37] Leslie Lamport. 2001. Paxos Made Simple. *ACM SIGACT News (Distributed Computing Column)* 32, 4 (Whole Number 121) (2001).
- [38] Leslie Lamport, Robert Shostak, and Marshall Pease. 1982. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems* (1982), 382–401.
- [39] Jacob Leshno and Philipp Strack. 2019. Bitcoin: An Impossibility Theorem for Proof-of-Work Based Protocols. *SSRN Electronic Journal* (2019).
- [40] Harry C. Li, Allen Clement, Edmund L. Wong, Jeff Napper, Indrajit Roy, Lorenzo Alvisi, and Michael Dahlin. 2006. BAR Gossip. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI)*, 191–204.
- [41] Maizi Liao, Wojciech Golab, and Seyed Majid Zahedi. 2023. IRS: An Incentive-compatible Reward Scheme for Algorand. [arXiv:2302.11178 \[cs.GT\]](https://arxiv.org/abs/2302.11178)
- [42] Ziyao Liu, Nguyen Cong Luong, Wenbo Wang, Dusit Niyato, Ping Wang, Ying-Chang Liang, and Dong In Kim. 2019. A survey on blockchain: A game theoretical perspective. *IEEE Access* 7 (2019), 47615–47643.
- [43] Silvio Micali, Michael Rabin, and Salil Vadhan. 1999. Verifiable random functions. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 120–130.
- [44] Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System.
- [45] Kartik Nayak, Srikanth Kumar, Andrew Miller, and Elaine Shi. 2016. Stubborn Mining: Generalizing Selfish Mining and Combining with an Eclipse Attack. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, 305–320.
- [46] Diego Ongaro and John Ousterhout. 2014. In Search of an Understandable Consensus Algorithm. In *Proceedings of the USENIX Conference on USENIX Annual Technical Conference (USENIX ATC)*, 305–320.
- [47] Rafael Pass and Elaine Shi. 2017. FruitChains: A Fair Blockchain. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 315–324.
- [48] Marshall Pease, Robert Shostak, and Leslie Lamport. 1980. Reaching agreement in the presence of faults. *Journal of the ACM (JACM)* 27, 2 (1980), 228–234.
- [49] QuantumMechanic. 2011. Proof of Stake Instead of Proof of Work. <https://bitcointalk.org/index.php?topic=27787.0>.
- [50] Team Rocket, Maofan Yin, Kevin Sekniqi, Robbert van Renesse, and Emin Gün Sirer. 2019. Scalable and Probabilistic Leaderless BFT Consensus through Metastability. <https://doi.org/10.48550/ARXIV.1906.08936>
- [51] Ayelet Sapirshstein, Yonatan Sompolskiy, and Aviv Zohar. 2016. Optimal Selfish Mining Strategies in Bitcoin. In *Proceedings of the International Conference on Financial Cryptography and Data Security (FC)*, 515–532.
- [52] Okke Schrijvers, Joseph Bonneau, Dan Boneh, and Tim Roughgarden. 2016. Incentive Compatibility of Bitcoin Mining Pool Reward Functions. In *Proceedings of the International Conference on Financial Cryptography and Data Security (FC)*, 477–498.
- [53] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.
- [54] Yonatan Sompolskiy and Aviv Zohar. 2015. Secure High-Rate Transaction Processing in Bitcoin. In *Financial Cryptography and Data Security*. Springer Berlin Heidelberg, Berlin, Heidelberg, 507–527.
- [55] Alistair Stewart and Eleftherios Kokoris-Kogias. 2020. GRANDPA: a Byzantine Finality Gadget. <https://doi.org/10.48550/ARXIV.2007.01560>
- [56] Gavin Wood et al. 2014. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper* 151, 2014 (2014), 1–32.
- [57] Shiquan Zhang, Kaiwen Zhang, and Bettina Kemme. 2020. A Simulation-Based Analysis of Multiplayer Selfish Mining. In *Proceedings of the IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 1–5.