

# Nonoutsourcable Scratch-Off Puzzles to Discourage Bitcoin Mining Coalitions

Andrew Miller

Ahmed Kosba  
University of Maryland  
{amiller, akosba, jkatz}@cs.umd.edu

Jonathan Katz

Elaine Shi  
Cornell University  
runting@gmail.com

## ABSTRACT

An implicit goal of Bitcoin’s reward structure is to diffuse network influence over a diverse, decentralized population of individual participants. Indeed, Bitcoin’s security claims rely on no single entity wielding a sufficiently large portion of the network’s overall computational power. Unfortunately, rather than participating independently, most Bitcoin miners join coalitions called *mining pools* in which a central pool administrator largely directs the pool’s activity, leading to a consolidation of power. Recently, the largest mining pool has accounted for more than half of network’s total mining capacity. Relatedly, “hosted mining” service providers offer their clients the benefit of economies-of-scale, tempting them away from independent participation. We argue that the prevalence of mining coalitions is due to a limitation of the Bitcoin proof-of-work puzzle – specifically, that it affords an effective mechanism for enforcing cooperation in a coalition. We present several definitions and constructions for “nonoutsourcable” puzzles that thwart such enforcement mechanisms, thereby deterring coalitions. We also provide an implementation and benchmark results for our schemes to show they are practical.

## 1. INTRODUCTION

Bitcoin [32] and subsequent decentralized cryptocurrencies have gained rapid popularity in recent years, and are often quoted as “a peek into the future financial and payment infrastructure”. Security of such cryptocurrencies is critical, and to ensure security the most fundamental assumption made by decentralized cryptocurrencies is that *no single entity or administration wields a large fraction of the computational resources in the network*. Violation of this assumption can lead to severe attacks such as history revision and double spending which essentially nullify all purported security properties that are widely believed today.

However, two recent trends in mining – namely, *mining pools* and *hosted mining* – have led to the concentration of mining power, and have cast serious doubt on the well-foundedness of these fundamental assumptions that underly the security of Bitcoin-like cryptocurrencies. Specifically, mining pools exist because solo miners wish to hedge mining risks and obtain rewards at a more stable, steady rate. At several times over the past two years, the largest

handful of mining pools have accounted for well over a third of the network’s overall computing effort [10]. For example, recently the largest mining pool, GHash.IO, has even exceeded 50% of the total mining capacity.<sup>1</sup> Currently, Hosted mining, on the other hand, allows individuals to outsource their mining effort to one or a few large service providers. Hosted mining services have already emerged, such as Alydian [9], whose “launch day pricing was \$65,000 per Terahash, and mining hosting contracts are available in 5 and 10 Th/sec blocks” [9]. Hosted mining is appealing because it can potentially reduce miners’ cost due to economies of scale. Henceforth we will refer to both mining pools and hosted mining as mining coalitions.

Such large mining coalitions present a potential lurking threat to the security of Bitcoin-like cryptocurrencies. To exacerbate the matter, several recent works [18, 24] showed that it may be incentive compatible for a mining coalition to deviate from the honest protocol – in particular, Eyal and Sirer [18] showed that a mining concentration of about 1/3 of the network’s mining power can obtain disproportionately large rewards by exhibiting certain “selfish mining” behavior.

While alternatives to centralized mining pools are well-known and have been deployed for several years, (such as *P2Pool*, [44] a decentralized mining pool architecture), these have unfortunately seen extremely low user adoption (at the time of writing, they account for less than 2% of the network). Fundamentally, the problem is that Bitcoin’s reward mechanism provides no particular incentive for users to use these decentralized alternatives.

Increasing understanding of these problems has prodded extensive and continual discussions in the broad cryptocurrency community, regarding how to deter such coalitions from forming and retain the decentralized nature of Bitcoin-like cryptocurrencies [27]. The community demands a technical solution to this problem.

### 1.1 Our Results and Contributions

Our work provides a timely response to this community-wide concern [27], providing the *first formally founded solution* to combat Bitcoin mining centralization. Our key observation is the following: an enabling factor in the growth of mining pools is a simple yet effective enforcement mechanism; members of a mining pool do not inherently trust one another, but instead submit cryptographic proofs (called “shares”) to the other pool members (or to the pool operator), in order to demonstrate they are contributing work that can only benefit the pool (e.g., work that is tied to the pool operator’s public key).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
CCS’15 October 12–16, 2015, Denver, CO, USA  
Copyright 2015 ACM ISBN 978-1-4503-3832-5/15/10 ... \$15.00  
DOI: <http://dx.doi.org/10.1145/2810103.2813621>.

<sup>1</sup>See <http://arstechnica.com/security/2014/06/bitcoin-security-guarantee-shattered-by-anonymous-miner-with-51-network-power/>

**Strongly nonoutsourcable puzzles.** Our idea, therefore, is to disable such enforcement mechanisms in a cryptographically strong manner. To this end, we are the first to propose *strongly nonoutsourcable puzzles*, a new form of proof-of-work puzzles which additionally guarantee the following:

*If a pool operator can effectively outsource mining work to a worker, then the worker can steal the reward without producing any evidence that can potentially implicate itself.*

Intuitively, if we can enforce the above, then any pool operator wishing to outsource mining work to an untrusted worker runs the risk of losing its entitled mining reward, thus effectively creating a disincentive to outsource mining work (either in the form of mining pools or hosted mining). Our nonoutsourcable puzzle is broadly powerful in that it renders unenforceable even *external contractual agreements* between the pool operator and the worker. In particular, no matter whether the pool operator outsources work to the worker through a cryptocurrency smart contract or through an out-of-the-band legal contract, we guarantee that the worker can steal the reward without leaving behind evidence of cheating.

**Technical insights.** At a technical level, our puzzle achieves the aforementioned guarantees through two main insights:

- P1:** We craft our puzzle such that if a worker is doing a large part of the mining computation, it must possess a sufficiently large part of a “signing key” such that it can later sign over the reward to its own public key – effectively stealing the award from the pool operator;
- P2:** We offer a zero-knowledge spending option, such that a worker can spend the stolen reward in a way that reveals no information (including potential evidence that can be used to implicate itself).

As a technical stepping stone, we formulate a weaker notion of our puzzle referred to as a weakly nonoutsourcable puzzle. A weakly nonoutsourcable puzzle essentially guarantees property P1 above, but does not ensure property P2. As a quick roadmap, our plan is to first construct a weakly nonoutsourcable puzzle, and from there we devise a generic zero-knowledge transformation to compile a weakly nonoutsourcable puzzle into a strongly nonoutsourcable one. It turns out that a weakly nonoutsourcable puzzle is the implicit security notion adopted by the recent work of Permacoin [29] but without being formalized there. In Section 6, we argue that weakly nonoutsourcable puzzles alone are inadequate to defeat mining coalitions, and in particular hosted mining.

**Implementation and practical performance.** We show implementation and evaluation results to demonstrate the practical performance of our puzzles. Based on an instantiation using the succinct zero-knowledge option of Libsnark [5], we show that it would take a cheating worker only **14 seconds** (using approximately a thousand parallel cores) to successfully steal a block reward. Further, stealing a block’s reward in zero knowledge consumes only **\$10** worth of Amazon AWS compute-time, which is very small in comparison with the block’s reward – roughly **\$8,750**, based on Bitcoin’s current market price. Clearly this provides a sufficiently strong deterrent against mining coalitions. Note also that this zero-knowledge spending option is not normally incurred, since honest miners can simply adopt a cheap plaintext spending option whose cost is insignificant (see Section 6). For both the zero-knowledge and the cheap plaintext spending options, the block verification overhead is insignificant (at most 1.7 seconds) in comparison with the present Bitcoin epoch length (roughly 10 minutes).

**Deployment considerations.** For our nonoutsourcable puzzles to be practically deployed, it is also important to address several additional challenges, such as how to still allow miners to reduce mining uncertainty (i.e., the positive effects of mining pools), and how to simultaneously address various other design goals such as ASIC resistance, and lightweight clients. We give detailed explanations to address these practical issues related to deployment (Section 7.3). Notably, inspired by the design of state lottery games, we propose a new, multi-tier reward system that allows us to achieve the best of both worlds: ensure non-outsourcability of puzzles, and meanwhile allow smaller players to reduce payoff variance.

**Community demand and importance of formal security.** The community’s demand for a nonoutsourcable puzzle is also seen in the emergence of new altcoins [1, 41] that (plan to) adopt their own home-baked versions of nonoutsourcable puzzles. Their solutions, however, offer only weak nonoutsourcability, and do not provide any formal guarantees. The existence of these custom constructions further motivates our efforts, and demonstrates that it is non-trivial to both formalize the security notions as well as design constructions with provable security. To date, our work provides *the only formally-founded solution*, as well as the *first strongly nonoutsourcable puzzle construction*.

## 2. BITCOIN BACKGROUND

We define puzzles and nonoutsourcable puzzles as an independent concept, abstracting away the less relevant details about the Bitcoin protocol itself. Later, however, we will discuss how the puzzles we introduce can be integrated into a Bitcoin-like distributed digital currency. For this reason, as well as to understand the motivation behind our formal definitions, we first present some additional background on Bitcoin and its use of computational puzzles. For a more thorough explanation of the Bitcoin protocol, we refer the readers to [3, 7, 32].

**Puzzles, rewards, and epochs.** In Bitcoin, new money is printed at a predictable rate, through a distributed coin-minting process. At the time of writing, roughly speaking, 25 bitcoins are minted every 10 minutes (referred to as an epoch) on average. When an epoch begins, a public puzzle instance is generated by computing an up-to-date hash of the global transaction log (called the “blockchain”). Then, Bitcoin nodes race to solve this epoch’s puzzle. Whoever first finds an eligible solution to the puzzle can claim the newly minted coins corresponding to this epoch.

In slightly more detail, miners start with the puzzle instance  $\text{puz}$ , and construct a payload  $m$  which contains (a tree hash over) the miners public key and a new set of transaction to commit to the log during this epoch. He then searches for a nonce  $r$  such that  $\mathcal{H}(\text{puz}||m||r) < 2^{\lambda-d}$ , where  $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$  is a hash function and  $d$  is a difficulty parameter. The difficulty parameter is adjusted according to the total amount of computational resources devoted to mining to ensure that each epoch lasts 10 minutes on average.

In Section 4, we formally define a generalization of Bitcoin’s puzzle called *scratch-off puzzles*. More detail about the original Bitcoin puzzle construction can be found in the full online version of our paper.

**Consensus mechanism.** Bitcoin nodes reach consensus on the history of transactions by having nodes accept the blockchain with the largest total difficulty. Roughly speaking, this defeats history revision attacks, since to revise history would involve computing a blockchain that is more difficult than the known good chain. An adversary must therefore possess a significant fraction of the total

computational resources to successfully race against the rest of the network in extending the chain.

Bitcoin is novel in its use of computational puzzles as part of a consensus protocol for anonymous networks without any pre-established PKI. A related approach was earlier proposed by Aspin et al. [2], although their network model nonetheless retained a strong assumption about pre-established point-to-point channels.

### 3. SCRATCH-OFF PUZZLES

As introduced earlier, the Bitcoin protocol is built around a moderately hard computational puzzle. Bitcoin miners compete to solve these puzzles, and whoever solves a puzzle first in each epoch receives a reward. As there is no shortcut to solving this puzzle, for an attacker to dominate the network would require the attacker to expend more computational resources than the rest of the honest participants combined. Although the Bitcoin puzzle is commonly referred to as a *proof-of-work* puzzle, the requirements of the puzzle are somewhat different than existing definitions for proof-of-work puzzles [13, 16, 21, 42].

Before proceeding with our main contribution of nonoutsourcable puzzles, we first provide a formal definition of the basic requirements of the Bitcoin puzzle, which we call a *scratch-off puzzle*.<sup>2</sup> In particular, while a traditional proof-of-work puzzle [21] need only be solvable by a single sequential computation, a scratch-off puzzle must be solvable by several concurrent non-communicating entities.

In what follows, we let  $\lambda$  denote a security parameter. A scratch-off puzzle is parameterized by parameters  $(\underline{t}, \mu, d, t_0)$  where, informally speaking,  $\underline{t}$  denotes the amount of work needed to attempt a single puzzle solution,  $\mu$  refers to the maximum amount by which an adversary can speed up the process of finding solutions,  $d$  affects the average number of attempts to find a solution, and  $t_0$  denotes the initialization overhead of the algorithm. We typically assume that  $t_0 \ll 2^d \underline{t}$ , where  $2^d \underline{t}$  is the expected time required to solve a puzzle.

**DEFINITION 1.** *A scratch-off puzzle is parameterized by parameters  $(\underline{t}, \mu, d, t_0)$ , and consists of the following algorithms (satisfying properties explained shortly):*

1.  $\mathcal{G}(1^\lambda) \rightarrow \text{puz}$ : generates a puzzle instance.
2.  $\text{Work}(\text{puz}, m, t) \rightarrow \text{ticket}$ : The Work algorithm takes a puzzle instance  $\text{puz}$ , some payload  $m$ , and time parameter  $t$ . It makes  $t$  unit scratch attempts, using  $t \cdot \underline{t} + t_0$  time steps in total. Here  $\underline{t} = \text{poly}(\lambda)$  is the unit scratch time, and  $t_0$  can be thought of as the initialization and finalization cost of Work.
3.  $\text{Verify}(\text{puz}, m, \text{ticket}) \rightarrow \{0, 1\}$ : checks if a ticket is valid for a specific instance  $\text{puz}$ , and payload  $m$ . If ticket passes this check, we refer to it as a winning ticket for  $(\text{puz}, m)$ .

Intuitively, the honest Work algorithm makes  $t$  unit scratch attempts, and each attempt has probability  $2^{-d}$  of finding a winning ticket, where  $d$  is called the puzzle’s difficulty parameter. For simplicity, we will henceforth use the notation

$$\zeta(t, d) := 1 - (1 - 2^{-d})^t$$

to refer to the probability of finding a winning ticket using  $t$  scratch attempts. For technical reasons that will become apparent later, we additionally define the shorthand  $\zeta^+(t, d) := \zeta(t + 1, d)$ . For the

<sup>2</sup>The terms “scratch-off puzzle” and “winning ticket” are motivated by the observation that Bitcoin’s coin minting process resembles a scratch-off lottery, wherein a participant expends a unit of effort to learn if he holds a winning ticket.

remainder of the paper, we assume that the puzzle’s difficulty parameter  $d$  is fixed, hence we omit the  $d$  and write  $\zeta(t)$  and  $\zeta^+(t)$  for simplicity. We also define the algorithm  $\text{WorkTillSuccess}(\text{puz}, m)$  as  $\text{Work}(\text{puz}, m, \infty)$ ; i.e., this algorithm runs until it finds a winning ticket for the given instance and payload.

A scratch-off puzzle must satisfy three requirements:

1. **Correctness.** For any  $(\text{puz}, m, t)$ , if  $\text{Work}(\text{puz}, m, t)$  outputs ticket  $\neq \perp$ , then  $\text{Verify}(\text{puz}, m, \text{ticket}) = 1$ .
2. **Feasibility and parallelizability.** Solving a scratch-off puzzle is feasible, and can be parallelized. More formally, for any  $\ell = \text{poly}(\lambda)$ , for any  $t_1, t_2, \dots, t_\ell = \text{poly}(\lambda)$ , let  $t := \sum_{i \in [\ell]} t_i$ .

$$\Pr \left[ \begin{array}{l} \text{puz} \leftarrow \mathcal{G}(1^\lambda), \\ m \leftarrow \{0, 1\}^\lambda, \\ \forall i \in [\ell] : \text{ticket}_i \leftarrow \text{Work}(\text{puz}, m, t_i) : \\ \exists i \in [\ell] : \text{Verify}(\text{puz}, m, \text{ticket}_i) \end{array} \right] \geq \zeta(t) - \text{negl}(\lambda).$$

Intuitively, each unit scratch attempt, taking time  $\underline{t}$ , has probability  $2^{-d}$  of finding a winning ticket. Therefore, if  $\ell$  potentially parallel processes each makes  $t_1, t_2, \dots, t_\ell$  attempts, the probability of finding one winning ticket overall is  $\zeta(t) \pm \text{negl}(\lambda)$  where  $t = \sum_{i \in [\ell]} t_i$ .

3.  **$\mu$ -Incompressibility.** Roughly speaking, the work for solving a puzzle must be incompressible in the sense that even the best adversary can speed up the finding of a puzzle solution by at most a factor of  $\mu$ . More formally, a scratch-off puzzle is  $\mu$ -incompressible (where  $\mu \geq 1$ ) if for any probabilistic polynomial-time adversary  $\mathcal{A}$  taking at most  $t \cdot \underline{t}$  steps,

$$\Pr \left[ \begin{array}{l} \text{puz} \leftarrow \mathcal{G}(1^\lambda), \\ (m, \text{ticket}) \leftarrow \mathcal{A}(\text{puz}) : \\ \text{Verify}(\text{puz}, m, \text{ticket}) = 1 \end{array} \right] \leq \zeta^+(\mu t) \pm \text{negl}(\lambda).$$

Note that  $\zeta^+(t) = 1 - (1 - 2^{-d})^{t+1}$  is roughly the probability of outputting a winning ticket after  $t$  unit scratch attempts, though we additionally allow the adversary to make a final guess at the end (as in [42]), and hence the  $t + 1$  in the exponent instead of just  $t$ . Ideally, we would like the compressibility factor  $\mu$  to be as close to 1 as possible. When  $\mu = 1$ , the honest Work algorithm is the optimal way to solve a puzzle.

This definition implies, in particular, that solutions to previous puzzles do not help in solving a freshly generated puzzle unseen ahead of time.

#### 3.1 Non-Transferability

For a practical scheme we could integrate into Bitcoin, we should require that the payload of a ticket is non-transferable, in the following sense: if an honest party publishes a ticket attributed to a payload  $m$  (e.g., containing a public key belonging to the party to whom the reward must be paid), the adversary should not gain any advantage in obtaining a puzzle solution attributed to some different payload  $m^*$  for the same  $\text{puz}$ . This is because in Bitcoin, each epoch is defined by a globally known, unique puzzle instance  $\text{puz}$ ; at most one winning ticket for  $\text{puz}$  and a payload message is accepted into the blockchain; and a user who solves a puzzle only receives the reward if their message is the one that is attributed. If an adversary can easily modify a victim’s winning ticket to be attributed to a different payload of its choice, then the adversary can listen for when the victim’s ticket is first announced in the network, and then immediately start propagating the modified ticket (e.g., containing its own public key for the reward payment) and attempt to outrace the victim. It is possible that the network

will now deem the adversary as the winner of this epoch—this is especially true if the adversary has better network connectivity than the victim (as described in [18]). For simplicity in developing our constructions and nonoutsourcable definition, we define this non-transferability requirement separately below. Intuitively, non-transferability means that seeing a puzzle solution output by an honest party does not help noticeably in producing a solution attributed to a different payload  $m^*$ .

**DEFINITION 2.** Let  $\delta$  be a nonnegative function of  $\ell$ . A scratch-off puzzle is  $\delta$ -non-transferable if it additionally satisfies the following property:

For any  $\ell = \text{poly}(\lambda)$ , and for any adversary  $\mathcal{A}$  taking  $t \cdot \underline{t}$  steps,

$$\Pr \left[ \begin{array}{l} \text{puz} \leftarrow \mathcal{G}(1^\lambda) \\ m_1, m_2, \dots, m_\ell \leftarrow \mathcal{A}(1^\lambda) \\ \forall i \in [\ell] : \text{ticket}_i \leftarrow \text{WorkTillSuccess}(\text{puz}, m_i), \\ (\text{puz}, m^*, \text{ticket}^*) \leftarrow \mathcal{A}(\text{puz}, \{m_i, \text{ticket}_i\}_{i=1}^\ell) : \\ \text{Verify}(\text{puz}, \text{ticket}^*, m^*) \wedge (\forall i \in [\ell] : m^* \neq m_i) \end{array} \right] \leq \zeta^+(\mu + \delta)t + \text{negl}(\lambda)$$

## 4. OUTSOURCED MINING AND WEAKLY NONOUTSOURCEABLE PUZZLES

The Bitcoin scratch-off puzzle described in the previous section is amenable to secure outsourcing, in the sense that it is possible for one party (the *worker*) to perform mining work for the benefit of another (the *pool operator*) and to *prove* to the pool operator that the work done can only benefit the pool operator.

To give a specific example, let  $m$  be the public key of the pool operator; if the worker performs  $2^{d'}$  scratch attempts, on average it will have found at least one value  $r$  such that  $\mathcal{H}(\text{puz}||m||r) < 2^{\lambda-d'}$ . The value  $r$  can be presented to the pool operator as a “share” (since it represents a portion of the expected work needed to find a solution); intuitively, any such work associated with  $m$  cannot be reused for any other  $m^* \neq m$ . This scheme is an essential component of nearly every Bitcoin mining pool to date [39]; the mining pool operator chooses the payload  $m$ , and mining participants are required to present shares associated with  $m$  in order to receive participation credit. The rise of large, centralized mining pools is due in large part to the effectiveness of this mechanism.

We now formalize a generalization of this outsourcing protocol, and then proceed to construct puzzles that are *not* amenable to outsourcing (i.e., for which no effective outsourcing protocol exists).

### 4.1 Notation and Terminology

**Pool operator and Worker.** We use the terminology *pool operator* and *worker* referring respectively to the party outsourcing the mining computation and the party performing the mining computation. While this terminology is natural for describing mining pools, we stress that our results are intended to simultaneously discourage both mining pools and hosted mining services. In the case of hosted mining, the roles are roughly swapped; the cloud server performs the mining work, and the individuals who hire the service receive the benefit and must be convinced the work is performed correctly. We use this notation since mining pools are more well-known and widely used today, and therefore we expect the mining-pool oriented terminology to be more familiar and accessible.

**Protocol executions.** A protocol is defined by two algorithms  $\mathcal{S}$  and  $\mathcal{C}$ , where  $\mathcal{S}$  denotes the (honest) worker, and  $\mathcal{C}$  the (honest)

pool operator. We use the notation  $(o_S; o_C) \leftarrow (\mathcal{S}, \mathcal{C})$  to mean that a pair of interactive Turing Machines  $\mathcal{S}$  and  $\mathcal{C}$  are executed, with  $o_S$  the output of  $\mathcal{S}$ , and  $o_C$  the output of  $\mathcal{C}$ .

In this paper we assume the pool operator executes the protocol program  $\mathcal{C}$  correctly, but the worker may deviate arbitrarily.<sup>3</sup> We use the notation  $(\mathcal{A}, \mathcal{C})$  to denote an execution between a malicious worker  $\mathcal{A}$  and an honest pool operator  $\mathcal{C}$ . Note that protocol definition always uses the honest algorithms, i.e.,  $(\mathcal{S}, \mathcal{C})$  denotes a protocol or an honest execution; whereas  $(\mathcal{A}, \mathcal{C})$  represents an execution.

## 4.2 Definitions

**Outsourcing protocol.** We now define a generalization of outsourced mining protocols, encompassing both mining pools and hosted mining services. Our definition of outsourcing protocol is broad – it captures any form of protocol where the pool operator and worker may communicate as interactive Turing Machines, and at the end, the pool operator may obtain a winning ticket with some probability. The protocol is parametrized by three parameters  $t_C, t_S$ , and  $t_e$ , which roughly models the pool operator’s work, honest worker’s work, and the “effective” amount of work during the protocol.

**DEFINITION 3.** A  $(t_S, t_C, t_e)$ -outsourcing protocol for scratch-off puzzle  $(\mathcal{G}, \text{Work}, \text{Verify})$ , where  $t_e < t_S + t_C$  and  $t_c < t_e$ , is a two-party protocol,  $(\mathcal{S}, \mathcal{C})$ , such that

- The pool operator’s input is  $\text{puz}$ , and the worker’s input is  $\perp$ .
- The pool operator  $\mathcal{C}$  runs in at most  $t_C \cdot \underline{t}$  time, and the worker  $\mathcal{S}$  in at most  $t_S \cdot \underline{t}$  time.
- $\mathcal{C}$  outputs a tuple  $(\text{ticket}, m)$  at the end, where  $\text{ticket}$  is either a winning ticket for payload  $m$  or  $\text{ticket} = \perp$ . Further, when interacting with an honest  $\mathcal{S}$ ,  $\mathcal{C}$  outputs a  $\text{ticket} \neq \perp$  with probability at least  $\zeta(t_e) - \text{negl}(\lambda)$ .

Formally,

$$\Pr \left[ \begin{array}{l} \text{puz} \leftarrow \mathcal{G}(1^\lambda) \\ (\cdot; \text{ticket}, m) \leftarrow (\mathcal{S}, \mathcal{C}(\text{puz})) : \\ \text{Verify}(\text{puz}, m, \text{ticket}) \end{array} \right] \geq \zeta(t_e) - \text{negl}(\lambda).$$

The parameter  $t_e$  is referred to as the *effective billable work*, because the protocol  $(\mathcal{S}, \mathcal{C})$  has the success probability of performing  $t_e$  unit scratch attempts. Note that it must be the case that  $t_e < \mu(t_S + t_C)$ . Intuitively, an outsourcing protocol allows effective outsourcing of work by the pool operator if  $t_e \gg t_C$ .

Note that this definition does not specify how the payload  $m$  is chosen. In typical Bitcoin mining pools, the pool operator chooses  $m$  so that it contains the pool operator’s public key. However, our definition also includes schemes where  $m$  is jointly computed during interaction between  $\mathcal{S}$  and  $\mathcal{C}$ , for example.

**Weak nonoutsourcability.** So far, we have formally defined what an outsourcing protocol is. Roughly speaking, an outsourcing protocol generally captures any possible form of contractual agreement between the pool operator and the worker. The outsource protocol defines exactly what the worker has promised to do for the pool operator, i.e., the “honest” worker behavior. If a worker is malicious, it need not follow this honest prescribed behavior. The notion of weak non-outsourcability requires that no matter what the prescribed contractual agreement is between the pool operator and the worker— as long as this agreement “effectively” outsources work to

<sup>3</sup>This is without loss of generality, and does not mean that we assume the mining pool operator is honest, since the protocol  $(\mathcal{S}, \mathcal{C})$  may deviate from “honest” Bitcoin mining.

the worker—there exists an adversarial worker that can always steal the pool operator’s ticket should the pool operator find a winning ticket during the protocol. Effectiveness is intuitively captured by how much effective work the worker performs vs. the work performed by the pool operator in the honest protocol. Note that there always exists a trivial, ineffective outsourcing protocol, where the pool operator always performs all the work by itself – in this case, a malicious worker will not be able to steal the ticket. Therefore, the weak non-outsourcability definition is parametrized by the effectiveness of the honest outsourcing protocol.

More specifically, the definition says that the adversarial worker can generate a winning ticket associated with a payload of its own choice, over which the pool operator has no influence. In a Bitcoin-like application, a natural choice is for an adversarial worker to replace the payload with a public key it owns (potentially a pseudonym), such that it can later spend the stolen awards. Based on this intuition, we now formally define the notion of a *weakly nonoutsourcable* scratch-off puzzle.

**DEFINITION 4.** A *scratch-off-puzzle* is  $(t_S, t_C, t_e, \alpha, p_s)$ -weakly nonoutsourcable if for every  $(t_S, t_C, t_e)$ -outsourcing protocol  $(\mathcal{S}, \mathcal{C})$ , there exists an adversary  $\mathcal{A}$  that runs in time at most  $t_S \cdot t + \alpha$ , such that:

- Let  $m^* \xleftarrow{\$} \{0, 1\}^\lambda$ . Then, at the end of an execution  $(\mathcal{A}(\text{puz}, m^*), \mathcal{C}(\text{puz}))$ , the probability that  $\mathcal{A}$  outputs a winning ticket for payload  $m^*$  is at least  $p_s \zeta(t_e)$ . Formally,

$$\Pr \left[ \begin{array}{l} \text{puz} \leftarrow \mathcal{G}(1^\lambda); m^* \xleftarrow{\$} \{0, 1\}^\lambda \\ (\text{ticket}^*; \text{ticket}, m) \leftarrow (\mathcal{A}(\text{puz}, m^*), \mathcal{C}(\text{puz})); \\ \text{Verify}(\text{puz}, \text{ticket}^*, m^*) \end{array} \right] \geq p_s \zeta(t_e).$$

- Let  $\text{view}_h$  denote the pool operator’s view in an execution with the honest worker  $(\mathcal{S}, \mathcal{C}(\text{puz}))$ , and let  $\text{view}^*$  denote the pool operator’s view in an execution with the adversary  $(\mathcal{A}(\text{puz}, m^*), \mathcal{C}(\text{puz}))$ . Then,

$$\text{view}^* \stackrel{c}{\equiv} \text{view}_h.$$

When  $\mathcal{C}$  interacts with  $\mathcal{A}$ , the view of the pool operator  $\text{view}^*$  is computationally indistinguishable from when interacting with an honest  $\mathcal{S}$ .

Later, when proving that puzzles are weakly nonoutsourcable, we typically construct an adversary  $\mathcal{A}$  that runs the honest protocol  $\mathcal{S}$  until it finds a ticket for  $m$ , and then transforms the ticket into one for  $m^*$  with probability  $p_s$ . For this reason, we refer to the adversary  $\mathcal{A}$  in the above definition as a *stealing adversary* for protocol  $(\mathcal{S}, \mathcal{C})$ . In practice, we would like  $\alpha$  to be small, and  $p_s \leq 1$  to be large, i.e.,  $\mathcal{A}$ ’s run-time is not much different from that of the honest worker, but  $\mathcal{A}$  can steal a ticket with high probability.

If the pool operator outputs a valid ticket for  $m$  and the worker outputs a valid ticket for  $m^*$ , then there is a race to determine which ticket is accepted by the Bitcoin network and earns a reward. Since the  $\mu$ -incompressibility of the scratch-off puzzle guarantees the probability of generating a winning ticket associated with either  $m$  or  $m^*$  is bounded above by  $\zeta^+(\mu(t_S + t_C))$ , the probability of the pool operator outputting a ticket — but not the worker — is bounded above by  $\zeta^+(\mu(t_S + t_C)) - p_s \zeta(t_e)$ .

Note that weak nonoutsourcability does not imply that the puzzle is transferable. In other words, a puzzle can be simultaneously non-transferable and weakly nonoutsourcable. This is so because the stealing adversary  $\mathcal{A}$  may rely on its view of the entire outsourcing protocol when stealing the ticket for its own payload  $m^*$ , whereas the adversary for the non-transferability game is only given winning tickets as input (but no protocol views).

As we mentioned in the beginning of this section, the prevalence of Bitcoin mining pools can be attributed in part to the effective outsourcing protocol used to coordinate untrusted pool members – in other words, the Bitcoin puzzle is not nonoutsourcable. We state and prove a theorem to this effect in the full online version of this paper.

## 5. A WEAKLY NONOUTSOURCEABLE PUZZLE

In this section, we describe a weakly nonoutsourcable construction based on a Merkle-hash tree construction. We prove that our construction satisfies weak nonoutsourcability (for a reasonable choice of parameters) in the random oracle model. Informally, our construction achieves the following:

For any outsourcing protocol that can effectively outsource a fixed constant fraction of the effective work, an adversarial worker will be able to steal the puzzle with at least constant probability.

Our construction is inspired by the Floating Preimage Signature (FPS) scheme used in Permacoin [29], which is a puzzle integrated with a proof-of-retrievability. However, Permacoin [29] only described the issue of nonoutsourcability informally, and made no attempt to formalize the definition nor to discuss nonoutsourcability beyond the context of archival storage. Our construction is formally defined in our online full version [30], but here we provide an informal explanation of the intuition behind it.

**Intuition.** To solve a puzzle, a node first builds a Merkle tree with random values at the leaves; denote the root by digest. Then the node repeatedly samples a random value  $r$ , computes  $h = \mathcal{H}(\text{puz} || r || \text{digest})$ , and uses  $h$  to select  $q$  leaves of the Merkle tree and their corresponding branches (i.e., the corresponding Merkle proofs). It then hashes those branches (along with  $\text{puz}$  and  $r$ ) and checks to see if the result is less than  $2^{\lambda-d}$ .

Once successful, the node has a value  $r$  what was “difficult” to find, but is not yet bound to the payload message  $m$ . To effect such binding, a “signing step” is performed in which  $h' = \mathcal{H}(\text{puz} || m || \text{digest})$  is used to select a set of  $4q'$  leaf nodes (i.e., using  $h'$  a seed to a pseudorandom number generator). Any  $q'$  of these leaves, along with their corresponding branches, constitute a signature for  $m$  and complete a winning ticket.

Intuitively, this puzzle is weakly nonoutsourcable because in order for the worker to perform scratch attempts, it must

- either know a large fraction of the leaves and branches of the Merkle tree, in which case it will be able to sign an arbitrary payload  $m^*$  with high probability – by revealing  $q'$  out of the  $4q'$  leaves (and their corresponding branches) selected by  $m^*$ ,
- or incur a large amount of overhead, due to aborting scratch attempts for which it does not know the necessary leaves and branches,
- or interact with the pool operator frequently, in which case the pool operator performs a significant fraction of the total number of random oracle queries.

To formally prove this construction is weakly nonoutsourcable, we assume that the cost of the Work algorithm is dominated by calls made to random oracles. Thus, for simplicity, in the following theorems we equate the running time with the number of calls to the random oracle. However, the theorem can be easily generalized (i.e., relaxing by a constant factor) as long as the cost of the rest of the computation is only a constant fraction of the random-oracle calls.

The construction sketched above, and formally defined in our online full version, is a scratch-off puzzle. We defer this proof to our online full version [30].

Let  $q, q' = O(\lambda)$ . Let the number of leaves  $L \geq q + 8q'$ . Suppose  $d > 10$  and  $t_e \cdot 2^{-d} < 1/2$ . Under the aforementioned cost model, the above construction is a  $(t_S, t_C, t_e, \alpha, p_s)$  weakly non-outsourceable puzzle, for any  $0 < \gamma < 1$  s.t.  $t_C < \gamma t_e, p_s > \frac{1}{2}(1 - \gamma) - \text{negl}(\lambda)$ , and  $\alpha = O(\lambda^2)$ ; and is 0-non-transferable. In other words, if the pool operator's work  $t_C$  is a not a significant fraction of  $t_e$ , i.e., work is effectively outsourced, then an adversarial worker will be able to steal the pool operator's ticket with a reasonably big probability, and without too much additional work than the honest worker.

The proof that this puzzle is weakly nonoutsourcable can be found in our online full version [30], but we sketch the main idea here. Informally, to “effectively” outsource work to the worker, the worker must know more than a constant fraction (say, 1/3) of the leaves before calling the random oracle to determine whether an attempt is successful. However, if the worker knows more than 1/3 fraction of the leaves, due to a simple Chernoff bound, it will be able to easily steal the solution should one be found. To make this argument formally is more intricate. For lack of space, we defer (to the full online version of our paper) the proof that this puzzle is non-transferable.

## 6. STRONGLY NONOUTSOURCEABLE PUZZLES

In the previous section, we formally defined and constructed a scheme for weakly nonoutsourcable puzzles, which ensure that for any “effective” outsourcing protocol, there exists an adversarial worker that can steal the pool operator's winning ticket with significant probability, should a winning ticket be found. This can help deter outsourcing when individuals are expected to behave selfishly.

One critical drawback of the weakly nonoutsourcable scheme (and, indeed, of Permacoin [29]) is that a stealing adversary may be detected when he spends his stolen reward, and thus might be held accountable through some external means, such as legal prosecution or a tainted public reputation.

For example, a simple detection mechanism would be for the pool operator and worker to agree on a  $\lambda/2$ -bit prefix of the nonce space to serve as a watermark. The worker can mine by randomly choosing the remaining  $\lambda/2$ -bit suffix, but the pool operator only *accepts* evidence of mining work bearing this watermark. If the worker publishes a stolen puzzle solution, the watermark would be easily detectable.

Ideally, we should enable the stealing adversary to evade detection and leave no incriminating trail of evidence. Therefore, in this section, we define a “strongly nonoutsourcable” puzzle, which has the additional requirement that a stolen ticket cannot be distinguished from a ticket produced through independent effort.

**DEFINITION 5.** A puzzle is  $(t_S, t_C, t_e, \alpha, p_s)$ -strongly nonoutsourcable if it is  $(t_S, t_C, t_e, \alpha, p_s)$ -weakly nonoutsourcable, and additionally the following holds:

For any  $(t_S, t_C, t_e)$ -outsourcing protocol  $(S, C)$ , there exists an adversary  $\mathcal{A}$  for the protocol such that the stolen ticket output by  $\mathcal{A}$  for payload  $m^*$  is computationally indistinguishable from a honestly computed ticket for  $m^*$ , even given the pool operator's view in the execution  $(\mathcal{A}, C)$ . Formally, let  $\text{puz} \leftarrow \mathcal{G}(1^\lambda)$ , let  $m^* \xleftarrow{\$} \{0, 1\}^\lambda$ . Consider a protocol execution  $(\mathcal{A}(\text{puz}, m^*), C(\text{puz}))$ : let  $\text{view}^*$  denote the pool operator  $C$ 's view and  $\text{ticket}^*$  the stolen

Let NIZK be a non-interactive zero-knowledge proof system. Also assume that  $\mathcal{E} = (\text{Key}, \text{Enc}, \text{Dec})$  is a CPA-secure public-key encryption scheme.

Let  $(\mathcal{G}', \text{Work}', \text{Verify}')$  be a weakly nonoutsourcable scratch-off puzzle scheme. We now construct a strongly nonoutsourcable puzzle scheme as below.

- $\mathcal{G}(1^\lambda)$ : Run the puzzle generation of the underlying scheme  $\text{puz}' \leftarrow \mathcal{G}'(1^\lambda)$ . Let  $\text{crs} \leftarrow \text{NIZK.Setup}(1^\lambda)$ ; and let  $(\text{sk}_\mathcal{E}, \text{pk}_\mathcal{E}) \leftarrow \mathcal{E}.\text{Key}(1^\lambda)$ . Output  $\text{puz} \leftarrow (\text{puz}', \text{crs}, \text{pk}_\mathcal{E})$
- $\text{Work}(\text{puz}, m, t)$ :
  - Parse  $\text{puz} := (\text{puz}', \text{crs}, \text{pk}_\mathcal{E})$ .
  - $\text{ticket}' \leftarrow \text{Work}'(\text{puz}', m, t)$ .
  - Encrypt  $c \leftarrow \text{Enc}(\text{pk}_\mathcal{E}; \text{ticket}'; r)$ .
  - Set  $\pi \leftarrow \text{NIZK.Prove}(\text{crs}, (c, m, \text{pk}_\mathcal{E}, \text{puz}'), (\text{ticket}', r))$  for the following NP statement:
    - $\text{Verify}'(\text{puz}', m, \text{ticket}') \wedge c = \text{Enc}(\text{pk}_\mathcal{E}; \text{ticket}'; r)$
  - Return  $\text{ticket} := (c, \pi)$ .
- $\text{Verify}(\text{puz}, m, \text{ticket})$ :
  - Parse  $\text{puz} := (\text{puz}', \text{crs}, \text{pk}_\mathcal{E})$ , and parse  $\text{ticket}$  as  $(c, \pi)$ .
  - Check that  $\text{Verify}(\text{crs}, (c, m, \text{pk}_\mathcal{E}, \text{puz}'), \pi) = 1$ .

Figure 1: A generic transformation from any *weakly nonoutsourcable* scratch-off puzzle to a *strongly nonoutsourcable* puzzle.

ticket output by  $\mathcal{A}$  in the execution. Let  $\text{ticket}_h$  denote an honestly generated ticket for  $m^*$ , ( $\text{ticket}_h := \text{WorkTillSuccess}(\text{puz}, m^*)$ ), and let  $\text{view}_h$  denote the pool operator's view in the execution  $(S, C(\text{puz}))$ . Then,

$$(\text{view}^*, \text{ticket}^*) \stackrel{c}{\equiv} (\text{view}_h, \text{ticket}_h)$$

Recall that in Bitcoin, the message payload  $m$  typically contains a Merkle root hash representing a set of new transactions to commit to the blockchain in this round, including the public key to which the reward is assigned. Thus to take advantage of the strongly nonoutsourcable puzzle, the stealing worker should bind its substituted payload  $m^*$  to a *freshly generated* public key for which it knows the corresponding private key. It can then spend its stolen reward anonymously, for example by laundering the coins through a mixer [8].

In Figure 1, we present a generic transformation that turns any weakly nonoutsourcable puzzle into a strongly nonoutsourcable puzzle. The strengthened puzzle is essentially a *zero-knowledge* extension of the original – a ticket for the strong puzzle is effectively a proof of the statement “I know a solution to the underlying puzzle.”

If  $(\text{GenKey}', \text{Work}', \text{Verify}')$  is a  $(t_S, t_C, t_e, \alpha, p_s)$  weakly nonoutsourcable puzzle, then the puzzle described in Figure 1 is a  $(t_S, t_C, t_e, \alpha + t_{\text{enc}} + t_{\text{NIZK}}, p_s - \text{negl}(\lambda))$  strongly nonoutsourcable puzzle, where  $t_{\text{enc}} + t_{\text{NIZK}}$  is the maximum time required to compute the encryption and NIZK in the honest  $\text{Work}$  algorithm.

We next state a theorem that this generic transformation essentially preserves the non-transferability of the underlying puzzle. If the underlying puzzle  $(\mathcal{G}', \text{Work}', \text{Verify}')$  is  $\delta'$ -non-transferable, then the derived puzzle through the generic transformation is  $\delta$  non-transferable for

$$\mu + \delta' \leq \frac{(\mu + \delta)t}{t \cdot \underline{t} + (t_{\text{enc}} + t_{\text{nizk}})\ell}$$

where  $t_{\text{enc}}$  and  $t_{\text{nizk}}$  are the time for performing each encryption and NIZK proof respectively.

Again, due to space restrictions, we defer the proof of this theorem to the full online version of our paper.

**Cheap plaintext option.** Although we have shown it is plausible for a stealing worker (with parallel resources) to compute the zero-knowledge proofs, this would place an undue burden on honest independent miners. However, it is possible to modify our generic transformation so that there are *two* ways to claim a ticket: the first is with a zero-knowledge proof as described, while the second is simply by revealing a plaintext winning ticket for the underlying weakly nonoutsourcable puzzle.

## 7. IMPLEMENTATION AND MICRO-BENCHMARKS

In order to demonstrate the practicality of our schemes, we implemented both our weakly nonoutsourcable and strongly nonoutsourcable puzzle schemes and provide benchmark results below.

**Metrics.** We are concerned with two main performance criteria. First, the size of a ticket and cost of verifying a ticket should be minimal, since each participant on the network is expected to verify every ticket independently. Second, in order for our scheme to be an effective deterrent, the cost and *latency* required to “steal” a ticket should be low enough that it is at least *plausible* for an outsourcing worker to compute a stolen ticket *and* propagate it throughout the network before any other solution is found.

When comparing the verification cost of our schemes to that of the current Bitcoin protocol, we include both the cost due to the puzzle itself, as well as the total cost of validating a block including transactions. At present, there are over 400 transactions per block on average;<sup>4</sup> we assume each transaction carries at least 1 ECDSA signature that must be verified. In general, the computational cost of validating blocks in Bitcoin is largely dominated by verifying the ECDSA signatures in transactions rather than verifying puzzle solutions. We measured that the time to verify an ECDSA signature on a 2.4GHZ Intel CPU is 1.7ms.<sup>5</sup> On average, at the time of writing, a block contains about 350 kilobytes of data and 600 transactions, each with an average of two signatures.<sup>6</sup>

### 7.1 Our Weakly Nonoutsourcable Puzzle

The weakly nonoutsourcable puzzle is straightforward to implement, and its overhead relative to the Bitcoin puzzle consists only of  $\lambda \log \lambda$  additional hashes; we implemented this in unoptimized Python and discuss its performance later on. In contrast, the strongly nonoutsourcable puzzle requires much more care in implementation due to the NIZK proof, which we will describe shortly.

We used the SHA-1 hash function throughout our implementation, since this has a relatively efficient implementation as an arithmetic circuit [36]. We restricted our focus to the following puzzle parameters: the signature tree consists of  $2^h = 2^{10}$  leaves, and the number of leaves revealed during a scratch attempt and a claim is  $q = q' = 10$ . This provides roughly 50 bits of security for the non-transferability property.

**Performance results.** In Table 1 (first and second rows), we show that if we replace Bitcoin’s puzzle with our weakly nonoutsourcable puzzle, the slowdown for the block verification operation will

<sup>4</sup>Average number of transactions per block: <https://blockchain.info/charts/n-transactions-per-block>

<sup>5</sup>Unless otherwise noted, we conducted our measurements over at least 1000 trials, and omit the error statistic if the standard deviation is within  $\pm 1\%$ .

<sup>6</sup> Average block size: <https://blockchain.info/charts/>

be only 2%. More specifically, while our puzzle verification itself is over a thousand times more expensive than the Bitcoin puzzle, puzzle verification only accounts for a very small percentage of the overall verifier time. Therefore, the overall performance slowdown is insignificant for practical purposes. Likewise, while the size of the ticket in our scheme is almost a hundred times larger than that of Bitcoin, the ticket is a small fraction of the total size of a block when transactions are included.

An adversarial worker can steal a ticket in a marginal amount of time (only one additional hash in expectation, for example, assuming the worker knows at least a third of the Merkle tree branches used during scratch attempts). This cost is insignificant compared to the expected time for solving a puzzle.

### 7.2 Our Strongly Nonoutsourcable Puzzle

We next describe more details of our instantiation and implementation of our strongly nonoutsourcable puzzle, followed by evaluation.

We implemented our puzzle twice, each time using a different NIZK libraries: Pinocchio [36] and Libsnark [5]. Both are implementations of a generic [20] NIZK scheme.<sup>7</sup> Pinocchio includes a compiler that generates an arithmetic circuit from high-level C code, while Libsnark [5] provides a library C++ for composing systems of equations. We used a combination of hand-tuned and generated-from-C-code arithmetic circuits, and developed an adapter for Libsnark to use Pinocchio’s arithmetic circuit files.

In the full online version of our paper we discuss a concrete parameterization of our scheme. We implemented an optimization to improve the parallel running time of the prover. Essentially, we break the overall statement into many substatements, all of which can be proven concurrently; the overall proof consists of a proof for one “Type II” statement, and proofs for some number of “Type I” statements. The number of Type I statements is determined by a parameter  $C$  (smaller  $C$  means a larger number of smaller circuits). We discuss this in more detail in the full online version.

**Performance results.** The prover and verifier costs for our strongly nonoutsourcable implementation are presented in Tables 1 and 2. Each of the bottom four rows of Table 1 and top four rows of Table 2 corresponds to a different setting of the parameter  $C$ , the number of 160-bit blocks (of the underlying ticket) checked by each substatement (smaller values of  $C$  indicate higher degrees of parallelism). The total number of substatements required ( $\#$ ) is reported along with computing time *per circuit* for the prover and verifier. We also report the total verification time over all the statements, as well as the total proof size. Note that our benchmarks are for a sequential verifier, although verification could also be parallelized. The bottom row is for the second type of statement, which does not depend on  $C$ . Due to the longer time required to compute these proofs, the quantities reported are averaged over only three trials.

The reader may immediately notice the vast improvement in prover performance using Libsnark rather than Pinocchio for our implementation; in particular the speedup is much greater than previous reports (i.e. several orders of magnitude vs one order of magnitude) [5]. This is readily explained with reference to the highly sequential nature of our statements, which yields deep and highly-connected circuits. Profiling reveals that the cost of generating our proofs in Pinocchio is dominated by the polynomial interpolation

<sup>7</sup> Libsnark [5] implements several optimizations over the original GGPR [20] scheme. The version we used includes an optimization that turns out to be unsound. [35] Libsnark has since been patched to restore soundness; the patch is reported only to incur an overhead of 0.007% on typical circuits, hence we report our original figures.

Table 1: Estimated puzzle and block verification costs for various schemes

Scheme	Puzzle only				Transactions Included			
	Verif.	Ratio	Size	Ratio	Verif.	Ratio	Size	Ratio
Bitcoin	11.7 $\mu$ s	1	80B	1	0.68s	<b>1</b>	350KB	<b>1</b>
Weak	15.1ms	1.3e3	6.6KB	83	0.70s	<b>1.02</b>	357KB	<b>1.02</b>
Strong ( $C = 4$ )	0.48s	4.1e+04	10.8KB	135	1.16s	1.71	361KB	<b>1.03</b>
Strong ( $C = 3$ )	0.62s	5.3e+04	17.0KB	213	1.30s	1.91	367KB	<b>1.05</b>
Strong ( $C = 2$ )	0.93s	8e+04	23.2KB	290	1.61s	2.37	373KB	<b>1.07</b>
Strong ( $C = 1$ )	1.68s	1.4e+05	29.4KB	368	2.36s	3.47	379KB	<b>1.08</b>

step, which greatly exceeds that of simpler circuits with comparable number of gates [36].

Keeping in mind our goal is to prove it is *plausible* for a worker to produce stolen ticket proof with low latency, we believe it is reasonable to assume that such a worker has access to parallel computing resources. Using Libsnark, the combination of our statement-level parallelism and the parallel SNARK implementation leads to proof times in **under 15 seconds** at the  $C = 2$  setting. Since the average time between puzzle solutions in the Bitcoin network is 10 minutes, this can be a wholly plausible deterrent. At this setting, verification of an entire proof takes under one second. Since approximately 144 Bitcoin puzzle solutions are produced each day, it would take approximately two minutes for a single-threaded verifier to validate a day’s worth of puzzle solutions.

Assuming computational power can be rented at \$1.68 per hour (based on Amazon EC2 prices for the c3.8xlarge used in our trials, which provides 32 cores), it would cost an attacker less than \$10 in total to produce a stolen ticket proof within 20 seconds. This is vastly less than reward for a puzzle solution, which at the current time is approximately \$8,750.

### 7.3 Cryptocurrency Integration

We now discuss several practical aspects of integrating nonoutsourcable puzzles within existing cryptocurrency designs.

**Integrating the puzzle with Bitcoin-like cryptocurrencies.** In our definitions, we indicate that  $\text{GenPuz}(1^\lambda)$  must be a random function that generates a puzzle instance, and in all of our schemes  $\text{GenPuz}(1^\lambda)$  simply returns a uniform random string. However, in the actual Bitcoin protocol, the next puzzle instance is generated by applying a hash function to the solution of the previous puzzle. Our approach is likewise to determine each next puzzle instance from the hash of the previous solution and message,  $\text{puz}' := \mathcal{H}(\text{puz}||m||\text{ticket})$ .

**Further Integration Issues.** In the full online version of our paper, we further discuss how our nonoutsourcable puzzles can be combined with other proposals for complementary properties, such as faster blocks [26, 40], support for lightweight mobile clients [32], and either ASIC-resistance [43] or backward-compatibility with existing mining equipment [17].

## 8. MULTI-TIER BLOCK REWARDS

We want to arrive at a cryptocurrency design that simultaneously discourages centralized mining pools and hosted mining services, yet encourages participation from individual miners and provides similar overall functionality and security as Bitcoin today. To achieve this, there are two major remaining challenges.

**Challenge 1: Lower variance rewards for individual miners.** Individual miners should not have to wait an unreasonable amount of time to earn a Bitcoin reward. Intuitively, we can achieve this by decreasing the average time between blocks, so that rewards are

given out much more frequently. We are constrained, however, by the latency of network propagation, and the time it takes to compute the zero knowledge proofs used in the strongly nonoutsourcable puzzle.

**Challenge 2: Discourage statistical enforcement over time.** Our definition of nonoutsourcable puzzles essentially describes a one-shot game, and ensures that the worker can steal a single puzzle solution from the pool operator and evade detection. However, this definition does not immediately eliminate statistical enforcement techniques over time. For example, pool operators could monitor the output of a hosted service provider and punish it (e.g., through legal prosecution) if it underperforms significantly. Intuitively, we should address this by giving out larger rewards much less frequently, so that the worker can steal solutions over some reasonable timeframe and plausibly claim it was just unlucky.

**Conflicting requirements.** What we need is a reward structure that simultaneously answers the above challenges. Challenge 1 desires paying out small rewards rapidly, whereas Challenge 2 clearly favors paying out large rewards less frequently. Further, to satisfy Challenge 1, if we reduce the inter-block time to the order of seconds, we phase another challenge: since it takes at least 14 seconds to generate the zero-knowledge proof, it would be infeasible for a miner to steal a block this way. The miner could choose to steal the reward using the plaintext option, but since the reward at stake is low, the mining pool could require a small collateral deposit to discourage such blatant stealing.

### 8.1 Proposed Multi-Tier Reward Structure

We propose to satisfy both of these properties by designing a reward structure with multiple possible prizes. Our multi-tier design is inspired by the payoff structure of state lottery games, which often have several consolation prizes as well as large, less frequent jackpots [34, 38]. The effectiveness of such lotteries at encouraging wide participation has long been proven in practice. Our implicit assumption is that miners will tolerate a high-variance payoff overall, as long as they earn *some* reward fairly frequently.

In Table 3, we provide a concrete example of such a multi-tier reward schedule, the rationale for which we discuss below. Each attempt at solving a puzzle yields some chance of winning each of three possible prizes (in contrast with Bitcoin today, in which every block earns the same reward). The prizes are not only associated with different reward values, but also count with varying weight towards the blockchain “difficulty” scoring function. The first two columns indicate the average time between rewards of a given type, along with their relative frequency (adding up to 1). The middle two columns indicate the prize value (in btc), along with the real-time contribution to the total expected reward value (adding up to 1). Note that the overall expected payout rate is the same as in Bitcoin today (25btc every 10 minutes, on average). The final pair of columns indicates the blockchain difficulty weight associated with each reward type along with their expected relative contributions



Table 2: Proof and verification micro-benchmarks for strongly nonoutsourcable puzzles

Type I Statements										
C	#	Gates	Pinocchio [36]			Libsnark [5]				Size
			Prove	Verify	Total	Prove (Single-core)	Prove (Multi-core)	Verify	Total	
1	220	213k	268.2s	11ms	2.42s	16.33s	9.84s	7.6ms	1.672s	29.4KB
2	120	280k	578.4s	11ms	1.32s	20.29s	<b>13.90s</b>	7.7ms	<b>0.924s</b>	23.2KB
3	80	392k	1002.9s	11ms	0.80s	26.92s	17.18s	7.7ms	0.616s	17.0KB
4	60	467k	1242.1s	11ms	0.66s	32.88s	20.71s	7.8ms	0.468s	10.8KB
Type II Statements										
#	Gates	Prove	Verify	Total	Prove (Single-core)	Prove (Multi-core)	Verify	Total	Size	
1	282K	508.5s	10ms	0.01s	19.42s	<b>13.34s</b>	7.8ms	0.008s	<1KB	

towards the overall difficulty of a blockchain (again, adding up to 1).

**Low-value consolation prize: provides low-variance rewards to solo-miners.** The consolation prize is awarded the most frequently (e.g., once every three seconds). The prize is small, less than a tenth of a Bitcoin ( $\approx \$21$  at the time of writing), but the small prizes contribute overall to 70% of the total expected value.

**Medium-value main prize: ensures block confirmations arrive regularly.** The medium-value prizes are necessary to ensure that the log of transactions approximately as quickly and securely as in Bitcoin. They are given out at the same rate as ordinary Bitcoin blocks; they carry a larger reward than the low-value prize, but contribute much less to the overall expected payout. However, the medium-value blocks account for nearly 75% of the total difficulty, and a miner who finds one of these blocks has an average of 7.5 minutes to propagate her solution before it would become stale. The 14 seconds it takes to compute a zero-knowledge proof is relatively small in comparison.

**High-value jackpot: defends against statistical detection of cheating hosted mining services.** The jackpot prize is very rare, and accounts for a small, yet *disproportionately large* fraction of the total expected value. The role of this reward is to engender distrust of hosted mining providers; they would profit greatly by stealing these rewards, but it would be hard to obtain statistical evidence that they have done so.

**Implementing Multi-Tier Rewards.** It is straightforward to implement multi-tier rewards on top of any known scratch-off puzzle constructions (i.e., Bitcoin or our nonoutsourcable constructions). Recall that in these constructions, the critical step of a mining attempt is to compare a hash value to a threshold,  $\mathcal{H}(\text{puz}||\text{ticket}) < T$ , where the threshold  $T = 2^{\lambda-d}$  is parameterized by the difficulty  $d$ . To implement three reward tiers, we introduce two additional thresholds,  $T_{\text{medium}} < T_{\text{high}} < T$ . If the hash value lies between  $T_{\text{high}}$  and  $T$ , then this attempt earns a high-value reward; if it is between  $T_{\text{medium}}$  and  $T_{\text{high}}$  it earns a medium-value reward; and otherwise it earns a low-value reward. These thresholds must be set according to the desired frequency of each reward type.

## 8.2 Economic Analysis of Multi-Tier Rewards

We argue that our proposed reward structure would simultaneously satisfy the necessary properties.

First, our scheme offers small payoff variance. At the time of writing, the overall Bitcoin hashpower is over  $3.5 \cdot 10^{17}$  hashes per second. The most cost effective entry-level Bitcoin ASIC we know of is the  $8.0 \cdot 10^{14}$  ASICMiner BE Tube, which costs \$320. Using this device to solo-mine, the expected time to find a block would be over 8.3 years. However, under our proposed scheme and typical parameters, over a 60 day period, the mining rig mentioned earlier has a better than 98% chance of winning at least one of these prizes.

Table 3: Reward schedules for Bitcoin &amp; our scheme. We give a typical parametrization for the multi-tier reward structure. Parameters can be tuned based on different scenarios.

	Time	Freq	Prize (btc)	(rel)	Weight	(rel)
Bitcoin	10m	1	25	1	1	1
Low	3s	.995	8.8E-2	0.7	1	0.2499
Med	10m	.005	5	0.2	600	0.7497
High	3mo	3.8E-7	3.3E4	0.1	1800	3.4E-4

Next we argue that the payoff structure is also effective at preventing temporal statistical detection. Suppose a large hosted mining provider controls 25% of the network hashpower. Under the original Bitcoin reward structure, it should expect to mine 6574 blocks during a six-month period, and the chance of it mining fewer than 6429 blocks ( $\approx 98\%$ ) is less than one in a thousand. Hence, it could expect steal at most 145 puzzle solutions (worth \$900,000, at today’s price) over this time period before being implicated with high confidence – and even less before generating considerable suspicion. On the other hand, under our proposed scheme, even if the service provider is honest, it has a better than 60% chance of failing to find *any* jackpot during the same time period. Thus if it *does* steal one, it would arouse no suspicion, yet the expected value of this strategy is over \$4.1 million USD.

Finally, we explain that as in the strawman scheme, it is plausible that if miners joined pools requiring small collateral deposits that they would prefer not to defect when they find low-value blocks. However, since the low-value blocks account for only 25% of the blockchain’s total difficulty, even if a coercer influences all the transactions in these blocks, this would be insufficient to enforce a blacklist policy, for example. The difficulty weight of the main prize is high enough to provide ample time to steal the puzzle solution (7.5 minutes) and adequate incentive to do so (thousands of dollars worth). Since these blocks account for 75% of the overall difficulty, pool members would be encouraged at least to steal these blocks. Note that the jackpot blocks count more towards the blockchain weight than the other blocks (so that a miner who finds a jackpot block has a long time window (45 minutes on average) before it becomes stale), but contributes very little to the total difficulty of a chain. This prevents an attacker from revising a large span of history by finding a single jackpot block.

## 9. DISCUSSION

We have proposed a technical countermeasure against the consolidation of mining power that threatens the decentralization of Bitcoin and other cryptocurrencies. Although we have presented a formal definition that captures the security guarantees of our construction and described how it can be practically integrated into a cryptocurrency, due to the difficulty involved in modifying an in-use cryptocurrency (i.e., via a “hard fork” upgrade [7]) and the high

stakes involved in cryptocurrencies generally, the bar for adopting a new design is set very high. Our work provides a significant step in this direction by providing a sound and practical approach to discouraging centralization. However, in order for our solution to be deployed we must provide a thorough and compelling argument that this solution is *fully effective, preferable to all alternatives, and does not conflict with other aspects of the system*. Towards this end, we address several typical objections we have encountered in the past, from academic reviewers and the Bitcoin community alike:

**“Mining pools are good because they lower the variance for solo miners. Therefore, nonoutsourcable puzzles are not well-motivated.”** In Section 1 we describe the severe consequences that can occur due to the concentration of mining power — basically all purported security properties of decentralized cryptocurrencies can be broken if mining coalitions with significant mining power misbehave (and in some cases it may be in their best interest to misbehave [18]). Further, our multi-tier reward system design (see Section 8.2) achieves the best of both worlds, (i.e., ensuring low variance for solo mining as well as discouraging mining coalitions).

**“Can miners still use smart contracts or legal contracts to enforce mining coalitions in spite of the nonoutsourcable puzzles?”** Our definition of nonoutsourcable puzzles prevents the enforcement of contractual mechanisms *including* smart contracts or legal contracts. An enforcement mechanism, such as seizing collateral deposits or legal prosecution, is only effective if it can be applied with few false positives. The worker can steal the puzzle solutions without being held accountable, since the zero-knowledge spending option ensures the worker can spend stolen coins without revealing any evidence that can later be used to implicate it.

**“What about collecting statistical evidence cheating workers?”** Suppose a pool operator monitors the puzzle production rate of a worker over time, to detect if the worker is potentially cheating. One enforcement mechanism might be for the pool to require that a worker submit a deposit to join, such that in case the worker is not producing solutions at the expected rate, the deposit can be confiscated and redistributed.

As mentioned earlier in Section 8.2, our puzzle definition is by nature a one-shot game. Although our nonoutsourcable puzzle alone does not prevent the collection of statistical evidence, in Section 8.2, we argued that *by combining our puzzle with a multi-tier reward system, we effectively make it highly costly or unreliable to accumulate statistical evidence over time*. In particular, a worker can opt to steal only the “jackpot prize” (which happens only infrequently but offers a large reward), while behaving honestly when it finds a “consolation prize” which is of much smaller amount but paid off at a frequent interval. Such an attack cannot be reliably detected within a reasonable of timeframe (e.g., several years).

**“Can coalitions be prevented by other, simpler solutions that do not require zero knowledge proofs?”** The Bitcoin community has put forth two main alternative approaches to ours. First, we could promote the use of P2Pool and other forms of “responsible” mining, so that users can join pools without ceding full control of their resources to a central authority. This has been unsuccessful so far. At the time of writing, P2Pool accounts for less than 2% of the total hashpower; and while some pools support a protocol (called “getblocktemplate”) that allows pool members to see the contents of the blocks they are assigned to work on (and, hence, could leave if they detect the pool is applying some disagreeable policy), the top six pools (which account for more than two thirds of the total hashpower) do not. A second approach is to monitor large pools and

apply social pressure to limit their size. However, pools have been accused of *hiding* their bandwidth to avoid backlash. In any case, we make an analogy to coercion-resistance in electronic voting (see Section 10): although social deterrents to undesired behavior may in some cases be effective, greater confidence can be derived from a technical and economic deterrent.

**“It’s too late to change Bitcoin; and regardless, large miners wouldn’t support this change.”** While we have described our design as a proposed modification to Bitcoin, this is primarily for ease of presentation; our design is also applicable as the basis for a new cryptocurrency, or as a modification to any of the hundreds of Bitcoin-like “altcoins” [7] which compete with Bitcoin (though, at the time of writing, Bitcoin remains far-and-away the most popular). Indeed altcoins have already begun to experiment with (weak) nonoutsourcable puzzles.

It seems unlikely our proposed design will soon be adopted by Bitcoin. Due to the coordination involved and the risk of splintering the network, there is (understandably) considerable political resistance within the Bitcoin community to adopting “hard fork” protocol changes, except in extreme cases [7]. However, such changes have occurred in the past, and could occur again. Though miners are influential and it would be unwise to adopt a new policy that causes them to leave, they aren’t unilaterally responsible for Bitcoin governance [7]; instead stakeholders include payment processors and services, operators of “full nodes” that may not mine, and developers of popular clients. Additionally, as we mentioned in Section 7.3 and explain in detail in our extended online paper, our nonoutsourcable puzzle constructions can be made backward compatible with existing Bitcoin mining equipment, lessening the impact on established miners [17]. Finally, even if our design is not adopted, the mere public knowledge of a viable coalition-resistant design alternative that the community *could* adopt — if necessary — may already serve as a deterrent against large coalitions.

## 10. RELATED WORK

**Computational puzzles.** Moderately hard computational puzzles, often referred to as “proofs of work,” were originally proposed for the purpose of combating email spam [16] (though this application is nowadays generally considered impractical [25]). Most work on computational puzzles has focused instead on “client puzzles,” which can be used to prevent denial-of-service attacks [23]. Recently, several attempts have been made to provide formal security definitions for client puzzles [13, 21, 42].

**Theoretical and economic understanding of Bitcoin.** Although a purely digital currency has been long sought by researchers [11, 12, 14], Bitcoin’s key insight is to frame the problem as a consensus protocol and to provide an incentive for users to participate. Although Bitcoin’s security has initially been proven (informally) in the “honest majority” model [19, 31, 32], this assumption is unsatisfying since it says nothing about whether the incentive scheme indeed leads to an honest majority. An economic analysis of Bitcoin by Kroll et al. [24] showed that honest participation in Bitcoin may be incentive compatible under assumptions such as a homogeneous population of miners and a limited strategy space. More recently, Eyal and Sirer [18] showed that with a more realistic strategy space, when a single player (or coalition) comprises more than a third of the network’s overall strength, the protocol is not incentive compatible (and in fact the threshold is typically much less than one-third, depending on other factors involving network topology). This result underscores the importance of discouraging the formation of Bitcoin mining coalitions.

**Decentralized Mining Pools.** While most mining pools (including the largest) are operated by a central administrator, P2Pool [44] is a successful protocol for *decentralized* mining pools that achieve the desired effect (lower payout variance for participants) that does not require an administrator. It is possible that engineering efforts to improve P2Pool’s performance and usability and public awareness campaigns may steer more users to P2Pool rather than centralized mining pools (at the time of writing<sup>8</sup>, P2Pool accounts for only 1% of the total mining capacity while the two largest pools together account for 49%). However, as P2Pool inherently requires more overhead than a centralized pool, we believe it is wiser to directly discourage coalitions through the built-in reward mechanism.

**Altcoins.** Numerous attempts have been made to tweak the incentive structure by modifying Bitcoin’s underlying puzzle. The most popular alternative, Litecoin<sup>9</sup> uses a script-based [37] puzzle intended to promote the use of general purpose equipment (especially CPUs or GPUs) rather than specialized equipment (e.g., Bitcoin mining ASICs). Another oft-cited goal is to make the puzzle-solving computation have an intrinsically useful side effect (this is discussed, for example, in [24]). To our knowledge, we are the first to suggest deterring mining coalitions as a design goal.

Zerocoin [28], Zerocash [4], and PinocchioCoin [15] focus on making Bitcoin transactions anonymous by introducing a public cryptographic accumulator for mixing coins. Spending a coin involves producing a zero-knowledge proof that a coin has not yet been spent. Although our zero-knowledge proof construction may bear superficial resemblance to this approach, our work addresses a completely different problem.

**Coercion-resistance in Electronic Voting.** The approach we take is inspired by notions of *coercion-resistance* in electronic voting. Vote buying (as well as other forms of coercion) is illegal in all US state and federal elections [22]. While the threat of legal prosecution already poses a deterrent against such behavior, electronic voting schemes have been designed to provide technical countermeasures as well [6, 33]. In short, such schemes ensure that voters are unable to obtain any *receipt* which could demonstrate how they voted to a coercive attacker. We draw an analogy between vote buying and what we call *outsourcing schemes*; analogous to receipt-freeness, the (strongly) nonoutsourcable property prevents a worker from proving to a pool operator how its hashpower is used.

**Most Closely Related Work.** In Permacoin [29], Miller et al. proposed a Bitcoin-like system that achieves decentralized data storage as a useful side effect of mining. As part of their development, they implicitly developed a weakly nonoutsourcable puzzle that deters consolidation of storage capacity. Our Merkle-tree-based weakly nonoutsourcable puzzle construction is directly inspired by the construction in Permacoin. However, Permacoin does not make any attempt to formalize the notion of (weakly) nonoutsourcable puzzles, nor to consider the goal of deterring outsourcing outside the context of archival storage. Our paper provides the first formal treatment of nonoutsourcable puzzles. Additionally, we introduce a new notion of *strongly* nonoutsourcable puzzles, which *repairs a critical flaw in Permacoin* (namely, that weakly nonoutsourcable puzzles provide no deterrence against hosted mining providers with valuable reputations, or against pools who collect collateral deposits from their members).

In independent work, Eyal and Sirer [17] developed a technique for combining a weakly nonoutsourcable puzzle with an arbitrary

<sup>8</sup>According to <https://blockchain.info/pools> retrieved on August 1, 2014

<sup>9</sup><https://litecoin.org/>

scratch-off puzzle, resulting in a puzzle that retains the best properties of the constituents; this can be used, for example, to create a nonoutsourcable puzzle that is backward-compatible with existing Bitcoin mining equipment. We discuss applications of this technique in the full online version of our paper.

## 11. CONCLUSION

The prevalence of Bitcoin mining coalitions (including both mining pools and hosted mining services), which lead to consolidation of power and increased systemic risk to the network, are a result of a built-in design limitation of the Bitcoin puzzle itself – specifically, that it admits an effective coalition enforcement mechanism. To address this, we have proposed formal definitions of *nonoutsourcable* puzzles for which no such enforcement mechanism exists. We have contributed two constructions: a weak nonoutsourcable puzzle provable in the random oracle model, and a generic transformation from any weak nonoutsourcable puzzle to a strong one. The former may already be a sufficient deterrent against mining pools, while the latter thwarts both hosted mining and mining pools. We have implemented both of our techniques and provide performance evaluation results showing these add only a tolerable overhead to the cost of Bitcoin blockchain validation. Overall, we are optimistic that our approach, combined with suitable modifications to the reward structure, could be used to guarantee that participation as an independent individual is the most effective mining strategy.

**Acknowledgments** We thank the readers and reviewers of earlier drafts of this paper. This work was supported in part by NSF awards #0964541, #1223623, and #1518765.

## 12. REFERENCES

- [1] ziftrcoin : A cryptocurrency to enable commerce. Whitepaper.
- [2] James Aspnes, Collin Jackson, and Arvind Krishnamurthy. Exposing computationally-challenged byzantine impostors. *Department of Computer Science, Yale University, New Haven, CT, Tech. Rep.*, 2005.
- [3] Simon Barber, Xavier Boyen, Elaine Shi, and Ersin Uzun. Bitter to better – how to make bitcoin a better currency. In *Financial Cryptography and Data Security*, pages 399–414. Springer, 2012.
- [4] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *Security and Privacy (SP), 2014 IEEE Symposium on*. IEEE, 2014.
- [5] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 781–796, San Diego, CA, August 2014. USENIX Association.
- [6] Josh Benaloh and Dwight Tuinstra. Receipt-free secret-ballot elections. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 544–553. ACM, 1994.
- [7] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua Kroll, and Edward W. Felten. Research perspectives on bitcoin and second-generation digital currencies. In *2015 IEEE Symposium on Security and Privacy*. IEEE, 2015.
- [8] Joseph Bonneau, Arvind Narayanan, Andrew Miller, Jeremy Clark, Joshua A Kroll, and Edward W Felten. Mixcoin:

- Anonymity for bitcoin with accountable mixes. In *Financial Cryptography and Data Security*, pages 486–504. Springer, 2014.
- [9] Danny Bradbury. Alydian targets big ticket miners with terahash hosting. <http://www.coindesk.com/alydian-targets-big-ticket-miners-with-terahash-hosting/>, August 2013.
- [10] Vitalik Buterin. Bitcoin network shaken by blockchain fork. <http://bitcoinmagazine.com/3668/bitcoin-network-shaken-by-blockchain-fork/>, 2013.
- [11] David Chaum. Blind signatures for untraceable payments. In *Crypto*, volume 82, pages 199–203, 1982.
- [12] David Chaum, Amos Fiat, and Moni Naor. Untraceable electronic cash. In *Advances in Cryptology—CRYPTO’88*, pages 319–327. Springer, 1990.
- [13] Liqun Chen, Paul Morrissey, Nigel P Smart, and Bogdan Warinschi. Security notions and generic constructions for client puzzles. In *Advances in Cryptology—ASIACRYPT 2009*, pages 505–523. Springer, 2009.
- [14] Wei Dai. b-money. <http://www.weidai.com/bmoney.txt>, 1998.
- [15] George Danezis, Cedric Fournet, Markulf Kohlweiss, and Bryan Parno. Pinocchio coin: building zerocoin from a succinct pairing-based proof system. In *Proceedings of the First ACM workshop on Language support for privacy-enhancing technologies*, pages 27–30. ACM, 2013.
- [16] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *CRYPTO*, 1993.
- [17] Ittay Eyal and Emin Gün Sirer. How to disincentivize large bitcoin mining pools. Blog post: <http://hackingdistributed.com/2014/06/18/how-to-disincentivize-large-bitcoin-mining-pools/>, June 2014.
- [18] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*, pages 436–454. Springer, 2014.
- [19] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology—EUROCRYPT 2015*, pages 281–310. Springer, 2015.
- [20] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct nizks without pcps. In *Advances in Cryptology—EUROCRYPT 2013*, pages 626–645. Springer, 2013.
- [21] Bogdan Groza and Bogdan Warinschi. Cryptographic puzzles and dos resilience, revisited. *Designs, Codes and Cryptography*, pages 1–31, 2013.
- [22] Richard L Hasen. Vote buying. *California Law Review*, pages 1323–1371, 2000.
- [23] Ari Juels and John G Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*, volume 99, pages 151–165, 1999.
- [24] Joshua A Kroll, Ian C Davey, and Edward W Felten. The economics of bitcoin mining or, bitcoin in the presence of adversaries. *WEIS*, 2013.
- [25] Ben Laurie and Richard Clayton. Proof-of-work proves not to work; version 0.2. In *Workshop on Economics and Information, Security*, 2004.
- [26] Yoad Lewenberg, Yonatan Sompolinsky, and Aviv Zohar. Inclusive block chain protocols. *Financial Cryptography and Data Security*. Springer, 2015.
- [27] Jon Matonis. The bitcoin mining arms race: Ghash.io and the 51% issue. <http://www.coindesk.com/bitcoin-mining-detente-ghash-io-51-issue/>, July 2014.
- [28] Ian Miers, Christina Garman, Matthew Green, and Aviel D Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pages 397–411. IEEE, 2013.
- [29] Andrew Miller, Ari Juels, Elaine Shi, Bryan Parno, and Jonathan Katz. Permacoin: Repurposing bitcoin work for long-term data preservation. In *IEEE Symposium on Security and Privacy*, 2014.
- [30] Andrew Miller, Ahmed Kosba, Elaine Shi, and Jonathan Katz. Nonoutsourcable scratch-off puzzles to discourage bitcoin mining coalitions. (full version). [http://cs.umd.edu/~amiller/nonoutsourcable\\_full.pdf](http://cs.umd.edu/~amiller/nonoutsourcable_full.pdf), 2015.
- [31] Andrew Miller and Joseph J. LaViola, Jr. Anonymous byzantine consensus from moderately-hard puzzles: A model for bitcoin. UCF Tech Report. CS-TR-14-01.
- [32] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://bitcon.org/bitcoin.pdf>, 2008.
- [33] Valtteri Niemi and Ari Renvall. How to prevent buying of votes in computer elections. In *Advances in Cryptology—ASIACRYPT’94*, pages 164–170. Springer, 1995.
- [34] Emily Oster. Are all lotteries regressive? evidence from the powerball. *National Tax Journal*, June, 2004.
- [35] Bryan Parno. A note on the unsoundness of vntinyram’s snark. Cryptology ePrint Archive, Report 2015/437, 2015. <http://eprint.iacr.org/>.
- [36] Bryan Parno, Craig Gentry, Jon Howell, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *IEEE Symposium on Security and Privacy*, pages 238–252, 2013.
- [37] Colin Percival and Simon Josefsson. The script password-based key derivation function. 2012.
- [38] John Quiggin. On the optimal design of lotteries. *Economica*, 58(229):1–16, 1991.
- [39] Meni Rosenfeld. Analysis of bitcoin pooled mining reward systems. *arXiv preprint arXiv:1112.4980*, 2011.
- [40] Yonatan Sompolinsky and Aviv Zohar. Secure high-rate transaction processing in bitcoin. *Financial Cryptography and Data Security*. Springer, 2015.
- [41] Spreadcoin. <http://spreadcoin.net/files/SpreadCoin-WhitePaper.pdf>, October 2014.
- [42] Douglas Stebila, Lakshmi Kuppusamy, Jothi Rangasamy, Colin Boyd, and Juan Gonzalez Nieto. Stronger difficulty notions for client puzzles and denial-of-service-resistant protocols. In *Topics in Cryptology—CT-RSA 2011*, pages 284–301. Springer, 2011.
- [43] John Tromp. Cuckoo cycle: a new memory-hard proof-of-work system. *Bitcoin Research Workshop*, 2015.
- [44] Forrest Voight. p2pool: Decentralized, dos-resistant, hop-proof pool. <https://bitcointalk.org/index.php?topic=18313>, June 2011.