

REM: Resource-Efficient Mining for Blockchains

Fan Zhang^{*,§}

fanz@cs.cornell.edu

Ittay Eyal^{*,§}

ittay.eyal@cornell.edu

Robert Escriva^{*}

escriva@cs.cornell.edu

Ari Juels^{†,§}

juels@cornell.edu

Robbert van Renesse^{*,§}

rvr@cs.cornell.edu

^{*}Cornell University [†]Cornell Tech, Jacobs Institute

[§]Initiative for CryptoCurrencies & Contracts

Abstract

Blockchains show promise as potential infrastructure for financial transaction systems. The security of blockchains today, however, relies critically on Proof-of-Work (PoW), which forces participants to waste computational resources.

We present REM (Resource-Efficient Mining), a new blockchain mining framework that uses *trusted hardware* (Intel SGX). REM achieves security guarantees similar to PoW, but leverages the partially decentralized trust model inherent in SGX to achieve a fraction of the waste of PoW. Its key idea, *Proof-of-Useful-Work (PoUW)*, involves miners providing trustworthy reporting on CPU cycles they devote to *inherently useful workloads*. REM flexibly allows any entity to create a useful workload. REM ensures the trustworthiness of these workloads by means of a novel scheme of *hierarchical attestations* that may be of independent interest.

To address the risk of compromised SGX CPUs, we develop a statistics-based formal security framework, also relevant to other trusted-hardware-based approaches such as Intel’s Proof of Elapsed Time (PoET). We show through economic analysis that REM achieves less waste than PoET and variant schemes.

We implement REM and, as an example application, swap it into the consensus layer of Bitcoin core. The result is the first full implementation of an SGX-based blockchain. We experiment with four example applications as useful workloads for our implementation of REM, and report a computational overhead of 5 – 15%.

1 Introduction

Despite their imperfections [26, 35, 37, 65, 71], blockchains [38, 64, 67] have attracted the interest of the financial and technology industries [16, 25, 34, 47, 69, 74] as a way to build a transaction systems with distributed trust. Among the impediments to widespread adoption of decentralized or “permission-

less” blockchains, however, one is fundamental. Proofs-of-Work (PoWs) in blockchains are wasteful.

PoWs are nonetheless the most robust solution today to two fundamental problems in decentralized cryptocurrency design: How to select consensus leaders and how to apportion rewards fairly among participants. A participant in a PoW system, known as a *miner*, can only lead consensus rounds in proportion to the amount of computation she invests in the system. This prevents an attacker from gaining majority power by cheaply masquerading as multiple machines. The cost, however, is the above-mentioned waste. PoWs serve no useful purpose beyond consensus and incur huge monetary and environmental costs. Today the Bitcoin network uses more electricity than produced by a nuclear reactor, and is projected to consume as much as Denmark by 2020 [30].

We propose a solution to the problem of such waste in a novel block-mining system called REM. Nodes using REM replace PoW’s wasted effort with useful effort of a form that we call *Proof of Useful Work (PoUW)*. In a PoUW system, users can utilize their CPUs for any desired workload, and can simultaneously contribute their work towards securing a blockchain.

There have been several attempts to construct cryptocurrencies that recycle PoW by creating a resource useful for an external goal, but they have serious limitations. Existing schemes rely on esoteric resources [55], have low recycling rates [63], or are centralized [41]. Other consensus approaches, e.g., BFT or Proof of Stake, are in principle waste-free, but restrict consensus participation or have notable security limitations.

Intel [47] recently introduced a new approach to eliminating waste in distributed consensus protocols that relies instead on trusted hardware, specifically a new instruction set architecture extension in Intel CPUs called *Software Guard Extensions (SGX)*. SGX permits the execution of trustworthy code in an isolated, tamper-free environment, and can prove remotely that outputs represent the result of such execution. Leveraging this capability, Intel’s proposed Proof of Elapsed Time (PoET) is an in-

novative system with an elegant and simple underlying idea. A miner runs a trustworthy piece of code that idles for a randomly determined interval of time. The miner with the first code to awake leads the consensus round and receives a reward. PoET thus promises energy-waste-free decentralized consensus with security predicated on the tamper-proof features of SGX. PoET operates in a *partially-decentralized model*, involving limited involvement of an authority (Intel), as we explain below.

Unfortunately, despite its promise, as we show in this paper, PoET presents two notable technical challenges. First, in the basic version of PoET, an attacker that can corrupt a single SGX-enabled node can win every consensus round and break the system completely. We call this the *broken chip* problem. Second, miners in PoET have a financial incentive to power mining rigs with cheap, outmoded SGX-enabled CPUs used solely for mining. The result is exactly the waste that PoET seeks to avoid. We call this the *stale chip* problem.

REM addresses both the stale and broken chip problems. Like PoET, REM operates in a partially decentralized model: It relies on SGX to prove that miners are generating valid PoUWs. REM, however, avoids PoET’s stale chip problem by substituting PoUWs for idle CPU time, disincentivizing the use of outmoded chips for mining. Miners in a PoUW system are thus entities that use or outsource SGX CPUs for computationally intensive workloads, such as scientific experiments, pharmaceutical discovery, etc. All miners can concurrently mine for a blockchain while REM gives them the flexibility to use their CPUs for *any desired workload*.

We present a detailed financial analysis to show that PoUW successfully addresses the stale chip problem. We provide a taxonomy of different schemes, including PoW, PoET, novel PoET variants, and PoUW. We analyze these schemes in a model where agents choose how to invest capital and operational funds in mining and how much of such investment to make. We show that the PoUW in REM not only avoids the stale chip problem, but yields the smallest overall amount of mining waste. Moreover, we describe how small changes to the SGX feature set could enable even more efficient solutions.

Unlike PoET, REM addresses the broken chip problem. Otherwise, compromised SGX-enabled CPUs would allow an attacker to generate PoUWs at will, and both unfairly accrete revenue and disrupt the security of the blockchain [29, 75, 79]. Intel has sought to address the broken chip problem in PoET using a statistical-testing approach, but published details are lacking, as appears to be a rigorous analytic framework. For REM, we set forth a rigorous statistical testing framework for mitigating the damage of broken chips, provide analytic security bounds, and empirically assess its performance given the volatility of mining populations in real-world

cryptocurrencies. Our results also apply to PoET.

A further challenge arises in REM due to the feature that miners may choose their own PoUWs workloads. It is necessary to ensure that miner-specified mining applications running in SGX accurately report their computational effort. Unfortunately SGX lacks secure access to performance counters. REM thus includes a *hierarchical attestation* mechanism that uses SGX to attest to compilation of workloads with valid instrumentation. Our techniques, which combine static and dynamic program analysis techniques, are of independent interest.

We have implemented a complete version of REM, encompassing the toolchain that instruments tasks to produce PoUWs, compliance checking code, and a REM blockchain client. As an example use, we swap REM in for the PoW in Bitcoin core. As far as we are aware, ours is the first full implementation of an SGX-backed blockchain. (Intel’s Sawtooth Lake, which includes PoET, is implemented only as a simulation.) Our implementation supports trustworthy compilation of any desired workload. As examples, we experiment with four REM workloads, including a commonly-used protein-folding application and a machine learning application. The resulting overhead is about 5 – 15%, confirming the practicality of REM’s methodology and implementation.

Paper organization

The paper is organized as follows: Section 2 provides background on proof-of-work and Intel SGX. We then proceed to describe the contributions of this work:

- *PoUW and REM*, a low-waste alternative to PoW that maintains PoW’s security properties (§3).
- A *broken-chip countermeasure* consisting of a rigorous statistical testing framework that mitigates the impact of broken chips (§4 and Appendix §A).
- A *methodology for trustworthy performance instrumentation* of SGX applications using a combination of static and dynamic program analysis and SGX-backed trusted compilation (§5).
- *Design and full implementation of REM* as a resource-efficient PoUW mining system with automatic tools for compiling arbitrary code to a PoUW-compliant module. Ours is the first full implementation of an SGX-backed blockchain protocol (§5).
- A *model of consensus-algorithm resource consumption* that we use to compare the waste associated with various mining schemes. We overview the model and issues with previous schemes (§6) and defer the details to the appendix (§B and §C).

We discuss related work in §7 and conclude in §8.

2 Background

2.1 Blockchains

Blockchain protocols allow a distributed set of participants, called miners, to reach a form of consensus called Nakamoto consensus. Such consensus yields an ordered list of transactions. Roughly speaking, the process is as follows. Miners collect cryptographically signed transactions from system users. They validate the transactions' signatures and generate blocks that contain these transactions plus a pointer to a parent block. The result is a chain of blocks called (imaginatively) a blockchain.

Each miner, as it generates a block, gets to choose the block's contents, specifically which transactions will be included and in what order. System participants are connected by a peer-to-peer network that propagates transactions and blocks. Occasionally, two or more miners might nearly simultaneously generate blocks that have the same parent, forming two branches in the blockchain and breaking its single-chain structure. Thus a mechanism is used to choose which branch to extend, most simply, the longest chain available [64].¹

An attacker could naturally seek to generate blocks faster than everyone else, forming the longest chain and unilaterally choosing block contents. To prevent such an attack, a block is regarded as valid only if it contains proof that its creator has performed a certain amount of work, a proof known as a *Proof of Work* (PoW).

A PoW takes the form of a *cryptopuzzle*: In most cryptocurrencies, a miner must change an input (nonce) in the block until a cryptographic hash of the block is smaller than a predetermined threshold. The security properties of hash functions force a miner to test nonces by brute force until a satisfying block is found. Such a block constitutes a solution to the cryptopuzzle and is itself a PoW. Various hash functions are used in practice. Each type puts different load on the processor and memory of a miner's computing device [64, 63, 78].

The process of mining determines an exponentially distributed interval of time between the blocks of an individual miner, and, by extension, between blocks in the blockchain. The expected amount of work to solve a cryptopuzzle, known as its *difficulty*, is set according to a deterministic algorithm that seeks to enforce a static expected rate of block production by miners (e.g., 10 minute block intervals in Bitcoin). An individual miner thus generates blocks at a rate that is proportional to its *mining power*, its hashrate as a fraction of that in the entire population of miners. Compensation to miners is granted per block generated, leading to an expected miner revenue that is proportional to the miner's

¹There are alternatives to this protocol [37, 58, 73, 78], however the differences are immaterial to our exploration here.

hashrate.

As the mining power is invested in a cryptocurrency grows, the cryptocurrency's cryptopuzzle difficulty rises to keep the block generation rate stable. When compensation is sufficiently high, it is worthwhile for a large number of participants to mine, leading to a high difficulty requirement. This, in turn, makes it difficult for an attacker to mine a large enough fraction of blocks to perform a significant attack.

PoW properties. The necessary properties for PoW to support consensus in a blockchain, i.e., resist adversarial control, are as follows. First, a PoW must be tied to a unique block, and be valid only for that block. Otherwise, a miner can generate conflicting blocks, allowing for a variety of attacks. A PoW should be moderately hard [15], and its difficulty should be accurately tunable so that the blockchain protocol can automatically tune the expected block intervals. Validation of PoWs, on the other hand, should be as efficient as possible, given that it is performed by the whole network. (In most cryptocurrencies today, it requires just a single hash.) It should also be possible to perform by any entity with access to the blockchain — If the proofs or data needed for validation are made selectively available by a single entity, for instance, that entity becomes a central point of control and failure.²

2.2 SGX

Intel Software Guard Extensions (SGX) [45, 46, 48, 49, 13, 43, 62] is a set of new instructions available on recent-model Intel CPUs that confers hardware protections on user-level code. SGX enables process execution in a Trusted Execution Environment (TEE), and specifically in SGX in a protected address space known as an *enclave*. An enclave protects the confidentiality and the integrity of the process from certain forms of hardware attack and other processes on the same host, including privileged processes like operating systems.

An enclave can read and write memory outside the enclave region as a form of inter-process communication, but no other process can access enclave memory. Thus the isolated execution in SGX may be viewed in terms of an ideal model in which a process is guaranteed to execute correctly and with perfect confidentiality, but relies on a (potentially malicious) operating system for supporting services such as I/O, etc. This model is a simplification: SGX is known to expose some internal enclave

² The Bitcoin protocol is expected to soon allow for the so-called segregated witness architecture [22, 60]. Then, transaction signatures (witnesses) are kept in a data structure that is technically separate (segregated) from the blockchain data structure. Despite this separation of data structures, the data in both must be propagated to allow for distributed validation.

state to the OS [79]. Our basic security model assumes ideal isolated execution, but as we detail in Section 4, we have baked a defense against compromised SGX CPUs into REM.

Attestation SGX allows a remote system to verify the software running in an enclave and communicate securely with it. When an enclave is created, the CPU produces a hash of its initial state known as a *measurement*. The software in the enclave may, at a later time, request a report which includes a measurement and supplementary data provided by the process. The report is digitally signed using a hardware-protected key to produce a proof that the measured software is running in an SGX-protected enclave. This proof, known as a *quote*, is part of an *attestation* can be verified by a remote system.

SGX signs quotes in attestations using a group signature scheme called *Enhanced Privacy ID* or *EPID* [72]. This choice of primitive is significant in our design of REM, as Intel made the design choice that attestations can only be verified by accessing Intel’s Attestation Service (IAS) [50], a public Web service maintained by Intel whose primary responsibility is to verify attestations upon request.

REM uses attestations as proofs for new blocks, so miners need to access IAS to verify blocks. The current way in which IAS works forces miners to access IAS on every single verification, adding an undesirable round-trip time to and from Intel’s server to the block verification time. This overhead, however, is not inherent, and is due only to a particular design choice by Intel. As we suggest in Section 5.4, a simple modification, to the IAS protocol, which Intel is currently testing, can eliminate this overhead entirely.

Randomness As operating systems sit outside of the trusted computing base (TCB) of SGX, OS-served random functions such as `srand` and `rand` are not accessible to enclaves. SGX instead provides a hardware-protected random number generator (RNG) using the `rdrand` instruction. REM relies on the SGX RNG.

3 Overview of PoUW and REM

The basic idea of PoUW, and thus REM, is to replace the wasteful computation of PoW with arbitrary useful computation. A miner proves that a certain amount of useful work has been dedicated to a specific branch of the blockchain. Intuitively, due to the value of the useful work outside of the context of the blockchain supported by REM, the hardware and power are well spent, and there is no waste. A comprehensive analysis of the waste is deferred to Appendix B. Here we describe the security

model of REM and then give an overview of its system mechanics.

3.1 Security Model

A PoW solution embodies a statistical proof of an effort spent by the miner. With PoUW, however, a miner herself reports her own effort. The rational miner’s incentive is to lie, report more work than she actually performed, and monopolize the blockchain. In PoUW / REM, use of a TEE — Intel SGX in particular — prevents such attacks and enforces correct reporting of work. The resulting trust model is starkly different from that in traditional PoW.

PoET introduced, and we similarly use in REM, a *partially decentralized* blockchain model. The blockchain is *permissionless*, i.e., any entity can participate as a miner, as in a fully decentralized blockchain such as Bitcoin. It is only partially decentralized, though, in that it relies for security on two key assumptions about the hardware manufacturer’s behavior.

First, we must assume that Intel correctly manages identities, specifically that it assigns a signing key (used for attestations) only to a valid CPU. It follows that Intel does not forge attestations and thus mining work. Such forgery, if detected in any context, would undermine the company’s reputation and the perceived utility of SGX, costing far more than potential blockchain revenue. Second, we assume that Intel does not blacklist valid nodes in the network, rendering their attestations invalid when the IAS is queried. Such misbehavior would be publicly visible and similarly damaging to Intel if unjustified.

Even assuming trustworthy manufacturer behavior, though, a limited number of individual CPUs might be physically or otherwise compromised by a highly resourced adversary (or adversaries). Our trust model assumes the possibility of such an adversary and makes the strong assumption that she can learn the attestation (EPID signing) key for compromised machines and thus can issue arbitrary attestations for those machines. In particular, as we shall see, she can falsify random number generation and lie about work performed in REM.

Even this strong adversary, though, does have a key limitation: As signing keys are issued by the manufacturer, and given our first assumption above, it is not possible for an adversary to forge identities. We further assume that the signatures are linkable. In SGX, the EPID signature scheme for attestations has a linkable (pseudonymous) mode [50, 13, 72], which permits anyone to determine whether two signatures were generated by the same CPU. As a result, even a compromised node cannot masquerade as multiple nodes.

Outside the REM security model It is important to note that REM is a *consensus framework*, i.e., a means to generate blocks, not a *cryptocurrency*. REM can be integrated into a cryptocurrency, as we show by swapping it into the Bitcoin consensus layer. As REM has roughly the same exponentially distributed block-production interval, such integration need not change security properties above the consensus layer. For example, fork resolution, transaction validation, block propagation, etc., *remain the same in a REM-backed blockchain as in a PoW-based one*. Thus we do not expand the discussion of the security issues relevant to those elements in the REM security model.

3.2 REM overview

Figure 1 presents an architectural overview of REM.

There are three types of entities in the ecosystem of REM: A *blockchain agent*, one or more *REM miners*, and one or more *useful work clients*.

The useful work clients supply useful workloads to REM miners in the form of *PoUW tasks*, each of which encompass a *PoUW enclave* and some *input*. Any SGX-compliant program can be transformed into a PoUW enclave using the toolchain we developed. Note that a PoUW enclave has to conform to certain security requirements. The most important is that it meters effort correctly, something that can be efficiently verified by a compliance checker and a novel technique we introduce called *hierarchical attestation*. We refer readers to §5.2 and §5.3 for details.

The blockchain agent collects transactions and generates a block template, a block lacking the proof of useful work (PoUW). As detailed later, a REM miner will attach the required PoUW and return it to the agent. The agent then publishes the full block to the P2P network, making it part of the blockchain and receiving the corresponding reward.

A miner takes as input a block template and a PoUW task to produce PoUWs. It launches the PoUW enclave in SGX with the prescribed input and block template. Once the PoUW task halts, its results are returned to the useful work client. The PoUW enclave meters work performed by the miner and declares whether the mining effort is successful and results in a block. Effort is metered on a per-instruction basis. The PoUW enclave randomly determines whether the work results in a block by treating each instruction as a Bernoulli trial. Thus mining times are distributed in much the same manner as in proof-of-work systems. While in, e.g., Bitcoin, effort is measured in terms of executed hashes, in REM, it is the number of executed useful-work instructions. Intuitively, REM may be viewed as *simulating* the distribution of block-mining intervals associated with PoW, but

REM does so with PoUW, and thus eliminates wasted CPU effort.

When a PoUW enclave determines that a block has been successfully mined, it produces a PoUW, which consists of two parts: an SGX-generated attestation demonstrating the PoUW enclave’s *compliance* with REM and another attestation that a block was successfully mined by the PoUW enclave at a given *difficulty parameter*. The blockchain agent concatenates the PoUW to the block template, forming a full block, and publishes it to the network.

When a blockchain participant verifies a fresh block received on the blockchain network, in addition to verifying higher-layer properties (e.g., in a cryptocurrency such as Bitcoin, that transactions, previous block references, etc., are valid), the participant verifies the attestations in the associated PoUW.

Intel’s PoET scheme looks similar to REM in that its enclave randomly determines block intervals and attests to block production. PoET, however, lacks the production of useful work, an essential ingredient, as we explain later in the paper. We now discuss our strategy in REM for handling compromised nodes.

4 Tolerating Compromised SGX Nodes

SGX does not achieve perfect enclave isolation. While no real practical attack is known, researchers have demonstrated potentially dangerous side-channel attacks against applications [79] and even expressed concerns about whether an attestation key might be extracted [29].

Therefore, even if we assume SGX chips are manufactured in a secure fashion, some number of individual instances could be broken by well-resourced adversaries. A single compromised node could be catastrophic to an SGX-based cryptocurrency, allowing an adversary to create blocks at will and perform majority attacks on the blockchain. While she could not spend other people’s money, which would require access to their private keys, she could perform denial-of-service attacks, selectively drop transactions, or charge excessive transaction fees.

In principle, a broken attestation key can be revoked through the Intel Attestation Service (IAS), but this can only happen if the break is detected to begin with. Consequently, Intel has explored ways of detecting SGX compromise in PoET [9] by statistically testing for implausibly frequent mining by a given node (using a “z-test”). Details are lacking in published materials, however, and a rigorous analytic framework seems to be needed.

For REM, we explore compromise detection within a rigorous definitional and analytic framework. The centerpiece is what we call a *block-acceptance policy*, a flexibly defined rule that determines whether a proposed block in a blockchain is legitimate. As we show, defining




Figure 1: Architecture overview of REM

and analyzing policies rigorously is challenging, but we provide strong analytical and empirical evidence that a relatively simple statistical-testing policy (which we denote P_{stat}) can achieve good results. P_{stat} both limits an adversary’s ability to harvest blocks unfairly and minimizes erroneous rejection of honestly mined blocks.

4.1 Threat Model and Definitions

4.1.1 Basic notation

To model block-acceptance policies, let $M = \{m_1, \dots, m_n\}$ be the set of all miners, which we assume to be static. (Miners can join and leave the system; M includes all potential miners.) An adversary \mathcal{A} controls a static subset $M_{\mathcal{A}} \in M$, where $|M_{\mathcal{A}}| = k$. $\text{rate}(m_i)$ specifies the *mining rate* of m_i , the number of mining operations per unit time it performs.

We define a candidate *block* to be a tuple $B = (t, m, d)$, where t is a timestamp, $m \in M$ the identity of the CPU that mines the block, and d is the block difficulty. Difficulty d is defined as the win probability per mining operation in the underlying consensus protocol (e.g. a hash in Bitcoin, a unit time of sleep in PoET, an instruction in PoUW). \mathcal{B} denotes the set of possible blocks B .

A *blockchain* is an ordered sequence of blocks. At time τ , blockchain $C(\tau)$ is a sequence of accepted blocks $C(\tau) = \{B_1, B_2, \dots, B_n\}$ for some n . We drop τ where its clear from context. We let $r(\tau)$ denote the number of rejected blocks of honest miners, i.e., miners in $M - M_{\mathcal{A}}$, in the history of $C(\tau)$. (Of course, this is not and indeed cannot be recorded in a real blockchain system.) Let \mathcal{C} be the space of all possible blockchains C . Let C_m denote blockchain C restricted to blocks mined by miner $m \in M$.

In REM, a blockchain-acceptance policy is used to determine whether a block appears to come from a legitimate miner (CPU that hasn’t been compromised).

Definition 1. (*Blockchain-Acceptance Policy*) A *blockchain-acceptance policy* (or simply *policy*) $P: \mathcal{C} \times \mathcal{B} \rightarrow \{\text{reject}, \text{accept}\}$ is a function that takes as input a blockchain and a proposed block, and outputs whether the proposed block is legitimate.

4.1.2 Security and efficiency definitions

We model the consensus algorithm for the blockchain, the adversary \mathcal{A} , and honest miners respectively as (ideal) programs $\text{prog}_{\text{chain}}$, $\text{prog}_{\mathcal{A}}$, and prog_m . Together, they define what we call a *security game* $S(P)$ for a particular policy P .

We define security games and their constituent programs formally in Appendix A.2. Where clear from context in what follows, we use the notation S , rather than $S(P)$, i.e., omit P .

A security game S may itself be viewed as a probabilistic algorithm. Thus we may treat the blockchain resulting from execution of S for interval of time τ as a random variable $C_S(\tau)$.

Normalizing the revenue from mining a block to 1, we define the *payoff* for a miner m for a given blockchain C as $\pi_m(C) = |C_m|$.

An adversary \mathcal{A} seeks to maximize payoffs for its miners, as reflected in the following definition:

Definition 2. (*Advantage of \mathcal{A}*). For a given security game S , the *advantage of \mathcal{A} for time τ* is:

$$\text{Adv}_{\mathcal{A}}^S(\tau) = \frac{\mathbb{E}[\pi_{\hat{m}}(C_S(\tau))]}{\max_{m_j \in M - M_{\mathcal{A}}} \mathbb{E}[\pi_{m_j}(C_S(\tau))]},$$

for any $\hat{m} \in M_{\mathcal{A}}$. Note that $\mathbb{E}[\pi_{\hat{m}}(C_S(\tau))]$ is equal for all such \hat{m} , as they all use strategy $\Sigma_{\mathcal{A}}$ and can emit blocks as frequently as desired (ignoring $\text{rate}(\hat{m})$).

A policy that keeps $\text{Adv}_{\mathcal{A}}^S(\tau)$ low is desirable, but there’s a trade-off. A policy that rejects too many policies incurs high *waste*, meaning that it rejects many blocks from honest miners. We define waste as follows.

Definition 3. (*Waste of a policy*). For a given blockchain $C(\tau) = \{(B_1, B_2, \dots, B_n)\}$, the *waste* is defined as

$$\text{Waste}(C(\tau)) = \frac{r(\tau)}{n + r(\tau)}.$$

For security game S , the *waste at time τ* is defined as

$$\text{Waste}^S(\tau) = \mathbb{E}[\text{Waste}(C_S(\tau))].$$

```

 $P_{\text{stat}}^{\alpha, \text{rate}_{\text{best}}}(C, B):$ 
  parse  $B \rightarrow (\tau, m, d)$ 
  if  $|C_m| > F^{-1}(1 - \alpha, d\tau(\text{rate}_{\text{best}})):$ 
    output reject
  else
    output accept

```

Figure 2: $P_{\text{stat}}^{\alpha} \cdot F^{-1}(\cdot, \lambda)$ is the quantile function for Poisson distribution with rate λ .

Our exploration of policies focuses critically on the trade-offs between low $\text{Adv}_{\mathcal{A}}^S(\tau)$ and low $\text{Waste}^S(\tau)$. To illustrate the issue, we give a simple example in Appendix A.3 of a policy that allows any CPU to mine only one block over its lifetime. As $\tau \rightarrow \infty$, it achieves the optimal $\text{Adv}_{\mathcal{A}}^S(\tau) = 1$, but at the cost of $\text{Waste}^S(\tau) = 1$, i.e., 100% waste.

4.2 The REM policy: P_{stat}

REM makes use of a statistical-testing-based policy that we denote by P_{stat} . P_{stat} is compatible not just with PoUW, but also with PoET and potentially other SGX-based mining variants.

There are two parts to P_{stat} . First, P_{stat} estimates the rate of the fastest honest miner(s) (fastest CPU type), denoted by $\text{rate}_{\text{best}} = \max_{m \in \mathcal{M} - \mathcal{M}_{\mathcal{A}}} \text{rate}(m)$. There are various ways to accomplish this; a simple one would be to have an authority (e.g., Intel) publish specs on its fastest CPUs' performance. (In PoET, mining times are uniform, so $\text{rate}_{\text{best}}$ is just a system parameter.) We describe an empirical approach to estimating $\text{rate}_{\text{best}}$ in REM in Appendix A.1.

Given an estimate of $\text{rate}_{\text{best}}$, P_{stat} tests submitted blocks statistically to determine *whether a miner is mining blocks too quickly and may thus be compromised*. The basic principle is simple: On receiving a block B from miner m , P_{stat} tests the null hypothesis

$$H_0 = \{\text{rate}(m) \leq \text{rate}_{\text{best}}\}.$$

We use $|C_m(\tau)|$, the number of blocks mined by m at time τ , as the test statistic. Under H_0 , $|C_m|$ should obey a Poisson distribution with rate $d\tau(\text{rate}_{\text{best}})$, denoted as $\text{Pois}[d\tau(\text{rate}_{\text{best}})]$. P_{stat} rejects H_0 if $|C_m|$ is greater than the $(1 - \alpha)$ -quantile of the Poisson distribution. The false rejection rate for a single test is therefore at most α . We specify P_{stat} (for a given $\text{rate}_{\text{best}}$) in Figure 2.

An important property that differentiates P_{stat} from canonical statistical tests is that P_{stat} repeatedly applies a given statistical test to an accumulating history of samples. *The statistical dependency between samples makes the analysis non-trivial, as we shall show.*

4.3 Analysis of P_{stat}

We now analyze the average-case and worst-case waste and adversarial advantage of P_{stat} . We assume for simplicity that $\text{rate}_{\text{best}}$ is accurately estimated. We remove this assumption in the worst-case analysis below. We also assume that the difficulty $d(t)$ is stationary over the period of observation.

Waste Under P_{stat} , a miner generates blocks according to a Poisson process; whether a block is accepted or rejected depends on whether the miner has generated more blocks than a time-dependent threshold. This process is obviously not memoryless and thus not directly representable as a Markov process. We can, however, achieve a close approximation using a discrete-time Markov chain. Indeed, as we show, we can represent waste in P_{stat} using a discrete-time Markov chain that is *periodically identical* to the process it models, meaning that its expected waste is identical at any time $n\tau$, for $n \in \mathbb{Z}^+$ and τ a model parameter specified below. This Markov chain has a stationary distribution that yields an expression upper-bounding waste in P_{stat} . (We believe, and the periodic identical property suggests, that this bound is very tight.)

To construct the Markov Chain, we partition time into intervals of length τ ; we regard each such interval as a discrete timestep. Assuming that all honest miners mine at rate rate , let $\lambda = d\tau(\text{rate})$. Thus an honest miner generates an expected $\text{Pois}[\lambda]$ blocks in a given timestep i , which we may represent as a random variable Y_i . Without loss of generality, we may set $\tau = 1/(d \times \text{rate})$ and thus $\lambda = 1$ and $E[\text{Pois}[\lambda]] = 1$.

We represent the state of an honest miner at timestep n by a random variable $X_n = \sum_{i=1}^n (Y_i - E[Y_i]) = (\sum_{i=1}^n Y_i) - n$. Thus $X_n \in \mathbb{Z}$ is simply difference between the miner's actually mined blocks and the expected number.

Our Markov chain consists of a set of states $C = \mathbb{Z}$ representing possible values of X_n (we use the notation C here, as states represent $|C_m|$ for an honest miner m). Figure 3 gives a simple example of such a chain (truncated to only five states).

Our statistical testing regime may be viewed as rejecting blocks when a transition is made to a state whose value is above a certain threshold thresh . We denote the set of such states $C_{\text{rej}} = \{j \mid j \geq \text{thresh}\} \in C$ and depict corresponding nodes visually in our example in Figure 3 as red. P_{stat} sets thresh according to the statistical-testing regime we describe above and a desired false-rejection (Type-I) parameter α . Specifically,

$$C_{\text{rej}}[\alpha] = \{j \in \mathbb{Z} \mid j \geq F^{-1}(1 - \alpha, \tau \times \text{rate})\}. \quad (1)$$




Figure 3: Markov chain with states C representing P_{stat} . Red nodes show the rejection set $C_{\text{rej}} = \mathbb{Z}^+$, i.e., $\text{thresh} = 1$. Outgoing edges from 0 are omitted for clarity.

The transition probabilities in our Markov chain are:

$$P[i \rightarrow j | i \in C \setminus C_{\text{rej}}[\alpha]] = \begin{cases} P(j-i+1) & \text{if } j \geq i-1 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

$$P[i \rightarrow j | i \in C_{\text{rej}}[\alpha]] = \begin{cases} P(j+1) & \text{if } j \leq -1 \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

An example of transitions is given in Figure 3. For instance, from state -1 , the next state can be -2 if the miner doesn't produce any block in this step with probability $P(0)$, or state $-2+i$ if the miner produces $i+1$ blocks in this step, thus with probability $P(i+1)$.

Finally, an upper bound on the false rejection rate can be derived from the stationary probabilities of the Markov chain. Letting $q(s)$ denote the stationary probability of state s ,

$$\text{Waste}(P_{\text{stat}}^\alpha) = \sum_{s \in C_{\text{rej}}[\alpha]} sq(s). \quad (4)$$


We compare our analytic bounds with simulation results in below.

Adversarial Advantage We denote by Σ_{stat} the strategy of an adversary that publishes blocks as soon as they will be accepted by P_{stat} . In Appendix A.4, we show the following:

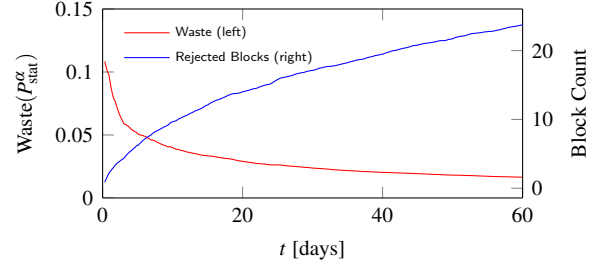
Theorem 1. *In a (non-degenerate) security game S where \mathcal{A} uses strategy Σ_{stat} ,*

$$\text{Adv}_{\mathcal{A}}^{S(P_{\text{stat}}^\alpha)} = \frac{1}{1 - \text{Waste}(P_{\text{stat}}^\alpha)}.$$

Simulation We simulate P_{stat} to explore its efficacy in both the average case and the worst case. Figure 4 shows the result of 1000 runs of a 60-day mining period simulation under P_{stat} . We set $\alpha = 0.4$. We present statistics with respect to the fastest (honest) CPU in the system, which for simplicity we assume mines one block



(a) Left y-axis: adversarial advantage of P_{stat} . Right y-axis: the number of blocks mined by a compromised CPU versus an honest CPU.



(b) Left y-axis: the waste of P_{stat} . Right y-axis: the number of rejected blocks.


Figure 4: 60-day simulation of P_{stat} . The fastest honest CPU mines one block per hour. The Markov chain analysis yields a long-term advantage upper bound of 1.006 and waste of 0.006.

per hour in expectation and refer to simply as “the honest miner.” The adversary uses attack strategy Σ_{stat} .

In Figure 4a, the solid blue line shows the average aggregate number of blocks mined by the adversary, and the dashed one those of the honest miner. The attacker's advantage is, of course, the ratio of the two values. Initially, the adversary achieves a relatively high advantage ($\approx 127\%$), but this drops below 110% within 55 blocks, and continues to drop. Our asymptotic analytic bound on waste (given below) implies an advantage of 100.6%.

Figure 4b shows the average waste of P_{stat} and absolute number of rejected blocks. The waste quickly drops below 10%. As blocks accumulate, the statistical power of P_{stat} grows, and the waste drops further. Analytically, we obtain $\text{Waste}(P_{\text{stat}}^\alpha) = 0.006$, or 0.6% from Eqn. 4.

Setting α Setting the parameter α imposes a trade-off on system implementers. As noted, α corresponds to the Type-I error for a single test in P_{stat} . As P_{stat} performs continuous testing, however, a more meaningful security measure is $\text{Waste}(P_{\text{stat}}^\alpha)$, the rate of falsely rejected blocks. Similarly there is no notion of Type-II error—particularly, as our setting is adversarial. $\text{Adv}_{\mathcal{A}}^{S(P_{\text{stat}}^\alpha)}$ captures the corresponding notion in REM. As shown in Figure 5, raising α results in a lower $\text{Adv}_{\mathcal{A}}^{S(P_{\text{stat}}^\alpha)}$, but higher



(a) The adversarial advantage of P_{stat} under different α (b) The waste of P_{stat} under different α

Figure 5: 60-day simulation of P_{stat} under various α . The fastest honest CPU mines an expected one block per hour.

Waste(P_{stat}^α), and vice versa.

5 Implementation Details

We have implemented a full REM prototype using SGX (§5.1), and as an example application swapped REM into the consensus layer of Bitcoin-core [23]. We explain how we implemented secure instruction counting (§5.2), and our hierarchical attestation framework (§5.3) that allows for arbitrary tasks to be used for work. We explain how to reduce the overhead of attestation due to SGX-specific requirements (§5.4). Finally (§5.5) we present two examples of PoUW and evaluate the overhead of REM.

5.1 Architecture

Figure 1 shows the architecture of REM. As discussed in §3.2, the core of REM is a miner program that does useful work and produces PoUWs. Each CPU instruction executed in the PoUW is analogous to one hash function computation in PoW schemes. That is, each instruction has some probability of successfully mining a block, and if the enclave determines this is the case, it produces a proof — the PoUW.

Pseudocode of the miner’s iterative algorithm is given in Algorithm 1. In a given iteration, it first takes a block template from the agent and calculates the previous block’s hash and difficulty. Then it reads the task to perform as useful work. Note that the enclave code has no network stack, therefore it receives its inputs from the miner untrusted code and returns its outputs to the miner untrusted code. The miner calls the TEE (SGX enclave) with the useful task and parameters for mining, and stores the result of the useful task. It also checks whether the enclave returned a successful PoUW; if so, it combines the agent-furnished template and PoUW into a legal block and sends it to the agent for publication.

Algorithm 1: Miner Loop. The green highlighted line is executed in a TEE (e.g., an SGX enclave).

```

1 while True do
2   template ← read from blockchain agent
3   hash, difficulty ← process(template)
4   task ← get from useful work client
5   outcome, PoUW ← TEE(task, hash, difficulty)
6   send outcome to useful work client
7   if PoUW ≠ ⊥ then
8     block ← formBlock(template, PoUW)
9     send block to blockchain agent

```

In REM, the miner untrusted layer is implemented as a Python script using RPC to access the agent.

To securely decide whether an instruction was a “winning” one, the PoUW enclave does the equivalent of generating a random number and checking whether it is smaller than value `target` that represents the desired system-wide block rate, i.e., difficulty. For this purpose, it uses SGX’s random number generator (SRNG). However, calling the SRNG and checking for a win after every single instruction would impose prohibitive overhead. Instead, we batch instructions by dividing useful work into subtasks of short duration compared to the inter-block interval (e.g. 10 second tasks for 10 minute average block intervals). We let each such subtask run to completion, and count its instructions. The PoUW enclave then calls the SRNG to determine whether at least one of the instructions has won, i.e., it checks for a result less than `target`, weighted by the total number of executed instructions. If so, the enclave produces an attestation that includes the input block hash and difficulty.

Why Count Instructions While instructions are reasonable estimates of the CPU effort, CPU cycles would have been a more accurate metric. However, although cycles are counted, and the counts can be accessed through the CPU’s performance counters, they are vulnerable to manipulation. The operating system may set their values arbitrarily, allowing a rational operator, who controls her own OS, to improve her chances of finding a block by faking a high cycle count. Moreover, counters are incremented even if an enclave is swapped out, allowing an OS scheduler to run multiple SGX instances and having them double-count cycles. Therefore, while instruction counting is not perfect, we find it is the best method for securely evaluating effort with the existing tools available in SGX.




Figure 6: REM toolchain to transfer a useful work to an PoUW-ready program. Everything in the diagram has been implemented besides existing tools such as ld and SGX signing tool.

Algorithm 2: PoUW Runtime

```

1 Function TEE(task, hash, diff)
2   outcome, n := task.run()
3   win := 0
4   PoUW := ⊥
   /* simulating n Bernoulli tests */
5    $l \leftarrow \mathcal{U}[0, 1]$  /* query SGX RNG */
6   if  $l \geq 1 - (1 - \text{diff})^n$  then
7     PoUW =  $\Sigma_{\text{intel}}[\text{hash} \mid \text{diff} \mid 1]$ 
8   return outcome, PoUW

```

5.2 Secure Instruction Counting

As we want to allow arbitrary useful work programs, it is critical to ensure that instructions are counted correctly even in the presence of malicious useful work programs. To this end, we adopt a hybrid method combining static and dynamic program analysis. We employ a customized toolchain that can instrument any SGX-compliant code with dynamic runtime checks implementing secure instruction counting.

Figure 6 shows the workflow of the PoUW toolchain. First, the useful work code (usefulwork.cpp), C / C++ source code, is assembled while reserving a register as the instruction counter. Next, the assembly code is rewritten by the toolchain such that the counter is incremented at the beginning of each basic block (a linear code sequence with no branches) by the number of instructions in that basic block. In particular, we use the LEA instruction to perform incrementing for two reasons. First, it completes in a single cycle, and second, it doesn’t change flags and therefore does not affect conditional jumps. The count is performed at the beginning of a block rather than its end to prevent a cheater from jumping to the middle of a block and gaining an excessive count.

Another challenge is how to ensure the result of instruction counting is used properly—we cannot rely on

the useful work programs themselves. The solution is to wrap the useful work with a predefined, trusted PoUW runtime, and make sure to the enclave can only be entered through the PoUW runtime. The logic of PoUW runtime is summarized in Algorithm 2, and it is denoted as PoUWruntime.so in Figure 6. The PoUW runtime serves as an “in-enclave” loader that launches the useful work program with proper input and collects the result of instruction counting. It takes the block hash and difficulty and starts mining by running the mining program. Once the mining program returns, the PoUW runtime extracts the instruction counter from the reserved register. Then it draws a random value from SRNG and determines whether a new block should be generated, based on the instruction counter and the current difficulty. If a block should be generated, the PoUW runtime produces an attestation recording the template hash that it is called with and the difficulty.

The last step of the toolchain is to compile the resultant assembly and link it (using linker GNU ld) with the PoUW runtime (PoUWruntime.so), to produce the PoUW enclave. Figure 21 in the Appendix shows a snippet of instrumented assembly code. This PoUW enclave is finally signed by an Intel SGX signing tool, creating an application PoUWEnclave.so that is validated for loading into an enclave.

The security of instruction counting relies on the assumption that once instrumented, the code cannot alter the its behavior. To realize this assumption in SGX, we need to require two invariants. First, code pages must be non-writable; second, the useful work program must be single threaded.

Enforcing Non-Writable Code Pages Writable code pages allow a program to rewrite itself at runtime. Although necessary in some cases (e.g. JIT), writable code opens up potential security vulnerabilities. In particular, writable code pages are not acceptable in REM because they would allow a malicious useful work program

```

.section data
ENCLAVE_MTX:
    .long 0

.section text
...
enclave_entry:
    xor %rax, %rax
    xchgl ENCLAVE_MTX(%rip), %rax
    cmp %rax, 0
    jnz enclave_entry

```

Figure 7: Code snippet: a spinlock to allow only the first thread to enter `enclave_entry`

to easily bypass the instrumentation. A general memory protection policy would be to require code pages to have $W \oplus X$ permission, namely to be either writable or executable, but not both. However, $W \oplus X$ permissions are not enforced by the hardware. Intel has in fact acknowledged this issue [7] and recommended that enclave code contain no relocation to enable the $W \oplus X$ feature.

REM thus explicitly requires code pages in the enclave code (`usefulwork.so`) to have $W \oplus X$ permission. This is straightforward to verify, as with the current implementation of the SGX loader, code page permissions are taken directly from the ELF program headers [6].

Enforcing Single Threading Another limitation of SGX is that the memory layout is largely predefined and known to the untrusted application. For example, the State Save Area (SSA) frames are a portion of stack memory that stores the execution context when handling interrupts in SGX. This also implies that the SSA pages have to be writable. The address of SSA frames for an enclave is determined at the time of initialization, as the Thread Control Structure (TCS) is loaded by the untrusted application through an EADD instruction. In other words, the address of SSA is always known to the untrusted application. This could lead to attacks on the instruction counting if a malicious program has multiple threads that interact via manipulation of the execution context in SSA. For example, as we will detail later, REM stores the counter in one of the registers. When one thread is swapped out, the register value stored in an SSA is subject to manipulation by another thread.

While more complicated techniques such as Address Space Layout Randomization (ASLR) for SGX could provide a general answer to this problem, for our purposes it suffices to enforce the condition that an enclave can be launched by at most one thread. As an SGX enclave has only one entry point, we can instrument the code with a spinlock to allow only the first thread to pass, as shown in Figure 7.

5.3 Hierarchical Attestation

A blockchain participant that verifies a block has to check whether the useful work program that produced the block’s PoUW followed the protocol and correctly counted its instructions. SGX attestations require such a verifier to obtain a fingerprint of the attesting enclave. As we allow arbitrary work, a naïve implementation would store all programs on the blockchain. Then a verifier that considers a certain block would read the program from the blockchain, verify it correctly counts instructions, calculate its fingerprint, and check the attestation. Beyond the computational effort, just placing all programs on the blockchain for verification would incur leads prohibitive overhead and enable DoS attacks via spamming the chain with overly large programs. The alternative of having an entity that verifies program compliance is also unacceptable, as it puts absolute blockchain control in the hands of this entity: it can authorize programs that deterministically win every execution.

To resolve this predicament, we form PoUW attestations with what we call *two-layer hierarchical attestations*. We hard-code only a single program’s fingerprint into the blockchain, a static-analysis tool called *compliance checker*. The compliance checker runs in a trusted environment and takes a user-supplied program as input. It validates that it conforms with the requirements defined above. First, it confirms the `text` section is non-writable. Then it validates the work program’s compliance by disassembling it and confirming that the dedicated register is reserved for instruction counting and that counts are correct and appear where they should. Next, it verifies that the PoUW runtime is correctly linked and identical to the expected PoUW runtime code. Finally, it verifies the only entry point is the PoUW runtime and that this is protected by a spinlock as shown in Figure 7. Finally, it calculates the program’s fingerprint and outputs an attestation including this fingerprint.

Every PoUW then includes two parts: The useful work program attestation on the mining success, and an attestation from the compliance checker of the program’s compliance (Figure 8). Note that the compliance attestation and the program’s attestation must be signed by the same CPU. Otherwise an attacker that compromises a single CPU could create fake compliance attestations for invalid tasks. Such an attacker could then create blocks at will from different uncompromised CPUs, circumventing the detection policy of Section 4.

In summary, the compliance enclave is verified through the hard-coded measurement in the blockchain agent. Its output is a measurement that should be identical to the measurement of the PoUW enclave `PoUWEnclave.so`. PoUW Enclave’s output should match the block template (namely the hash of the block




Figure 8: Block structure with a proof comprising the quotes from the compliance enclave and a work enclave.

prefix, up to the proof) and the prescribed difficulty.

Generalized Hierarchical Attestation The hierarchical attestation approach can be useful for other scenarios where participants need to obtain attestations to code they do not know in advance. As a general approach, one hard-codes the fingerprint of a *root compliance checker* that verifies its children’s compliance. Each of them, in turn, checks the compliance of its children, and so on, forming a tree. The leaves of the tree are the programs that produce the actual output to be verified. A hierarchical attestation therefore comprises a leaf attestation and a path to the root compliance checker. Each node attests the compliance of its child.

5.4 IAS access overhead

Verifying blocks doesn’t require trusted hardware. However, due to a design choice by Intel, miners must contact the IAS to verify attestations. Currently there is no way to verify attestations locally. This requirement, however, does not change the basic security assumptions. Moreover, a simple modification to the IAS protocol, which is being tested by Intel [5], could get rid of the reliance on IAS completely on verifiers’ side.

Recall that the IAS is a public web service that receives SGX attestations and responds with verification results. Requests are submitted to the IAS over HTTPS; a response is a signed “report” indicating the validation status of the queried platform [50]. In the current version of IAS, a report is not cryptographically linked with its corresponding request, which makes the report only trustworthy for the client initiating the HTTPS session. Therefore an IAS access is required for every block verification by every blockchain participant.

However, the following modification can eliminate this overhead: simply echoing the request in the body of the report. Since the report is signed by Intel using a published public key [50, 51], only one access to IAS would be needed globally for every new block. Other miners




Figure 9: REM Overhead

could use the resulting signed report. Such a change is under testing by Intel for future versions of the IAS [5].

5.5 Experiments

We evaluate the overhead of REM with four examples of useful work benchmarks in REM as mining programs: a protein folding algorithm [1], a Support Vector Machine (SVM) classifier [27], the zlib compression algorithm (iterated) [2], and the SHA3-256 hash algorithm (iterated) [10]. We evaluate each benchmark in three modes:

Native We compile with the standard toolchain.

SGX We port to SGX by removing system calls and replacing system libraries with SGX-compliant ones. Then we compile in SGX-prerelease mode and run with the SGX driver v1.7 [49].

REM After porting to SGX, we instrument the code using our REM toolchain. We then proceed as in the SGX mode.

We use the same optimization level (-O2) in all modes. The experiments are done on a Dell Precision Workstation with an Intel 6700K CPU and 32GB of memory. For more details on our experimental setup, see §D.1.

We compared the running time in three modes and the results are shown in Figure 9. The running time of the native mode is normalized to one as a baseline. For all four useful workloads, we observe an total overhead of 5.8% ~ 14.4% in REM relative to the native mode. Because the code is instrumented at control flow transfers, workloads with more jumps will incur more counting overhead. For example, SHA3-256 is highly iterative compared with the other workloads, so it incurs the most counting overhead.

We note that overhead for running in SGX is not uniform. For computation-bound workloads such as protein folding, zlib, and SHA3, SGX introduces little overhead (< 1%) because the cost of switching to SGX and obtaining attestations is amortized by the longer in-enclave execution time of the workload. In the shorter SVM benchmark, the cost of entering SGX is more significant.

In summary, we observe an overhead of roughly 5 – 15% for converting useful-work benchmarks into REM PoUW enclave.

6 Waste Analysis

To compare PoUW against PoET and alternative schemes, we explore a common game-theoretic model (with details deferred to the appendix). We consider a set of operators / agents that can either work locally on their own useful workloads or utilize their resource for mining. Based on the revenue from useful work and mining, and the capital and operational costs, we compute the equilibrium point of the system. We calculate the waste in this context as the ratio of the total resource cost (in U.S. dollars) spent per unit of useful work on a mining node compared with the cost when mining is not possible and all operators do useful work. We plug in concrete numbers for the parameters based on statistics we collected from public data sources.

Initial study of PoET identified a subtle pitfall involving miner’s ability to mine simultaneously on multiple blockchains, a problem solved by He et al. [42] in a scheme we call *Lazy-PoET*. Our analysis, however, reveals that even *Lazy-PoET* suffers from what we call the *stale-chip problem*. Miners are better off maintaining farms of cheap, outdated CPUs just for mining than using new CPUs for otherwise useful goals.

We consider instead an approach in which operators utilize their CPUs while mining, making newer CPUs more attractive due to the added revenue from the useful work done. We call this scheme *Busy PoET*. We find that it improves on *Lazy Poet*, but remains highly wasteful.

This observation leads to another approach, *Proof of Potential Work (PoPW)*. PoPW is similar to *Busy-PoET*, but reduces mining time according to the speed of the CPU (its potential to do work), and thus rewards use of newer CPUs. Although PoPW would greatly reduce waste, SGX does not allow an enclave to securely retrieve its CPU model, making PoPW theoretical only.

We conclude that PoUW incurs the smallest amount of waste among the options under study. For full details on our model, parameter choices, and analyses of the various mining schemes, we refer the reader to Appendices §B and §C.

7 Related Work

Cryptocurrencies and Consensus. Modern decentralized cryptocurrencies have stimulated strong interest in Proof-of-Work (PoW) systems [17, 33, 52] as well as techniques to reduce their associated waste.³

An approach similar to PoET [47], possibly originating with Dryja [31], is to limit power waste by so-called Proof-of-Idle. Miners buy mining equipment and get

³“Permissioned” systems, as supported in, e.g., Hyperledger [25] and Stellar [61], avoid waste by using traditional consensus protocols at the cost of avoiding decentralization.

paid for proving that their equipment remains idle. Beyond the technical challenges, as in PoET, an operator with a set budget could redirect savings from power to purchase more idle machines, producing capital waste.

Alternative approaches, like PoUW, aim at PoW producing work useful for a secondary goal. Permacoin [63] repurposes mining resources as a distributed storage network, but recycles only a small fraction of mining resources. Primecoin [55] is an active cryptocurrency whose “useful outputs” are Cunningham and Bi-twin chains of prime numbers, which have no known utility. Gridcoin [41, 40], an active cryptocurrency whose miners work for the BOINC [14] grid-computing network, relies on a central entity. FoldingCoin [70] rewards participants for work on a protein folding problem, but as a layer atop, not integrated with, Bitcoin.

Proof-of-Stake [77, 19, 54, 21] is a distinct approach in which miners gain the right to generate blocks by committing cryptocurrency funds. It is used in experimental systems such as Peercoin [56] and NXT [28]. Unlike PoW, however, in PoS, an attacker that gains majority control of mining resources for a bounded time can control the system forever. PoS protocols also require that funds, used as stake, remain frozen (and unusable) for some time. To remove this assumption, Bentov et al. [20] and Duong et al. [32] propose hybrid PoW / PoS systems. These works, and the line of hybrid blockchain systems starting with Bitcoin-NG [36, 57, 68], can all utilize PoUW as a low-waste alternative to PoW.

Another line of work on PoW for cryptocurrencies aims at PoWs that resist mining on dedicated hardware and prevent concentration of mining power, e.g., via memory-intensive hashing as in Script [59] and Ethereum [24]. Although democratization of mining power is not our focus here, PoUW in fact achieves this goal by restricting mining to general-use CPUs.

SGX. Due to the complexity of the x86-64 architecture, several works [29, 75, 79] have exposed security problems in SGX, such as side-channel attacks [79]. Tramer et al. [75] consider the utility of SGX if its confidentiality guarantees are broken. Similar practical concerns motivate REM’s tolerance mechanism of compromised SGX chips.

Ryoan [44] is a framework that allows a server to run code on private client data and return the output to the client. The (trusted) Ryoan service instruments the server operator’s code to prevent leakage of client data. In contrast, in REM, the useful-workload code is instrumented in an *untrusted* environment, and an attestation of its validity is produced within a trusted environment.

Haven [18] runs non-SGX applications by incorporating a library OS into the enclave. REM, in contrast, takes code amenable to SGX compilation and enforces cor-

rect instrumentation. In principle, Haven could allow for non-SGX code to be adapted for PoUW.

Zhang et al. [80] and Juels et al. [53] are the first works we are aware of to pair SGX with cryptocurrencies. Their aim is to augment the functionality of smart contracts, however, and is unrelated to the underlying blockchain layer in which REM operates.

8 Conclusion

We presented REM, which supports permissionless blockchain consensus based on a novel mechanism called Proof of Useful Work (PoUW). PoUW leverages Intel SGX to significantly reduce the waste associated with Proof of Work (PoW), and builds on and remedies shortcomings in Intel’s innovative PoET scheme. PoUW and REM are thus a promising basis for partially-decentralized blockchains, reducing waste given certain trust assumptions in a hardware vendor such as Intel.

Using a rigorous analytic framework, we have shown how REM can achieve resilience against compromised nodes with minimal waste (rejected honest blocks). This framework extends to PoET and potentially other SGX-based mining approaches.

Our implementation of REM introduces powerful new techniques for SGX applications, namely instruction-counting instrumentation and hierarchical attestation, of potential interest beyond REM itself. They allow REM to accommodate essentially any desired workloads, permitting flexible adaptation in a variety of settings.

Our framework for economic analysis offers a general means for assessing the true utility of mining schemes, including PoW and partially-decentralized alternatives. Beyond illustrating the benefits of PoUW and REM, it allowed us to expose risks of approaches such as PoET in the use of stale chips, and propose improved variants, including Proof of Potential Work (PoPW). We found that small changes to the TEE framework would be significant for reduced-waste blockchain mining. In particular, allowing for secure instruction (or cycle) counting would reduce PoUW overhead, and a secure chip-model reading instruction would allow for PoPW implementation.

We reported on a complete implementation of REM, swapped in for the consensus layer in Bitcoin core in a prototype system. Our experiments showed minimal performance impact (5-15%) on example benchmarks. In summary, our results show that REM is practically deployable and promising path to fair and environmentally friendly blockchains in partially-decentralized blockchains.

Acknowledgements

This work is funded in part by NSF grants CNS-1330599, CNS-1514163, and CNS-1564102, ARO grant W911NF-16-1-0145, and IC3 sponsorship from Chain, IBM, and Intel.

References

- [1] A Genetic Algorithm for Predicting Protein Folding in the 2D HP Model. <https://github.com/alican/GeneticAlgorithm>. Accessed: 2016-11-11.
- [2] A Lossless, High Performance Implementation of the Zlib (RFC 1950) and Deflate (RFC 1951) Algorithm. <https://code.google.com/archive/p/miniz/>. Accessed: 2017-2-16.
- [3] Amazon EC2 instance pricing. <https://aws.amazon.com/ec2/instance-types//aws.amazon.com/ec2/pricing/on-demand/>. Accessed: 2016-10-29.
- [4] Amazon EC2 instance types. <https://aws.amazon.com/ec2/instance-types/>. Accessed: 2016-10-29.
- [5] Attestation Service for Intel Software Guard Extensions (Intel SGX): API Documentation, Revision 2.0, Section 4.2.2. <https://software.intel.com/sites/default/files/managed/7e/3b/ias-api-spec.pdf>. Accessed: 2017-2-21.
- [6] Intel(R) Software Guard Extensions for Linux OS. <https://github.com/01org/linux-sgx>. Accessed: 2017-2-16.
- [7] Intel Software Guard Extensions Enclave Writer’s Guide. <https://software.intel.com/sites/default/files/managed/ae/48/Software-Guard-Extensions-Enclave-Writers-Guide.pdf>. Accessed: 2017-2-16.
- [8] Passmark software. <https://www.cpubenchmark.net/>. Accessed: 2016-10-29.
- [9] Sawtooth-core source code (validator). https://github.com/hyperledger/sawtooth-core/tree/0-7/validator/sawtooth_validator/consensus/poet1. Accessed: 2017-2-21.
- [10] Single-file C implementation of the SHA-3 implementation with Init/Update/Finalize hashing (NIST FIPS 202). <https://github.com/brainhub/SHA3IUF>. Accessed: 2017-2-16.
- [11] Stale CPU dealers on Alibaba. <https://wefound.en.alibaba.com/>. Accessed: 2016-10-29.
- [12] UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml/datasets.html>.
- [13] ANATI, I., GUERON, S., JOHNSON, S., AND SCARLATA, V. Innovative technology for CPU based attestation and sealing. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy* (2013), vol. 13.
- [14] ANDERSON, D. P. Boinc: A system for public-resource computing and storage. In *Grid Computing, 2004. Proceedings. Fifth IEEE/ACM International Workshop on* (2004), IEEE, pp. 4–10.
- [15] ASPNES, J., JACKSON, C., AND KRISHNAMURTHY, A. Exposing computationally-challenged Byzantine impostors. *Department of Computer Science, Yale University, New Haven, CT, Tech. Rep* (2005).
- [16] AZURE, M. Blockchain as a service. <https://web.archive.org/web/20161027013817/https://azure.microsoft.com/en-us/solutions/blockchain/>, 2016.

- [17] BACK, A. Hashcash – a denial of service counter-measure. <http://www.cypherspace.org/hashcash/hashcash.pdf>, 2002.
- [18] BAUMANN, A., PEINADO, M., AND HUNT, G. Shielding applications from an untrusted cloud with Haven. *ACM Trans. Comput. Syst.* 33, 3 (Aug. 2015), 8:1–8:26.
- [19] BENTOV, I., GABIZON, A., AND MIZRAHI, A. Cryptocurrencies without proof of work. *CoRR abs/1406.5694* (2014).
- [20] BENTOV, I., LEE, C., MIZRAHI, A., AND ROSENFELD, M. Proof of activity: Extending bitcoin’s proof of work via proof of stake. Cryptology ePrint Archive, Report 2014/452, 2014. <http://eprint.iacr.org/2014/452>.
- [21] BENTOV, I., PASS, R., AND SHI, E. Snow white: Provably secure proofs of stake. Cryptology ePrint Archive, Report 2016/919, 2016. <http://eprint.iacr.org/2016/919>.
- [22] BITCOIN COMMUNITY. Bitcoin source. <https://github.com/bitcoin/bitcoin>, retrieved Nov. 2016.
- [23] BITCOIN COMMUNITY. Bitcoin source. <https://github.com/bitcoin/bitcoin>, retrieved Mar. 2015.
- [24] BUTERIN, V. A next generation smart contract & decentralized application platform. <https://www.ethereum.org/pdfs/EthereumWhitePaper.pdf/>, retrieved Feb. 2015, 2013.
- [25] CACHIN, C. Architecture of the Hyperledger blockchain fabric. In *Workshop on Distributed Cryptocurrencies and Consensus Ledgers* (2016).
- [26] CARLSTEN, M., KALODNER, H., WEINBERG, S. M., AND NARAYANAN, A. On the instability of bitcoin without the block reward. In *ACM CCS* (2016).
- [27] CHANG, C.-C., AND LIN, C.-J. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* 2 (2011), 27:1–27:27. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [28] COMMUNITY, T. N. Nxt whitepaper, revision 4. https://web.archive.org/web/20160207083400/https://www.dropbox.com/s/cbuwrorf672c0yy/NxtWhitepaper_v122_rev4.pdf, 2014.
- [29] COSTAN, V., AND DEVADAS, S. Intel SGX Explained. *Cryptology ePrint Archive* (2016).
- [30] DEETMAN, S. Bitcoin could consume as much electricity as Denmark by 2020. <https://web.archive.org/web/20160828092858/http://motherboard.vice.com/read/bitcoin-could-consume-as-much-electricity-as-denmark-by-2020>, March 2016.
- [31] DRYJA, T. Optimal mining strategies. SF Bitcoin-Devs presentation. <https://www.youtube.com/watch?v=QN2TPeQ9mna>, 2014.
- [32] DUONG, T., FAN, L., VEALE, T., AND ZHOU, H.-S. Securing bitcoin-like backbone protocols against a malicious majority of computing power. Cryptology ePrint Archive, Report 2016/716, 2016. <http://eprint.iacr.org/2016/716>.
- [33] DWORCK, C., AND NAOR, M. Pricing via processing or combatting junk mail. In *Annual International Cryptology Conference* (1992), Springer, pp. 139–147.
- [34] DWYER, J. P., AND HINES, P. Beyond the byzz: Exploring distributed ledger technology use cases in capital markets and corporate banking. Tech. rep., Celent and MISYS, 2016.
- [35] EYAL, I. The miner’s dilemma. In *IEEE Symposium on Security and Privacy* (2015), pp. 89–103.
- [36] EYAL, I., GENCER, A. E., SIRER, E. G., AND VAN RENESSE, R. Bitcoin-ng: A scalable blockchain protocol. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)* (2016), pp. 45–59.
- [37] EYAL, I., AND SIRER, E. G. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security* (2014).
- [38] GARAY, J. A., KIAYIAS, A., AND LEONARDOS, N. The Bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques* (2015), pp. 281–310.
- [39] GARWOOD, F. (i) fiducial limits for the poisson distribution. *Biometrika* 28, 3-4 (1936), 437.
- [40] GRIDCOIN. Gridcoin. <https://web.archive.org/web/20161013081149/http://www.gridcoin.us/>, 2016.
- [41] GRIDCOIN. Gridcoin (grc) – first coin utilizing boinc – official thread. <https://web.archive.org/web/20160909032618/https://bitcointalk.org/index.php?topic=324118.0>, 2016.
- [42] HE, W., SONG, D., AND MILUTINOVIC, M. SGX and smart contracts. Initiative for Cryptocurrencies and Contracts Retreat (presentation), 2016.
- [43] HOEKSTRA, M., LAL, R., PAPPACHAN, P., PHEGADE, V., AND DEL CUVILLO, J. Using innovative instructions to create trustworthy software solutions. In *Proceedings of the 2Nd International Workshop on Hardware and Architectural Support for Security and Privacy* (New York, NY, USA, 2013), HASP ’13, ACM, pp. 11:1–11:1.
- [44] HUNT, T., ZHU, Z., XU, Y., PETER, S., AND WITCHEL, E. Ryoan: A distributed sandbox for untrusted computation on secret data. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)* (GA, Nov. 2016), USENIX Association, pp. 533–549.
- [45] INTEL. *Intel Software Guard Extensions Programming Reference*, 2014.
- [46] INTEL. *Intel 64 and IA-32 Architectures Software Developer’s Manual: Volume 3 (3A, 3B, 3C & 3D): System Programming Guide*, 325384-059us ed., June 2016.
- [47] INTEL. Sawtooth lake – introduction. <https://web.archive.org/web/20161025232205/https://intelledger.github.io/introduction.html>, 2016.
- [48] INTEL CORPORATION. Intel® Software Guard Extensions SDK. <https://software.intel.com/en-us/sgx-sdk>, 2015.
- [49] INTEL CORPORATION. Intel SGX for Linux. <https://01.intel-softwareguard-extensions>, 2016.
- [50] INTEL CORPORATION. Intel Software Guard Extensions: Intel Attestation Service API. https://software.intel.com/sites/default/files/managed/3d/c8/IAS_1_0_API_spec_1_1_Final.pdf, 2016.
- [51] INTEL CORPORATION. Public Key for Intel Attestation Service. <https://software.intel.com/en-us/sgx/resource-library>, 2016.
- [52] JAKOBSSON, M., AND JUELS, A. Proofs of work and bread pudding protocols. In *Secure Information Networks*. Springer, 1999, pp. 258–272.
- [53] JUELS, A., KOSBA, A., AND SHI, E. The Ring of Gyges: Investigating the future of criminal smart contracts. In *ACM CCS* (2016), pp. 283–295.
- [54] KIAYIAS, A., KONSTANTINOU, I., RUSSELL, A., DAVID, B., AND OLIYNYKOV, R. A provably secure proof-of-stake blockchain protocol. Cryptology ePrint Archive, Report 2016/889, 2016. <http://eprint.iacr.org/2016/889>.
- [55] KING, S. Primecoin: Cryptocurrency with prime number proof-of-work. <https://web.archive.org/web/20160307052339/http://primecoin.org/static/primecoin-paper.pdf>, 2013.

- [56] KING, S., AND NADAL, S. PPcoin: Peer-to-peer crypto-currency with proof-of-stake. <https://web.archive.org/web/20161025145347/https://peercoin.net/assets/paper/peercoin-paper.pdf>, 2012.
- [57] KOGIAS, E. K., JOVANOVIĆ, P., GAILLY, N., KHOFFI, I., GASSER, L., AND FORD, B. Enhancing bitcoin security and performance with strong consistency via collective signing. In *25th USENIX Security Symposium (USENIX Security 16)* (2016), pp. 279–296.
- [58] LEWENBERG, Y., SOMPOLINSKY, Y., AND ZOHAR, A. Inclusive block chain protocols. In *Financial Cryptography* (Puerto Rico, 2015).
- [59] LITECOIN PROJECT. Litecoin, open source P2P digital currency. <https://litecoin.org>, retrieved Nov. 2014.
- [60] LOMBROZO, E., LAU, J., AND WUILLE, P. BIP141: Segregated witness (consensus layer). <https://web.archive.org/web/20160521104121/https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki>, 2015.
- [61] MAZIERES, D. The Stellar consensus protocol: A federated model for Internet-level consensus. <https://web.archive.org/web/20161025142145/https://www.stellar.org/papers/stellar-consensus-protocol.pdf>, 2015.
- [62] MCKEEN, F., ALEXANDROVICH, I., BERENZON, A., ROZAS, C. V., SHAFI, H., SHANBHOGUE, V., AND SAVAGAONKAR, U. R. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy* (2013), p. 10.
- [63] MILLER, A., SHI, E., JUELS, A., PARNO, B., AND KATZ, J. Permacoin: Repurposing Bitcoin work for data preservation. In *Proceedings of the IEEE Symposium on Security and Privacy* (San Jose, CA, USA, 2014), IEEE.
- [64] NAKAMOTO, S. Bitcoin: A peer-to-peer electronic cash system. <http://www.bitcoin.org/bitcoin.pdf>, 2008.
- [65] NAYAK, K., KUMAR, S., MILLER, A., AND SHI, E. Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. *IACR Cryptology ePrint Archive 2015* (2015), 796.
- [66] PANDE, V. Protein folding: The science. <https://web.archive.org/web/20161016120350/https://folding.stanford.edu/home/the-science/>.
- [67] PASS, R., SEEMAN, L., AND SHELAT, A. Analysis of the blockchain protocol in asynchronous networks. Tech. rep., Cryptology ePrint Archive, Report 2016/454, 2016.
- [68] PASS, R., AND SHI, E. Hybrid consensus: Efficient consensus in the permissionless model. Cryptology ePrint Archive, Report 2016/917, 2016. <http://eprint.iacr.org/2016/917>.
- [69] POPPER, N. Central banks consider bitcoin’s technology, if not bitcoin. *New York Times*, Oct 11 2016.
- [70] ROSS, R., AND SEWELL, J. Foldingcoin white paper. <https://web.archive.org/web/2016102232226/http://foldingcoin.net/the-coin/white-paper/>, 2015.
- [71] SAPIRSHTAIN, A., SOMPOLINSKY, Y., AND ZOHAR, A. Optimal selfish mining strategies in Bitcoin. *CoRR abs/1507.06183* (2015).
- [72] SIMON JOHNSON, VINNIE SCARLATA, CARLOS ROZAS, ERNIE BRICKELL, AND FRANK MCKEEN. Intel Software Guard Extensions: EPID Provisioning and Attestation Services, 2015.
- [73] SOMPOLINSKY, Y., AND ZOHAR, A. Accelerating Bitcoin’s transaction processing, fast money grows on trees, not chains. In *Financial Cryptography* (Puerto Rico, 2015).
- [74] SWIFT, AND ACCENTURE. Swift on distributed ledger technologies. Tech. rep., SWIFT and Accenture, 2016.
- [75] TRAMER, F., ZHANG, F., LIN, H., HUBAUX, J.-P., JUELS, A., AND SHI, E. Sealed-glass proofs: Using transparent enclaves to prove and sell knowledge. Cryptology ePrint Archive, Report 2016/635, 2016. <http://eprint.iacr.org/2016/635>.
- [76] UNGER, R., AND MOULT, J. Genetic algorithms for protein folding simulations. *Journal of Molecular Biology* 231, 1 (1993), 75 – 81.
- [77] USER “QUANTUMMECHANIC”. Proof of stake instead of proof of work. <https://web.archive.org/web/20160320104715/https://bitcointalk.org/index.php?topic=27787.0>.
- [78] WOOD, G. Ethereum: A secure decentralised generalised transaction ledger (EIP-150 revision). <https://web.archive.org/web/20161019105532/http://gavwood.com/Paper.pdf>, 2016.
- [79] XU, Y., CUI, W., AND PEINADO, M. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *Proc. IEEE Symp. Security and Privacy* (May 2015), pp. 640–656.
- [80] ZHANG, F., CECCHETTI, E., CROMAN, K., JUELS, A., AND SHI, E. Town crier: An authenticated data feed for smart contracts. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, 2016), CCS ’16, ACM, pp. 270–282.

A Tolerating Compromised SGX Nodes: Details

A.1 Mining Rate Estimation

We start by discussing how to statistically infer the power of a CPU from its blocks in the blockchain. Reading the difficulty of each block in the main chain and the rate of blocks from a specific CPU, we can estimate a lower bound of that CPU’s power – it follows directly from the rate of its blocks. It is an lower bound since the CPU might not be working continuously, and the estimate’s accuracy increases with the number of available blocks.

Recall C_{m_i} is the blocks mined by miner m_i so far. C_{m_i} may contain multiple blocks, perhaps with varying difficulties. Without loss of generality, we write the difficulty as a function of time, $d(t)$. The difficulty is the probability for a single instruction to yield a win. Denote the power of the miner, i.e., its mining rate, by $rate_i$. Therefore in a given time interval of length T , the number of blocks mined by a specific CPU obeys Poisson distribution (since CPU rates are high and the win probability is small, it’s appropriate to approximate a Binomial distribution by a Poisson distribution,) and with rate $rate_i T d(t)$. Further, under independence assumption, the mining process of a specific CPU is specified by a Poisson process with rate $\lambda_i(t) = rate_i d(t)$, the product of the probability and the miner’s rate $rate_i$.

There are many methods to estimate the mean of a Poisson distribution. For example, by [39] an interval estimation of λ_i can be derived from $n = |C_{m_i}|$, with $1 - \alpha$

confidence, as

$$\lambda_i \in \left(\frac{1}{2T} \chi_{\alpha/2, 2n}^2, \frac{1}{2T} \chi_{1-\alpha/2, 2(n+1)}^2 \right). \quad (5)$$

Therefore, let $D = \int_{t_0}^{t_f} d(t) dt$, the $1 - \alpha$ confidence interval of rate_i is

$$\left(\frac{1}{2D} \chi_{\alpha/2, 2n}^2, \frac{1}{2D} \chi_{1-\alpha/2, 2(n+1)}^2 \right). \quad (6)$$

Knowing rates for all miners, the rate of the strongest CPU ($\text{rate}_{\text{best}}$) can be estimated. The challenge here is to limit the influence of adversarial nodes. To this end, instead of finding the strongest CPU directly, we approximate $\text{rate}_{\text{best}}$ based on rate_ρ (e.g. $f_{90\%}$), namely the ρ -percentile fastest miner.

Claim 1. *Suppose there are n nodes and k out of them are adversarial. Let $\mathbf{r}_h = \{\text{rate}_1, \dots, \text{rate}_{n-k}\}$ be the rates of $n - k$ honest nodes, in monotonically non-decreasing order, and let $\hat{\mathbf{r}} = \{\hat{\text{rate}}_{1/n}, \dots, \hat{\text{rate}}_{100\%}\}$ be the rates of all nodes, in monotonically non-decreasing order, then $\text{rate}_{k+1} \leq \hat{\text{rate}}_{\frac{k+1}{n}} \leq \text{rate}_1$.*

By Claim 1, we can obtain $\text{rate}_{\text{best}}$ that is at least as high as rate_ρ .

Bootstrapping. During the launch of a cryptocurrency, it could be challenging to estimate the mining power of the population accurately, potentially leading to poisoning attacks by an adversary. At this early stage, it makes sense to hardwire a system estimate of the maximum mining power of honest miners into the system and set conditions (e.g., a particular mining rate or target date) to estimate $\text{rate}_{\text{best}}$ as we propose above. If the cryptocurrency launches with a large number of miners, an even simpler approach is possible before switching to $\text{rate}_{\text{best}}$ estimation: We can cap the total number of blocks that any one node can mine, a policy we illustrate below. (See P_{simple} .)

A.2 Security game definition

We model REM as an interaction among three entities: a blockchain consensus algorithm, an adversary, and a set of honest miners. Their behavior together defines a *security game*, which we define formally below. We characterize the three entities respectively as (ideal) programs $\text{prog}_{\text{chain}}$, $\text{prog}_{\mathcal{A}}$, and prog_m , which we now define.

Blockchain consensus algorithm ($\text{prog}_{\text{chain}}$). A consensus algorithm determines which valid blocks are added to a blockchain C . We assume that underlying




Figure 10: The program for a blockchain. We omit details here on how difficulty d is set, i.e., how d_0 and adjust are chosen.

consensus and fork resolution are instantaneous; loosening this assumption does not materially affect our analyses. We also assume that block timestamping is accurate. Timestamps can technically be forged at block generation, but in practice miners reject blocks with large skews [23], limiting the impact of timestamp forgery.

Informally, $\text{prog}_{\text{chain}}$ maintains and broadcasts and authoritative blockchain C . In addition to verifying that block contents are correct, $\text{prog}_{\text{chain}}$ appends to C only blocks that are valid under a policy P . We model the blockchain consensus algorithm as the (ideal) stateful program specified in Figure 10.

Adversary \mathcal{A} ($\text{prog}_{\mathcal{A}}$). In our model, an adversary \mathcal{A} executes a *strategy* $\Sigma_{\mathcal{A}}$ that coordinates the k miners $M_{\mathcal{A}}$ under her control to generate blocks. Specifically:

Definition 4. (*Adversarial Strategy*). *An adversarial strategy is a probabilistic algorithm $\Sigma_{\mathcal{A}}$ that takes in a set of identities, the current blockchain and the policy, and outputs a time-stamp and identity for block submission. Specifically, $(M_{\mathcal{A}}, C, t, P) \rightarrow (\hat{t}, \hat{m}) \in \mathbb{R}^+ \times M_{\mathcal{A}}$.*

In principle, $\Sigma_{\mathcal{A}}$ can have dependencies among individual node behaviors. In our setting, this would not benefit \mathcal{A} , however. As we don’t know $M_{\mathcal{A}}$ *a priori*, though, the only policies we consider operate on individual miner block-generation history.

As a wrapper expressing implementation by \mathcal{A} of $\Sigma_{\mathcal{A}}$, we model \mathcal{A} as a program $\text{prog}_{\mathcal{A}}$, specified in Figure 11.

Honest miners (prog_m). Every honest miner $m \in M - M_{\mathcal{A}}$ follows an identical strategy, a probabilistic algorithm denoted Σ_h . In REM, Σ_h may be modeled as a simple algorithm that samples from a probability distribution on block mining times determined by $\text{rate}(m)$ (specifically in our setting, an exponential distribution with rate

```

                                progA[ΣA]

On receive (C, P, d) from progchain
     $\hat{t}, \hat{m} \leftarrow \Sigma_{\mathcal{A}}(M_{\mathcal{A}}, C, P, d)$ 
    if  $\hat{t}$  is not  $\perp$ :
        wait until  $\hat{t}$ 
        send “submit” ( $\hat{t}, \hat{m}, d$ ) to progchain

```

Figure 11: The program for an adversary \mathcal{A} that controls k nodes $M_{\mathcal{A}} = \{m_{\mathcal{A}1}, \dots, m_{\mathcal{A}k}\}$.

```

                                progm[Σh]

On receive (C, P, d) from progchain
     $\hat{t} \leftarrow \Sigma_h(C, d)$ 
    Send “submit” ( $\hat{t}, m, d$ ) to progchain

```

Figure 12: The program for an honest miner. Σ_h is the protocol defined by prog_{chain} (e.g. PoET or PoUW).

$rate(m)$). We express implementation by honest miner m of Σ_h as a program prog_m[Σ_h] (Figure 12).

To understand the security of REM, we consider a *security game* that defines how an adversary \mathcal{A} interacts with honest miners, a blockchain consensus protocol, and a policy given the above three ideal programs. Formally:

Definition 5. (*Security Game*) For a given triple of ideal programs (prog_{chain}[P], prog_A[Σ_A], prog_m[Σ_h]), and policy P, a security game S(P) is a tuple $S(P) = ((M, M_{\mathcal{A}}, rate(\cdot)); (\Sigma_{\mathcal{A}}, \Sigma_h))$.

We define the *execution* of S(P) as an interactive execution of programs (prog_{chain}[P], prog_A[Σ_A], prog_m[Σ_h]) using the parameters of S(P). As P, Σ_A and Σ_h are randomized algorithms, such execution is itself probabilistic. Thus we may view the blockchain resulting from execution of S for interval of time τ as a random variable $C_S(\tau)$.

A *non-degenerate* security game S is one in which there exists at least one honest miner m with $rate(m) > 0$.

A.3 Warmup policy

As a warmup, we give a simple example of a potential block-acceptance policy. This policy just allows one block throughout the life of a CPU, as shown in Figure 13.

Clearly, an adversary can not do better than mining one block. Denote this simple strategy Σ_{simple} . For any non-degenerate security game S, therefore, the advantage $\text{Adv}_{\mathcal{A}}^{S(P_{\text{simple}}^{\alpha})}(\tau) = 1$ as $\tau \rightarrow \infty$. This policy is optimal in

```

Psimple(C, B):
    parse B → (τ, m, d)
    if |Cm| > 0:
        output reject
    else
        output accept

```

Figure 13: A simple policy that allows one block per CPU over its lifetime.

that an adversary can not do better than an honest miner unconditionally. However the asymptotic waste of this policy is 100%.

Another disadvantage of this policy is that it discourages miners from participating. Arguably, miner would stay if the revenue from mining is high enough to cover the cost of replacing a CPU. But though a CPU is still valuable in other context even if it is blacklisted forever in *this* particular system, repurposing it incurs operational cost. Therefore chances are this policy would cause a loss of mining power, especially when the initial miner population is small, rendering the system more vulnerable to attacks.

A.4 Adversarial advantage

A block-acceptance policy depends only on the number of blocks by the adversary since its first one. Therefore an adversary’s best strategy is simply to publish its blocks as soon as they won’t be rejected. Denote this strategy as Σ_{stat} .

Clearly, an adversary will submit $F^{-1}(1 - \alpha, td \cdot \text{rate}_{\text{best}})$ blocks within $[0, t]$. On the other hand, the strongest honest CPU with rate $\text{rate}_{\text{best}}$ mines $td \cdot \text{rate}_{\text{best}}$ blocks in expectation. Recall that according to our Markov chain analysis, P_{stat} incurs false rejections for honest miners with probability $w_h(\alpha)$, which further reduces the payoff for honest miners. For a (non-degenerate) security game S, in which \mathcal{A} uses strategy Σ_{stat} , the advantage is therefore:

$$\text{Adv}_{\mathcal{A}}^{S(P_{\text{stat}}^{\alpha})} = \lim_{t \rightarrow \infty} \frac{F^{-1}(1 - \alpha, td \cdot \text{rate}_{\text{best}})}{(1 - w_h(\alpha)) td \cdot \text{rate}_{\text{best}}} \quad (7)$$

Theorem 1. In a (non-degenerate) security game S where \mathcal{A} uses strategy Σ_{stat} ,

$$\text{Adv}_{\mathcal{A}}^{S(P_{\text{stat}}^{\alpha})} = \frac{1}{1 - \text{Waste}(P_{\text{stat}}^{\alpha})}.$$

Proof. Let $\lambda = td \cdot \text{rate}_{\text{best}}$. It is known that as λ for a Poisson distribution goes to infinity, it converges in the limit to a normal distribution with mean and variance λ .

Therefore,

$$\lim_{\lambda \rightarrow \infty} \frac{F^{-1}(1 - \alpha, \lambda)}{(1 - w_h(\alpha))\lambda} = \lim_{\lambda \rightarrow \infty} \frac{\lambda + \sqrt{\lambda} z_p}{(1 - w_h(\alpha))\lambda} = \frac{1}{1 - w_h(\alpha)}.$$

□

Early in a blockchain’s evolution, the potential advantage of an adversary is relatively high. The confidence interval is wide at this point, allowing the adversary to perform frequent generation without triggering detection. As the adversary publishes more blocks, the confidence interval tightens, forcing the adversary to reduce her mining rate. This is illustrated by our numerical simulation in Section 4.3.

B Resource Consumption Model

In this appendix, we present a general economic / game-theoretic model for modeling resource consumption of consensus schemes. This model guides us toward an understanding of optimal mining strategies for various consensus schemes and thus a basis for comparison among them. We detail the model in Section B.1. We estimate real-world parameter choices for this model in Section B.2 and in Section B.3 present *cost per unit of useful work*, our key metric of waste / resource consumption. We use this model in Appendix C to compare various SGX-based mining schemes show that REM results in less resource waste than alternatives such as PoET.

B.1 Model

In our model, we consider a set of N^{op} operators that choose to commit their CPUs either to mining or unrelated useful work. This reflects the fact that there are certain barriers to enter the mining industry and therefore not everyone will participate. For completeness we also discussed the implication of removing the limit of numbers of operators in §C.3.

Each operator has an annual budget of budget for purchasing hardware and for paying operating expenses. In our context, both expenses are a function of the CPUs the operator chooses to use. Denote by *age* the age in years of the CPUs maintained by the operator, which we assume for simplicity to be uniform. If an operator chooses to maintain new CPUs, she enjoys better performance and efficiency (computation per power unit), but the cost is higher. The latest CPU has an age of 0, and we arbitrarily set the oldest CPU available to 10, denoted $\text{age}_{\text{max}} = 10$. Choosing a higher age_{max} value or removing this limit strengthens our results. In some situations, the operator may choose to have the CPUs but not utilize them. Denote by $u \in [0, 1]$ the utilization level of a CPU, where 0 means idle and 1, fully utilized.

The annual capital cost of maintaining a CPU of age *age* is denoted by the function $C(\text{age})$. The function decreases with CPU age, as older CPUs are significantly cheaper from recycling marketplace. The annual energy cost of a single CPU is denoted by the function $E(u)$, which increases with u . Denote by $\eta(\text{age}) \in [0, 1]$ the performance slowdown of a CPU, normalized to that of the latest model in the same family. η increases with CPU age, as newer generations (of similar MSRP) tend to be more powerful than their antecedents. Beyond the CPU itself, there is an overhead for the platform on which it runs. Denote by O_{std} the annual cost overhead for running a CPU, including server, racking, cooling etc. In some schemes (e.g., PoW), an operator can reduce costs by placing the CPU in a farm— a dedicated platform stripped down to essential resources and thus usable only for mining. Denote by O_{farm} (naturally $O_{\text{farm}} < O_{\text{std}}$) the annual cost overhead for running a CPU in a dedicated mining farm.

Due to performance improvements, the annual income for useful work from a CPU is a function $R_w(u, \text{age})$ of both its age and utilization (or simply R_w when the parameters are clear from the context). The function increases with both parameters. We assume operators have unbounded useful work, enough to populate any and all hardware they can afford.

We let R_{annual} denote the annual total mining revenue of the system, which we assume is independent of any other variable and vary our analyses. This total revenue is divided among the participants in a manner that depends on the specific scheme used, and the details are given below.

Operators strive to maximize their net revenue by choosing from a space of three strategies that dictate hardware use. These are:

1. *No mining*: The operator chooses not to mine, using her CPUs solely for useful work. This is profitable as long as the income R_w offsets costs.
2. *Standard mining*: The operator uses standard servers and mines on them (if the scheme allows).
3. *Farming*: The operator uses farm machines for mining, reducing the per-CPU overhead, but losing the ability to perform useful work.

We adopt a population-based representation of the strategy choices made by operators. Specifically, for operators collectively implementing a hybrid strategy, we express their choice in terms of the aggregate fraction of resources devoted to each strategy by the full population.

To compare the resource waste associated with different consensus systems, we examine the optimal operation point for each. We model operators’ utility, and




Figure 14: CPUs’ price and speed as function of age.

identify the best strategy for a rational operator. We develop expressions for the operators’ revenues, and instantiate them using the values obtained as explained in Section B.2. We then calculate the equilibrium point of each scheme and the optimal operation point for miners in each scheme. We measure the waste of each scheme in terms of the cost for each unit of useful work, referred to as the *useful price*.

B.2 Parameter Values

To estimate the price, performance, and energy efficiency of a CPU as a function of its age, we investigate historical CPU data for a family of 7 Intel CPUs spanning ten years. From each generation of Intel CPUs since 2006, we picked the fastest desktop chip. For each CPU, we estimate its price as of 2016 according to several marketplaces worldwide [11] and its performance according to standard benchmarks [8]. We present CPU specifications as published by Intel. Long discontinued models can still be purchased from many suppliers, generally at very low per-unit costs for high-volume orders. All amounts are given in 2016 USD, denoted by \$.

Table 1: Sample CPUs and specifications. Benchmark scores are normalized by the score of i7-6700K.

| Generation | Sample CPU | Launch Date | Q3’16 Price | speed (normed) | Power at $u = 1$ |
|--------------|------------------|-------------|-------------|----------------|------------------|
| Cedarmill | Pentium 4 661 | Q1’06 | \$5 | -4 | 85W |
| Wolfdale | Core 2 Duo E8500 | Q3’08 | \$10 | 0.21 | 65W |
| Lynnfield | Core i7-880 | Q2’10 | \$35 | 0.52 | 95W |
| Sandy Bridge | Core i7-2700K | Q4’11 | \$50 | 0.80 | 95W |
| Ivy Bridge | Core i7-3770K | Q2’12 | \$75 | 0.87 | 77W |
| Haswell | Core i7-4770K | Q2’13 | \$150 | 0.89 | 84W |
| Skylake | Core i7-6700K | Q3’15 | \$290 | 1.00 | 91W |

Price and performance Table 1 summarizes CPU specs. From the data in Table 1 we can see clearly that the CPU price C drops exponentially as it ages. CPU performance manifests a segmented, exponential trend: It doubled each year between 2006 and 2011 (following Moore’s Law), but then improvements slowed to around

10% per year. Our model uses a logit function to approximate this trend, as is shown in Figure 14. As a result of the performance slowdown of older CPUs, R_w decreases with age in a similar way.

Useful work revenue We used the cost model of Amazon Web Services to approximate the useful-work revenue R_w . In particular, we assume a miner’s mining computer has performance similar to an m4.large EC2 instance, which has 2 CPU cores and 8 GiB of memory [4]. As of the time of writing, an m4.large instance costs 0.12 per hour [3].

B.3 Cost per unit of useful work

To compare the resource costs of the different schemes in Section C, we will compute for each its *cost per unit of useful work*.

As a baseline, we first compute the cost per unit of useful work on an optimal hardware setup devoted exclusively to useful work. The revenue of a single CPU for useful work is the income over a year from useful work (R_w) minus the expenses for hardware purchase (C), overhead (O_{std}), and power (E):

$$R_w(u, \text{age}_w) - (C(\text{age}_w) + O_{std} + E(u, \text{age}_w)). \quad (8)$$

The number of CPUs an operator can purchase is

$$N^{\text{cpu}}(\text{age}_w) = \frac{\text{budget}}{C(\text{age}_w) + O_{std} + E(u, \text{age}_w)}. \quad (9)$$

Therefore the total profit from useful work is

$$P_{\text{useful}}(u, \text{age}_w) = N^{\text{cpu}}(\text{age}_w) \times \left(R_w(u, \text{age}_w) - (C(\text{age}_w) + O_{std} + E(u, \text{age}_w)). \right) \quad (10)$$

For simplicity we assume power usage grows linearly with utilization, and so utilization of 1 is optimal to reduce the waste of the independent elements C and O_{std} . Plugging in values from Section B.2, we can optimize equation 10, resulting in an optimum CPU age of $\hat{\text{age}}_w = 3.48$ years and a corresponding *baseline cost per unit of work of \$391.25*.

For a population of miners, the overall cost per unit of useful work is equal to the total investment in the mining ecosystem over a year, divided by the number of units of useful work. Consequently, it is a function of the fraction of mining CPUs doing useful work. If 100% of CPU time is devoted to useful work, the cost per unit of useful work is \$391.25; 0% corresponds to infinite cost per unit of useful work. Smaller percentages of CPU investment in useful work correspond to higher costs.

Figure 15: Summary of revenue analysis result with parameters according to §B.2. Notation: U is useful work, S is standard mining, F is farming. f_{2nd} is the ratio of participants choosing the second option, which is the more wasteful one. Waste is the cost for one unit of useful work normalized to the baseline (the cost in a system with no mining): \$391.25 ($\text{age}_w = 3.48$).

| Schemes | Choices | age_S | age_F | f_{2nd} | Waste |
|-------------|---------|----------------|---------------------|-----------|-------|
| Useful work | U | - | - | - | 1.0 |
| PoW | U, F | - | 4.68 | 76% | 4.2 |
| Lazy-PoET | U, F | - | age_{\max} | 76% | 4.2 |
| Busy-PoET | S, F | 4.51 | age_{\max} | 42% | 2.5 |
| PoPW | S, F | 4.39 | 5.61 | 26% | 1.4 |
| PoUW | U, S | 4.39 | - | 100% | 1.1 |

C Comparative Analysis of Mining Schemes

Using our model from Appendix §B, we now present an analysis showing that PoUW incentivizes minimally wasteful mining behavior and / or achieves secure consensus more effectively than alternatives. An important part of this section is our presentation of a spectrum of five consensus schemes, including PoW, PoUW, and three other SGX-based schemes (two of them newly presented in this work). These schemes illustrate a range of technical approaches against which we compare PoUW as a means of validating our design choices.

We first present and analyze our four consensus schemes other than PoUW—what we call *strawman schemes*—in Section C.1. We then offer detailed revenue analysis of PoUW in Section C.2. Figure 15 compares the results for the different schemes at the parameter values calculated in Section B.2. We consider a conservative total cryptocurrency revenue of \$20m. At higher values the significance of PoUW would only increase. Figure 16 compares the schemes at varying total cryptocurrency revenue values.

C.1 Strawman Schemes

Here we present our spectrum of strawman consensus schemes that serve as points of comparison with and explanations of design choices in PoUW. We start with the popular Proof of Work (PoW), and then present PoET and a proposed minor variant on PoET, as well as two strawman solutions of our own that demonstrate the difficulty of achieving waste-limited SGX-based consensus. We analyze PoUW in Section C.2. The reader may choose to skip directly to Section C.2 and then flip back here to understand the alternatives and rationale for our design choices.




Figure 16: Revenue analysis: The x axis is the annual cryptocurrency revenue (R_{annual}) and the left y -axis is the waste factor, i.e. useful price normalized by the useful price without mining. The right y -axis is the ratio of participants choosing the more-wasteful option in their scheme.

C.1.1 Proof of Work

We start with PoW, as used in Bitcoin, Ethereum, and other common cryptocurrencies. For the purpose of this analysis, we consider only CPU mining. See below for a discussion of dedicated mining hardware.

Each operator therefore has two options: either buying CPUs for useful work, or buying CPUs for farmed mining (i.e. farming). For both scenarios, operators can choose freely to use CPUs of any age. Denote by $f^m \in [0, 1]$ the ratio of operators that choose to mine. Denote by age_w and age_F the optimal age for doing useful work and farming, respectively.

Due to symmetry, the optimal age of CPUs at all miners is the same, and so their expended mining power is the same. The mining revenue is therefore distributed uniformly among all miners, and a single miner’s income is simply R_{annual} divided by the total number of mining CPUs, leaving a total mining revenue of

$$P_m(u, \text{age}, f^m) = N^{\text{cpu}}(\text{age}) \cdot \left(\frac{R_{\text{annual}}}{f^m \cdot N^{\text{op}} \cdot N^{\text{cpu}}(\text{age}) \cdot u} - C(\text{age}) - O_{\text{farm}} - E(u, \text{age}) \right). \quad (11)$$

Again, utilization is optimized at 1. The stable operation point is when the revenue from mining equals that of honest work (12), where operators are not motivated

to switch sides,

$$P_{\text{useful}}(1, \text{age}_w) = P_m(1, \text{age}_w, f^m) . \quad (12)$$

Solving (12) we obtain $f^m(\text{age}_m)$ as a function of age_m .

We still need to find the optimal operation parameters for mining. To this end, we calculate the symmetric Nash equilibrium where if all operators who choose to mine at age age , a single operator cannot increase her revenue by operating at a different age. In fact, finding such an age amounts to finding an age that if all but one miner operate at age, then the other miner's optimal operation age must also be age, that is,

$$\text{age} = \max_{\text{age}'} P_m(u, \text{age}', f^m(\text{age})) . \quad (13)$$

With the numbers from §B.2, we find the optimal age for mining is $\text{age}_F = 4.68$ years. At the equilibrium, 76.0% of the operators would be mining. The cost for one unit of useful work is \$1643.25. The results are summarized in Figure 15.

Figure 16 shows how the annual cryptocurrency revenue (R_{annual}) affects the miner ratio at the equilibrium, f^m , and therefore the cost of useful work. As the value of the currency and hence the mining revenue increases, the ratio of operators choosing to mine increases linearly, and the cost of useful work increases exponentially. As the miner ratio reaches 100%, the useful price goes to infinity as nobody is doing useful work.

Dedicated Hardware. Our analysis above assumes use of a CPU for mining, of course. But PoW functions that have been used for cryptocurrencies for extended periods, namely double-SHA256 (for Bitcoin) and Scrypt (for Litecoin), led to the development of dedicated hardware, particularly Application-Specific Integrated Circuits (ASIC). By design, these are far more efficient than any generic hardware for the purpose of the given PoW, but cannot be used for anything else. Arguably, for any PoW function, dedicated hardware can outperform generic hardware.

The equilibrium point in our analysis, however, is not actually affected by the performance or cost of dedicated hardware. The reason is that no matter what mining devices are used, a fixed amount of mining revenue is evenly distributed among the mining operators at the equilibrium, leaving the mining revenue unchanged. Therefore, advances in dedicated hardware do not affect the conclusions of our analysis.

C.1.2 Proof of Elapsed Time (PoET / Lazy-PoET)

As explained above, Intel proposed PoET as a waste-free PoW replacement for a permissionless blockchain. In

PoET, each CPU draws a random number r and sleeps for r time. Whichever gets the smallest number wakes up first and becomes leader for the next consensus epoch. Building on top of Intel's trusted hardware SGX, PoET makes use of the trusted random source protected by hardware, prohibiting selfish actors from increasing the frequency of their blocks, but with minimal computation. But the vanilla PoET proposed by Intel cannot be directly employed. The most critical issue is that it costs nothing for a miner to mine on multiple branches of a blockchain. As has been studied in the context of Proof of Stake, being able to work on multiple branches forces a strong assumption, namely that a majority of the miners blindly follow the protocol, even if each of the members of this majority is not individually motivated to do so. He, Song, and Milutinovic [42] proposed to fix this issue by using SGX's monotonic counters. They argue that depleting all 256 SGX counters ensures that the CPU is tied to a single branch.

We call their patched PoET scheme that maintains the CPU idle while mining *Lazy-PoET*. The point of *Lazy-PoET* is to allow for mining without any energy waste. In the *Lazy-PoET* scheme, an operator can use her hardware for either useful work, with fully-utilized CPUs, or for mining, with idle CPUs. Idle CPUs are cheaper to operate in farming mode, and the mining revenue of an old CPU is the same as that of a new and expensive CPU. Hence, the operator choices are reduced to either farming, or useful-work.

The revenue expression for *Lazy-PoET* is the same as with Proof of Work (11), but here the optimal age is age_{max} , since there is no benefit in using newer CPUs, and old CPUs are cheaper. Therefore, the revenue is only a function of the miner ratio,

$$P_m^{\text{Lazy}}(f^m) \triangleq P_m(0, \text{age}_{\text{max}}, f^m) . \quad (14)$$

Equating the revenues of mining (14) and useful work (10), we obtain again a miner ratio of 76% and the cost per unit of useful work is \$1643.4, i.e. a waste of 4.2. The results are summarized in Figure 15. Note that the numbers are identical to those of PoW, as the formulas are the same modulo age_{max} . The waste here is due to what we call the stale chip problem — farming with extremely old chips, insufficient for any useful work, yields high revenue. This optimum is robust to mining income changes. Mining is always optimal on stale chips. The annual mining revenue determines the ratio of miners; if it is too small, useful work simply becomes preferable to any sort of mining.

C.1.3 Busy Proof of Elapsed Time (Busy-PoET)

With Lazy-PoET, mining CPUs are kept idle. However, we note that the depleted SGX counters are allocated per-enclave, and not for the CPU itself: Different enclave code accesses different SGX counters. This is sufficient to ensure that the SGX is tied to a specific blockchain branch, since the mining enclave for a blockchain is unchangeable, preventing two from running on the same SGX and depleting the same counters. However, the implication is that the SGX and the CPU itself remain available for any other purpose, while mining and providing proofs of elapsed time as necessary. We call a system with proof-of-elapsed time where miners can concurrently use their hardware for useful work *Busy-PoET*.

Since the mining revenue adds no overhead in Busy-PoET, useful work with mining (henceforth *offhand mining*) is preferable to useful work without it. Nonetheless, farming remains a viable option because stale chips are significantly cheaper than recent ones.

As in previous schemes, the annual revenue R_{annual} is split among all mining CPUs. Denote by f^s , the ratio of operators that are working in standard mode. Unlike the previous schemes, here the CPU count includes the $f^s \times N^{\text{op}}$ operators in standard mode and the $(1 - f^s) \times N^{\text{op}}$ operators in farming mode. It is therefore a function of the CPU ages of both standard mining, age_S , and farming, age_F . We defer the development of this expression until after the discussion of both operator modes.

Standard mining Distinguished from PoW, the revenue from standard mining in Busy-PoET has an additional item of useful work revenue R_u :

$$R_w(u, \text{age}_S) + R_m - (C(\text{age}_S) + O_{\text{std}} + E(u, \text{age}_S)). \quad (15)$$

The total number of CPUs a standard operator can afford, $N_{\text{std}}^{\text{cpu}}$, has the same expression as in useful work (9), yielding a total annual revenue of (the dependency in f^s is through R_m)

$$P_S(u, \text{age}_S, f^s) = N_{\text{std}}^{\text{cpu}} \times (R_w(u, \text{age}_S) + R_m - (C(\text{age}_S) + O_{\text{std}} + E(u, \text{age}_S))). \quad (16)$$

Farming The expression for revenue from farmed mining is similar to that of previous schemes (the dependency in f^s is through R_m),

$$P_F(u, \text{age}_F, f^s) = N_{\text{farm}}^{\text{cpu}} \times (R_m - (C(\text{age}_F) + O_{\text{farm}} + E(u, \text{age}_F))). \quad (17)$$

The equilibrium analysis for Busy-PoET is not as simple now, as (16) and (17) are interdependent through f^s .

An equilibrium is a pair $(\text{age}_S, \text{age}_F)$ where an operator cannot improve her revenue by changing her strategy. This amounts to two conditions.

First, she cannot improve her revenue by changing her CPU age (Equation 13) where all other operators maintain their strategy. This is expressed for the two operation modes as the equation system

$$\begin{cases} \text{age}_S = \max_{\text{age}'_S} P_S(1, \text{age}'_S, f^s(\text{age}_S, \text{age}_F)) \\ \text{age}_F = \max_{\text{age}'_F} P_F(0, \text{age}'_F, f^s(\text{age}_S, \text{age}_F)). \end{cases} \quad (18)$$

Second, she cannot improve her revenue by changing her operation mode. We take this condition,

$$P_{\text{std}}(u, \text{age}_S, \text{age}_F) = P_{\text{farming}}(u, \text{age}_S, \text{age}_F),$$

to obtain an expression of f^s as a function of age_S and age_F . The resulting expression is too complex to be placed here.

We numerically solve the equation system and observe that, as with previous schemes, the cost of useful work grows quickly with the annual cryptocurrency revenue. Figure 16 shows that the ratio of standard miners decreases, as miners prefer farming with many cheap CPUs over performing useful work – another instance of the stale chip problem.

C.1.4 Proof of Potential Work

In order to mitigate the stale chip problem of the PoET variants, we propose a direct solution: Grant more block rewards to CPUs with better performance. Since with this approach a miner proves her CPU power, though she might not be utilizing it, we call it *Proof of Potential Work* (PoPW).

Technically, an SGX program can determine the CPU model by calling the `cpuid` instruction. Based on the CPU model, one can determine the value of a CPU by looking up in a public table, hard-coded in the blockchain protocol. The time for the mining enclave to return is therefore chosen with an exponential distribution parameter that is proportional to the CPU's power. The mining revenue is therefore linear in the slow-down factor of the CPU.

Note that PoPW makes stronger security assumptions than any of the other schemes we discuss. The principal that determines the values of the ever-changing CPU value table has significant power over the blockchain; for example, he can attribute high values to CPUs to which it has better access than other operators.

Standard mining As in Busy-PoET, standard operation dominates useful work, as the revenue for mining

comes without any overhead. Recalling that the slow-down of a CPU is $\eta(\text{age})$, the expression for the CPU revenue in standard mode is

$$R_w(u, \text{age}_S) + R_m \times \eta(\text{age}_S) - (C(\text{age}_S) + O_{\text{std}} + E(u, \text{age}_S)). \quad (19)$$

The total number of CPUs an operator can afford has the same expression as in PoW, yielding a total annual revenue of

$$N_{\text{std}}^{\text{cpu}} \times (R_w(u, \text{age}) + R_m \times \eta(\text{age}) - (C(\text{age}_S) + O_{\text{std}} + E(u, \text{age}_S))). \quad (20)$$

Farming The CPU revenue for farming is

$$R_m \times \eta(\text{age}_F) - (C(\text{age}_F) + O_{\text{std}} + E(u, \text{age}_F)). \quad (21)$$

The expression for number of CPUs is the same as before.

As before, the annual mining revenue is distributed among all mining CPUs, though now proportionally to their slow-down. The equilibrium analysis is similarly to that of Busy-PoET, resulting in a quick increase of waste as the annual cryptocurrency revenue grows, as show in Figure 16.

As shown in Figure 15, we observe that indeed the stale-chip problem is resolved: now the optimal farming age is 5.61, as opposed to age_{max} in Lazy-PoET and Busy-PoET, but the farming problem remains. Starting at some point, when the annual revenue is high enough, it becomes more profitable to farm, keeping idle CPUs, then to spend that amount on power for useful work. Also as discussed before, the trust model of PoPW is arguably too strong for a decentralized system.

C.2 Proof of Useful Work

Our solution, *Proof of Useful Work (PoUW)*, avoids the issues with previous schemes by directly counting work done towards mining effort. Since mining revenue is only granted when work is done, farming without useful work means the CPUs must be processing useless work in this mode. Therefore, standard operation dominates farming, as the revenue from useful work comes at no additional cost.

The useful work analysis and optimal age calculation remain unchanged, resulting in a revenue as expressed in (10).

For standard operation, the mining revenue now depends on the utilization of the CPU as well as on its age. The useful work utilization suffers a decrease due to the overhead of online effort monitoring, denoted O_{counting} .

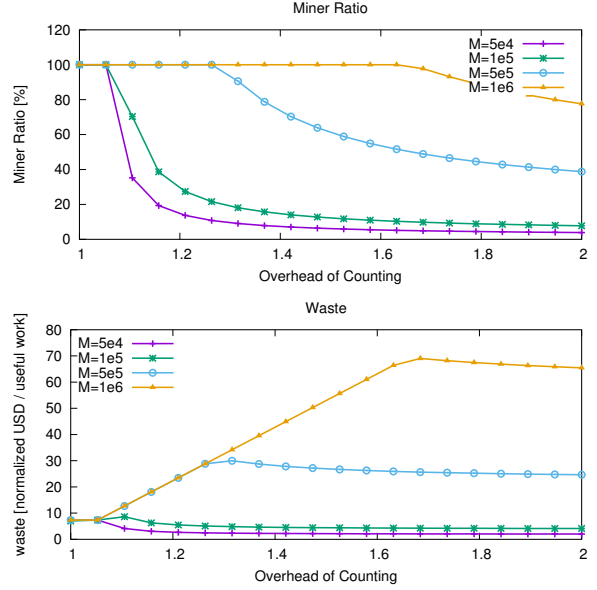


Figure 17: O_{counting} affects the miner ratio and the waste in PoUW.

The mining revenue does not suffer such a reduction, as the overhead is taken into account when calculating the mining effort.

Denote the optimal age for standard operation by age_S . The expression for the number of CPUs is as usual (9), at $\text{age} = \text{age}_S$. The standard mining revenue is the product of the CPU count and (22). This expression is a function of R_m and age_S ,

$$R_w\left(\frac{u}{O_{\text{counting}}}, \text{age}_S\right) + R_m \times \eta(\text{age}_S) \times u - (C(\text{age}_S) + O_{\text{std}} + E(u, \text{age}_S)). \quad (22)$$

As with PoW, the value of R_m is chosen such that the total revenue distributed is R_{annual} ,

$$R_m = \frac{R_{\text{annual}}}{f^s \times N^{\text{op}} \times N_{\text{std}}^{\text{cpu}} \times u \times \eta}. \quad (23)$$

The equilibrium point is where the standard and useful-work revenues are the same. Solving the equation we obtain an expression for the ratio of standard miners at equilibrium, which is a function of age_S . We proceed to find a symmetric Nash equilibrium as is done for PoW.

With the parameter values of Section B.2, we find that all operators choose to work in standard mode. The optimal CPU age is 4.39 years, rather than the 3.48 optimal age for useful-work operation. The average cost for useful work is \$430.1, i.e. a waste of 1.1. The results are summarized in Figure 15. As shown in Figure 16, PoUW has the lowest waste among the four schemes we

compared. Unlike other schemes, the waste in PoUW is not affected by the annual cryptocurrency revenue. The reason for that is PoUW effectively encourages all of the participants to mine as long as the mining revenue isn't critically low. With any reasonable annual cryptocurrency revenue, all of PoUW participants end up mining, yielding a fixed waste per useful work unit.

However, PoUW does introduce the overhead of secure instrumentation (O_{counting}) that is not present in other schemes. The impact of O_{counting} with different annual cryptocurrency revenue is shown in Figure 17. The top graph shows how O_{counting} impacts the miner ratio at equilibrium. Given an annual cryptocurrency revenue, higher O_{counting} discourages participants from mining as doing useful work would be more profitable. So the ratio of miners at equilibrium decreases with O_{counting} . The bottom graph shows how O_{std} impacts the waste. Because waste is only incurred by mining, higher annual cryptocurrency revenue increases miner ratio, leading to more waste. Meanwhile, given an annual cryptocurrency revenue, increasing O_{counting} will first increase the waste until all miners are forced to do useful work, where no waste is present.

The conclusion is high O_{counting} could diminish the security of PoUW because of a loss of miner power and an extra waste. As our implementation suggests, REM only incurs a O_{counting} of about 5 – 15%, i.e. 1.05 – 1.15 in Figure 17, allowing for high miner participation and low waste.

C.3 Unbounded analysis

In previous analysis, we assumed a bounded number of operators in the system. This assumption is grounded in the fact that there are certain barriers to enter the mining business, For completeness, we now propose an alternative model where the number of operators is unbounded. In this model there will be infinitely many CPUs doing useful work, and any CPU can switch to that at any point. Weakening this assumption only strengthens the results.

Admittedly, it is tricky to propose a perfect model for such a dynamic and pluralistic system. So we provided a first-order approximation based on realistic but simplifying assumptions. For example, the low electricity cost in China is quite important to the dynamics of Bitcoin mining. but the details of price distribution are unknown. Therefore we assumed the cost of CPUs and electricity is the same for all of the operators.

Just as before, participant has two options besides working on useful work: mining on herself or joining a farm. The age of CPU still plays an important rule as the price and performance vary significantly with the age. A equilibrium is defined similarly as before, satisfying two conditions: 1) revenue for participants is the same as if

they were useful workers; and 2) no single operator can earn more by working at a different CPU age.

The analysis for PoW and Lazy-PoET remains basically unchanged because in both schemes operators have in fact only one option besides useful work. We refer readers to the Appendix for more details.

Busy-PoET We note that the incremental overhead of Busy-PoET over useful work and the cost of farming are key to the analysis. Taking it to an extreme where PoET incurs zero additional overhead, every CPU doing useful work will mine along with the useful work, earning some mining revenue at no incremental cost. On the other hand, if the cost incurred by PoET is high, fewer participant will perform PoET at equilibrium, hence leaving some profit margin for farming. We define the inefficiency of PoET as the incremental overhead over doing useful work normalized by useful work revenue. The efficiency of PoET is one minus that.

Fig. 18 shows the number of standard mining CPUs at equilibrium, with PoET efficiency ranging from 60% to 100%, assuming no farming. We argue that in the case where PoET is very efficient, unbounded analysis is actually more realistic, as the equilibrium point derived from the unbounded model requires an impractically large amount of CPUs.

Approximating the situations in Bitcoin, if we assume there are farmers (or attackers) with access to cheap stale chips and electricity, the presence of them could skew the equilibrium significantly. At each point in Fig. 18, standard miners operate at the margin where adding one single CPU renders the standard mining a worse option than not participating. Therefore, a single farmer with cheap resources would expel a large amount of standard mining CPUs from the mining pool. We argue that this is a major drawback because it allows attacks at a relatively low cost.

Another factor that would change the equilibrium is the farming cost. In the extreme case where PoET incurs zero overhead, farming is no longer a viable option, unless it's free, because most of the mining revenue would have been harvested by standard miners, leaving farmers too little to cover the farming cost. But on the other hand, if PoET incurs non-negligible overhead, farming becomes possible if the farming cost is low enough. Fig. 19 shows the thresholds for farming cost as functions of PoET efficiency. For any given PoET efficiency, if farming cost is below than the lower threshold, all operators will end up farming at equilibrium; above the higher threshold, all operators will end up performing standard mining. If the farming cost is between the two, then equilibrium will involve a mixture of both.

PoUW As discussed previously, PoUW renders farming irrelevant, eliminating a major source of waste. How-




Figure 18: The number of standard mining CPUs at equilibrium, as a function of PoET efficiency. R_{annual} denotes the total annual cryptocurrency revenue. As a reference, Bitcoin yielded a total annual revenue of approximately 330 millions in 2015. As PoET efficiency goes to 1, the number of mining CPUs tends to infinity.




Figure 19: Farming cost affects the number of farmers at equilibrium. If the farming cost is below threshold low, all operators will join farms at equilibrium; above threshold high, all operators will do standard mining. Useful work cost is \$391.25.

ever, PoUW does introduce an overhead O_{counting} from instructions counting. Fig. 20 shows the number of CPUs in PoUW at equilibrium under different PoUW efficiency. Note that the equilibrium won't be skewed by farmers as farming in PoUW is strictly inferior, irrespective of farming cost.

The conclusion here is that if PoET and PoUW are sufficiently efficient, PoET can get rid of the farming issue unless farmers can operate at a very low cost. In this case, however, the bounded model is more realistic because equilibrium in the unbounded model requires an impractically large amount of participating CPUs. On the other hand, if PoET incurs non-negligible overhead, unbounded analysis draws a similar conclusion as the bounded one does: farming remains a viable option and because of that, an attacker can accrue old CPUs to attack at a lower cost.

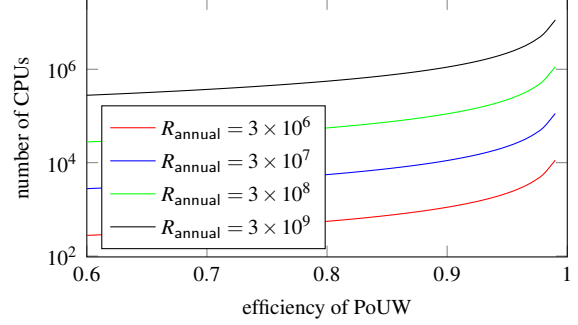


Figure 20: The number of standard mining CPUs at equilibrium, as a function of PoUW efficiency (i.e. $1 - O_{\text{counting}}$). R_{annual} denotes the total annual cryptocurrency revenue. As a reference, Bitcoin yielded a total annual revenue of approximately 330 millions in 2015. As PoUW efficiency goes to 1, the number of mining CPUs tends to infinity.

D Implementation Details

D.1 Experiments Setup Details

Computational protein folding enables prediction of the physical structure of chains of amino acids in protein molecules. It has various uses in the study of diseases such as Alzheimer's and cystic fibrosis as well as various types of cancer [66]. In our first experiment, we used an implementation [1] of the protein folding algorithm proposed by [76]; this software was a C++ program of ~ 800 LoC. To achieve SGX compliance, we removed all system calls. We instrumented the code automatically with our compilation toolchain. Finally we compared the execution of the instrumented and non-instrumented code to evaluate the overhead. The results are given in Figure 9.

In a second experiment, we ported a widely used SVM implementation `libsvm` [27] (with ~ 4300 LoC), again manually extracting system calls, and also replacing its RNG calls with those for the SGX RNG. Then we trained an SVM classifier on the popular Mushroom dataset from the UCI Machine Learning Repository [12] and deployed it in the enclave. The useful work in this case is to classify inputs. We queried the classifier on a test set of $\sim 8k$ points and measured the execution time.

In the third experiment, we ported an implementation [2] of `zlib` algorithm (the DEFLATE algorithm) to SGX and to REM. The program consists of ~ 5000 lines of C code and we removed its reliance on standard libraries manually. The useful workload in this experiment is to compress a random string of 50 MB. The experiment is repeated for 1000 times.

In the last experiment, we ported a C implementation [10] of the SHA-3 message digesting algorithm. The useful workload is to calculate SHA3 digests for a 100MB buffer of random bytes. The experiment is re-

```

...
.LEHB0:
    leaq    1(%r15), %r15    # added by PoUW
    call   _ZN11stlpmx_std12basic_stringIcNS...
.LEHE0:
    .loc 7 70 0 is_stmt 0 discriminator 2
    leaq    3(%r15), %r15    # added by PoUW
    leaq   -80(%rbp), %rax #, tmp94
    movq    %rax, %rsi # tmp94,
    movq    %rbx, %rdi # _4,
.LEHB1:
    leaq    1(%r15), %r15    # added by PoUW
    call   _ZN11stlpmx_std12out_of_rangeC1ER...
.LEHE1:
...

```

Figure 21: Code snippet: instrumented assembly code of a protein folding program.

peated for 1000 times.

D.2 Code Snippet

Figure 21 shows a snippet of assembly code instrumented with the REM toolchain. Register `r15` is the reserved instruction counter; it is incremented at the beginning of each basic block in the lines commented added by PoUW.