

A scalable verification solution for blockchains

Jason Teutsch
TrueBit Foundation
jt@truebit.io

Christian Reitwießner
Ethereum Foundation
chris@ethereum.org

March 7, 2017

Abstract

Bitcoin and Ethereum, whose miners arguably collectively comprise the most powerful computational resource in the history of mankind, offer no more power for processing and verifying transactions than a typical smart phone. The system described herein bypasses this bottleneck and brings scalable computation to Ethereum. Our new system consists of a financial incentive layer atop a dispute resolution layer where the latter takes form of a versatile “verification game.” In addition to secure outsourced computation, immediate applications include decentralized mining pools whose operator is an Ethereum smart contract, a cryptocurrency with scalable transaction throughput, and a trustless means for transferring currency between disjoint cryptocurrency systems.

Contents

1	Securing computations with economics	2
1.1	Outsourced computation	3
1.2	Practical impact	4
1.3	Smart contracts	5
2	How TrueBit works	5
2.1	System properties	7
2.2	Assumptions	8
2.3	Attacker model	9
3	Dispute resolution layer	10
3.1	Bottleneck: The Verifier’s Dilemma	10

3.2	Solution: The verification game	12
3.3	Detailed protocol	13
3.4	Runtime and security analysis	15
4	Incentive layer	17
4.1	Jackpots	18
4.2	Taxes	19
4.3	Deposits	20
4.4	Generating forced errors	21
4.5	Solver and Verifier election	22
4.6	Protocol overview	23
4.7	Sanity check	26
5	Defenses	26
5.1	Pairwise Sybil attacks	28
5.2	The trifecta	30
5.3	Collusion pools	32
5.4	On low-hanging fruit	34
5.5	A cash equivalence problem	35
6	Implementation	36
7	Applications	36
7.1	Practical decentralized pooled mining	37
7.2	Dogecoin–Ethereum bridge	38
7.3	Scalable transaction throughput	38
7.4	Towards a big data system	38

1 Securing computations with economics

Every Nakamoto consensus-based cryptocurrency (e.g. Bitcoin or Ethereum) offers something that everyone in the world agrees on: a definitive, public ledger of financial transactions, or *blockchain*. This consensus technology enables basic Bitcoin transactions, which transfer currency from one party to another, while Ethereum transactions perform financial and database operations contingent on the evaluation of more complex computational scripts. Anonymous miners, who freely join and leave cryptocurrency networks, determine the validity of these transactions and thereby establish who owns which coins. Remarkably, this verification scheme requires no central authority to operate securely.

In practice, miners can successfully maintain blockchain integrity so long as the computational burden of verification remains minimal. Indeed, Bitcoin and Ethereum, whose miners arguably collectively comprise the most powerful computational resource in the history of mankind, offer no more power for processing and verifying transactions than a typical smart phone. One cannot simply increase the volume or complexity of transactions flowing into the blockchain without risking inclusion of invalid transactions due to the so-called *Verifier’s Dilemma* [42] (see Section 3.1). In 2015, the Verifier’s Dilemma manifested itself in the form of the July 4 Bitcoin fork [28] which temporarily derailed the entire network. The Ethereum community also witnessed a related exploit in its 2016 denial-of-service attacks [15].

A *consensus computer* [42] permits users to outsource computations to the Ethereum network and receive *correct* answers in exchange for payment so long as the effort required to verify solutions does not exceed the threshold induced by the Verifier’s Dilemma. Thus the trustless consensus computer offers a small but reliable kernel of semantic truth.

Our contribution. We present here a system, called *TrueBit*, which amplifies the consensus computer’s capabilities. TrueBit enables trustless smart contracts, in theory, to securely perform any computation task. Moreover, TrueBit vastly reduces the number of redundant network node computations used in traditional Ethereum smart contracts. Presently every Ethereum miner has to independently replicate each smart contract action in its entirety, whereas TrueBit outsources most computation work to a handful of entities. In this way, TrueBit makes secure computing affordable.

1.1 Outsourced computation

Let us ponder for a moment why cryptocurrencies might offer an especially convenient framework for secure outsourced computation, the core application of TrueBit. In the traditional cloud models, users must trust that the machine hardware, software, and cloud administrator all perform as expected. A wide range of things can go wrong, particularly when one wishes to tie the results of such computations to monetized entities such as smart contracts. Proper economic incentives, the cornerstone of any cryptocurrency, can deter many types of errors from occurring in ways that simple task repetition cannot. Furthermore, in contrast to a cloud whose configuration details may not be visible to users, any systemic network errors that might occur on a blockchain-based system like TrueBit would appear in plain

sight to the entire community. Cryptocurrencies indeed provide an excellent starting point as they already achieve several desirable properties.

1. As witnessed by the Ethereum consensus computer, Nakamoto consensus grants networks the ability to trustlessly perform small computations correctly.
2. Blockchain public ledgers provide perfect transparency and immutability. Ethereum smart contracts inherit these characteristics.
3. Currency incentives, which in cryptocurrencies are inextricably tied to computational processes, can be used to recruit and reward participants.

In general, economic forces greatly reduce the scope of possible network errors. We can safely assume that errors which incur economic penalties are much less likely to occur than those which do not.

1.2 Practical impact

A market for work related to computationally-intensive smart contracts already exists. The Golem Project, which crowdfunded 820,000 ETH on the morning of November 11, 2016 (~\$8.6 million on that date), already cites TrueBit as a verification mechanism for their forthcoming outsourced computation network [13]. TrueBit has the potential to support many practical new applications. Here are some examples.

- *Outsourced computation.* TrueBit functions as a worldwide computation market. Anyone in the world can post a computational task, and anyone else can receive a reward for completing it. The system's incentive structure guarantees correctness of returned solutions.
- *Truly decentralized mining.* Centralized mining pools pose security threats. For any Nakamoto consensus-based cryptocurrency, one can use TrueBit to build a practical, efficient, and trustless mining pool managed by an Ethereum smart contract (see also [43]).
- *Trustless cryptocurrency exchange.* Currently users must find trusted exchange partners to transfer currency (or messages) between blockchains. TrueBit can facilitate secure transfer, for example, of dogecoins to the Ethereum blockchain and back (modulo some new opcodes for dogecoin) [7].

- *Scalable blockchain.* By decoupling verification for miners into a separate protocol, we can achieve high transaction throughput without facing a Verifier’s Dilemma.
- *Scalable “on-chain” storage.* Swarm [22] is developing a platform for incentivized, peer-to-peer storage. TrueBit can make Swarm data accessible to Ethereum smart contracts.

We discuss these ideas further in Section 7.

1.3 Smart contracts

Smart contracts, introduced in 1994 by Szabo [52] and first realized in the Ethereum [10] cryptocurrency in 2015, are uncensorable programs that live on Ethereum’s blockchain and have their own executable code and internal states, including storage for variable values and ether currency balance. Proposed applications for smart contracts include outsourced computation and storage [12, 22, 42], prediction markets [1], decentralized venture funds [23], and Bitcoin mining pools [21, 43]. Ethereum’s variable `gasLimit` restricts the number of computation steps that smart contracts can perform.

TrueBit itself is an Ethereum smart contract which allows users to call *TrueBit contracts* for trusted, computationally-intensive applications. Traditional Ethereum smart contracts can call a TrueBit contract as a subroutine, thereby effectively bypassing Ethereum’s `gasLimit`. In other words, TrueBit increases the per block work that the Ethereum network can process correctly.

2 How TrueBit works

TrueBit’s primary purpose is to realize correct, trustless computations despite miners’ limited computation bandwidth. Intuitively, we wish to reward participants who correctly perform computational tasks, but who decides whether these tasks were done correctly? In absence of a dispute, the party who performs a computational task on behalf of a TrueBit contract simply receives a reward. On the other hand, if a dispute does occur, we must rely on the only trusted resource, the limited network of miners, to resolve it. Dispute resolution occurs as a “verification game” subroutine in TrueBit, and we will discuss its details in Section 3.

Who initiates a dispute in the case of an incorrect computation, and how can we ensure that such a dispute will actually occur? One option

is to assume that each party with some stake in the contract brings their own trusted (although not necessarily mutually trusted) verifier, and flags for a challenge if needed. This trusted verifier approach suffices for some applications, but in general parties may need to rely on the results of a previously executed TrueBit contract *ex post facto* even if they themselves did not exist at the time of its creation. Like Ethereum smart contracts, TrueBit contracts must have universal validity.

Since there exist no trusted parties on Ethereum’s network, by symmetry we must allow any party to be hired to solve any computational task, and similarly anyone should be able to challenge a Solver’s outcome of a computational task. The latter requirement ensures that TrueBit operates by unanimous consensus. Assuming that an honest and motivated Challenger exists who will challenge any wrong answer given by the Solver, the system’s incentives for the Solver are straightforward: reward correct solutions. But how can we make sure that such a motivated Challenger exists?

Idea 1. Offer rewards for checking computations.

While this strategy may encourage participation, it provides no incentive for correct verification. We have no *a priori* reason to trust that the Verifier will substantively inspect the computation task, and we gain no additional assurance by increasing her reward for checking.

Idea 2. Offer rewards for finding bugs.

While a bug bounty may incentivize a Verifier to correctly perform a check, it only does so when the Verifier has some reason to believe that a bug might actually exist. Unless a potential Verifier believes that she has some real chance to find a bug, in practice we cannot expect her to participate in the system. In summary, we need both an incentive to partake in verification and an incentive to perform checks correctly. This leads us to the following proposition.

Idea 3. Offer a bug bounty *and* provide expectation that bugs will exist.

TrueBit takes this third approach by occasionally forcing Solvers to submit incorrect solutions. During “forced errors” TrueBit reverses the normal system incentives: the Solver gets paid for submitting an incorrect solution and penalized for submitting a correct one. We discuss the details of TrueBit’s forced error protocol in Section 4.

Let us return to our dispute resolution subroutine. At the heart of TrueBit’s protocol lies an interactive “verification game” which decides whether or not a contested computational task was performed correctly (see

Section 3). The verification game proceeds through a series of rounds, where each round recursively checks a smaller and smaller subset of the computation. A trusted network, in our case the Ethereum platform [10], merely enforces the rules of the game and therefore does not bear the bulk of the verification burden. Anyone on the network can generate tasks, compute, or verify in exchange for rewards. TrueBit does not require “honest” or “altruistic” nodes for correct operation but rather runs under the assumption that each node wishes to maximize its own profit. We shall discuss the details of TrueBit’s incentives for participation in Section 4.

2.1 System properties

Any outsourced computation system should be *fair* in the sense that parties who perform computational tasks indeed receive compensation for their work and *reliable* in the sense that paid computations are performed correctly. In addition to these properties, TrueBit also ensures *accessibility* through easy joining or exiting of the verification ecosystem. Any *anonymous* party can play the role of Task Giver, Solver, or Verifier, where a Verifier is a party who checks solutions and becomes a Challenger whenever she reports an error. In particular, TrueBit does not trust or rely on the reputation of its participants. Everyone who puts a deposit into the system has a fair chance to be hired for a given computational task.

TrueBit offers several novel advantages over traditional cloud computing and verifiable computable models. Verifiable computing ensures a correct answer for an outsourced computation by forcing the cloud to provide a short proof which witnesses the correctness of its computation. The idea is that this “short proof” should be much easier to check than performing the computation oneself. Researchers have achieved much progress on this method in recent years, however the cryptographic setup costs and computational overhead for the cloud in state-of-the-art systems make these methods unsuitable for most real-world applications. Moreover, many of the proof-based systems to-date, including Zaatar, Pinocchio, Ginger, and TinyRAM, require one to run thousands of instances of a single function before breaking even on the overall verification time for a 128×128 matrix multiplication puzzle [55]. The new cryptocurrency Zcash [27] successfully makes use of verifiable computing machinery, albeit only for a very restricted class of computation tasks. Below we contrast TrueBit with verifiable computing and traditional cloud computing.

1. *Incentives.* Unlike traditional cloud computing, where the user simply

trusts the cloud to provide correct answers, TrueBit provides financial incentives to ensure correctness.

2. *Transparency.* The entire inner workings of TrueBit’s interpreter sit on the blockchain and are open for inspection (see Section 6). Furthermore, the user community can democratically update the interpreter as needed.
3. *Efficiency.* Solvers in TrueBit have low computational overhead and minimal initial setup costs. The verification game (Section 3) does introduce some extra work, but in practice, due to high penalties for wrong answers and bogus challenges, we expect participants to appeal to the verification game only rarely, if at all.
4. *Simplicity.* TrueBit’s operation is relatively straightforward. Unlike traditional verifiable computing, TrueBit avoids deep probabilistically checkable proofs (PCPs), succinct non-interactive arguments of knowledge (SNARKs) [30], and exotic cryptographic assumptions (e.g. those used in zkSNARKs [31]). The standard cryptography used in TrueBit, namely hash functions and digital signatures, sits in the underlying blockchain network and does not surface in the TrueBit protocol itself.
5. *Adaptability.* TrueBit runs on top of Ethereum’s current protocol without impacting functionality.
6. *Keyless entry.* Participants do not need to manage cryptographic keys beyond those used in Ethereum. TrueBit establishes identities through financial deposits alone, and therefore the system avoids risks from cryptographic trapdoors.

2.2 Assumptions

Traditional distributed systems models focus on tolerance against arbitrary, Byzantine errors. In Ethereum and other Nakamoto consensus-based cryptocurrencies, however, we have no reason to expect arbitrary errors—nodes generally deviate from protocol due to financial incentives. Our system model makes two basic assumptions about the underlying network.

- (I) There exists a trusted network (i.e. Ethereum) that correctly performs very small computation tasks.
- (II) Participants are rational in the sense that they act to maximize their individual profits. In particular, we expect that CPUs will show up to

do computational tasks if and only if they expect fair compensation for their work.

The consensus computer [42] shows how one can use the incentives in Ethereum to establish assumption (I) above both in theory and in practice. The task of our new system will be to amplify this small amount of correct computation in order to handle larger tasks using assumption (II). Even though solo Ethereum miners rarely earn mining rewards, expected long-term returns suffice to merit their participation in the system (via mining pools). As we shall discuss in Section 3.1, anonymous parties may not necessarily perform correct computations when economic incentives, including conservation of CPU cycles, pressure them to do otherwise. This observation justifies assumption (II). We emphasize that we do not find the assumption that there exists a single honest (and non-lazy) participant [38] to be realistic. A person who loyally contributes CPU cycles without incentive to do so most likely does not exist in a large, general purpose, public system.

This paper takes a somewhat simplified view of Nakamoto consensus. In some cryptocurrencies, such as Bitcoin, miners can selectively censor a class of transactions by deciding not to include its members in their blocks. In the case of TrueBit, censorship of challenges (see Section 4) or transactions in the verification game (Section 3.2) could result in catastrophic, bogus solutions being accepted on the blockchain. In Ethereum, however, censoring individual transactions is difficult as miners cannot easily see where an Ethereum transaction might call without executing it. Indeed Ethereum users can obfuscate the function of their transactions, making it computationally costly for miners to hand-pick transactions for their blocks. Hence we shall assume that censorship of individual transactions does not happen. In Section 5.1, we analyze an attack on TrueBit in which miners withhold entire blocks and show that it is not profitable.

2.3 Attacker model

TrueBit’s security relies on two basic assumptions about the behavior and capabilities of potential attackers.

- (I) *Attackers cannot subvert the underlying consensus computer.* The Judges in the verification game (Section 3) always adjudicate correctly, as do the Referees in the incentive layer (Section 4).
- (II) *Attackers are rational.* We assume that attackers have limited financial resources and will not deviate from protocol unless it is profitable for them to do so.

While an adversary with sufficient computational power can derail a Nakamoto consensus-based cryptocurrency [47], assumption (I) above neatly sweeps this issue outside our domain. TrueBit itself does not utilize proof-of-work. Although we have chosen to implement TrueBit in the Ethereum platform for convenience, we could easily replace the “Judges” from (I) with any other universally trusted computing mechanism without affecting the system’s functionality.

Due to anonymity of all participants, *Sybil attacks*, in which a single party deviantly assumes multiple identities on the network, pose special threats to TrueBit (see Section 5.1 and Section 5.2). We will assume that an attacker may create as many identities as she likes, including simultaneously playing the roles of Task Giver, Solver, and Verifier, appearing as two distinct Verifiers, or managing multiple identities via pooled financial resources (see Section 5.3). Using assumption (II) above, we shall argue that TrueBit resists such attacks in Section 5.

The “forced errors” described in the beginning of this section (Section 2) pose a special challenge to TrueBit due to opportunistic attacks. As the prize for discovering forced errors is necessarily enormous, attackers may seek ways to obtain these large prizes prizes without performing the intended verification tasks (Section 5.4). For this reason, the occurrence of forced errors must appear truly random to all prospective Verifiers. Moreover, Solvers must be incentivized from prematurely sharing the fact that their error was forced or claiming that an error was forced when it really wasn’t. In general, we want consistent, not sporadic, verification. At least one Verifier should be present for every task. The attraction of forced errors represents the “Achilles’ heel” of TrueBit, and we devote the incentive structure in Section 4 to careful defending of this potential attack surface.

3 Dispute resolution layer

In this section we present and analyze TrueBit’s dispute resolution mechanism. TrueBit relies on network “Judges” with limited computational power who adjudicate all disputes, and we show how such Judges can correctly resolve disputes over complex computations.

3.1 Bottleneck: The Verifier’s Dilemma

Ethereum’s incentive structure severely restricts the amount of computation work that smart contracts can accurately enforce. Users interact with smart contracts by broadcasting *transactions* to the network. Such transactions

contain program scripts which require computational work to verify their validity prior to miners' acceptance on the blockchain. Each miner must not only locally verify but also remember the state of every single smart contract. On the surface, this redundant work limits smart contracts' computational scope due to inefficiency, but the true bottleneck lies deeper at the incentive level.

Let us recall miners' incentives for participating in Ethereum. In Nakamoto consensus-based cryptocurrencies [47], new transactions enter the blockchain through a process called *mining*. Participating *miners* maintain the integrity of the blockchain data by racing to solve computation-intensive, *proof-of-work* puzzles in exchange for rewards. The first miner who successfully broadcasts a solution to the current proof-of-work puzzle proves that she has spent the necessary computation power to merit appending her new set of transactions to the blockchain, and this step awards the miner a set of newly minted coins. Then the race begins again on top of this new block.

When an Ethereum user broadcasts a transaction to the network, the miner who appends her valid transaction to the blockchain receives a *transaction fee* for the service. On the other hand, Ethereum's mining protocol dictates that other miners should verify this transaction *gratis*, for the "common good" so to speak. More specifically, the mining protocol dictates that the longest chain consisting of valid transactions is the correct one, and miners should mine on top of it. If the time to verify new blocks is negligible, miners may indeed agree to check them, but if the verification work becomes substantial, then miners risk falling behind in the mining race by following protocol and actually verifying. In order to save time, rational miners might simply skip the verification step, leaving the blockchain open to invalid transactions. On the other hand, by skipping verification, a miner might also waste her CPU cycles by mining on top of a stale chain which the rest of the community deems invalid. Either way, the miner risks forfeiting rewards. A rational miner doesn't know what to do!

Any substantial verification burden placed on miners thus results in a *Verifier's Dilemma* [42]. Because of the Verifier's Dilemma, smart contracts whose verifications require non-trivial computational efforts will fail to execute correctly as rational miners, who may choose to deviate from the suggested protocol, stop donating their limited resources to thankless verification tasks. In short, Ethereum lacks a scalable verification mechanism and remains vulnerable to serious attacks [42] when miners have to verify more than a tiny amount.

TrueBit contracts securely execute computationally-intensive tasks for use in Ethereum smart contracts. The system's core consists of a smart

contract interface where a user can request a solution to a computational task or puzzle, and anyone can solve it in exchange for a reward. Our interactive verification game, described below, empowers users to both outsource arbitrary computations in a decentralized manner and receive back correct solutions. To achieve this goal, we will not only reduce the redundancy of verification work but also refine the incentive structure responsible for the Verifier’s Dilemma.

3.2 Solution: The verification game

The goal of TrueBit’s interactive verification game is to resolve a given dispute between a Solver, who has provided a solution to some computational task, and a Challenger, who disagrees with the solution and therefore calls the verification game to be played. The outer layer of TrueBit (Section 4) uses the verification game as a subroutine. The roles in the verification game include a Solver, who offers a solution to a given task, and a Challenger who disagrees with the Solver’s solution. The final party, the Judges, always perform computations correctly but possess extremely limited computational bandwidth. The Judges in TrueBit are the entire community of Ethereum miners who reach verdicts through Nakamoto consensus.

The verification game proceeds in a series of rounds, each of which narrows down the portion of the computation in dispute. In the first round, the Challenger forces the Solver to commit to a pair of computation steps delimiting some time interval. In the next round, the Challenger iteratively challenges a subset of the Solver’s computation over this time interval, and next she challenges over a subset of that subset, etc. until in the final round the final challenge is sufficiently trivial that the Judges can make a final ruling on whether the challenge was justified. The Judges also enforce that the Solver and Challenger follow the rules of the game. At the end of the verification game, either the cheating Solver will be discovered and punished in the outer layer of TrueBit (Section 4), or the Challenger will pay for the resources consumed by the false alarm.

Example (matrix multiplication [38]). We give an illustrative example of a verification game. Suppose that the given task is to compute the product of two matrices A and B . The Solver claims that $A \times B = C$. In the first round, the Challenger must claim an error in one of the entries in C , and outputs coordinates i, j such that

$$c_{i,j} \neq \sum_{k=1}^n a_{i,k} \cdot b_{k,j}$$

with corresponding evidence consisting of partial sums d_0, d_1, \dots, d_n where

$$d_m = \sum_{k=1}^m a_{i,k} \cdot b_{k,j}.$$

The Judges only verify that i, j are coordinates and that $d_n \neq c_{i,j}$ and that $d_0 = 0$. The Challenger loses if this fails to hold. In the second round, the Solver tries to defend himself by providing a k such that $d_k \neq d_{k-1} + a_{i,k} \cdot b_{k,j}$. If the Judges can verify this claim then the Solver wins the game, and otherwise the Challenger wins.

3.3 Detailed protocol

We can use a variation of the Example above to check arbitrary computations by creating a binary search to pinpoint the erroneous (or not erroneous) step in a computation. The idea for the following verification game is essentially due to Canetti, Riva, and Rothblum [32, 33] and was later independently discovered by the authors of this paper [38, 49]. Canetti, Riva, and Rothblum did not consider this game in the context of blockchains but rather a simpler scenario in which a user outsources a computation to k different clouds under the assumption that at least 1 out of k of the clouds will perform the computation correctly. The assumption of at least one trusted Solver/Challenger cloud is too strong for our purposes since in a purely rational network, such a Solver/Challenger may not exist. Section 4 describes our workaround for the smart contract setting.

For purposes of reducing the number of rounds of interaction and space demands on the blockchain, our protocol roughly follows [32, 33] which combines the (parametrized) binary search optimization from [38, Proposition 9] with the probabilistic Merkle tree construction from [49]. Recall that a *Merkle tree* is a binary tree in which each node is the hash of the concatenation of its children nodes. In general, the leaves of a Merkle tree will collectively contain some data of interest, and the root is a single hash value which acts as a certificate commitment for the leaf values in the following sense. If one knows only the root of a Merkle tree and wants to confirm that some data x sits at one of the leaves, the holder of the original data can provide a path from the root to the leaf containing x together with the children of each node traversed in the Merkle tree. Such a path is difficult to fake because one needs to know the child preimages for each hash in the path, so with high likelihood the data holder will supply a correct path if and only if x actually sits at the designated leaf.

For clarity, we will present the verification game in the familiar and standard language of Turing machines, but in actuality TrueBit uses the Google Lanai architecture to simulate steps in the computation (see Section 6). The verification game consists of three parties:

- a *Solver* who has allegedly performed some computational task and claims some solution,
- a *Challenger* who disputes the correctness of the Solver’s solution, and
- *Judges* with bounded computational power, who will rule on whether the output is correct.

TrueBit fills the roles of Solver and Challenger according to the incentive layer discussed in Section 4. The Judges are the set of miners in Ethereum who collectively record their decisions on the blockchain. The incentive layer (Section 4) determines the financial consequences of the Judges’ verdict.

To start, the Solver and Challenger each privately compile a tableau of Turing configurations mapping each time step of the task to its full internal state. Assume that the task runs in time t steps and space s , where s bits suffice to include the complete state of the Turing machine at any given time, including tape contents, head position, and machine state. The game protocol includes a fixed, integer parameter $c > 1$ which says how many configurations the Solver broadcasts to the blockchain in each round of the verification game. More configurations per round mean fewer total rounds needed, however it also means that the total number of configurations sent to the blockchain increases. In particular, we reach diminishing returns as the data sent per round reaches Ethereum’s per block capacity.

The verification game protocol proceeds as follows.

Main loop. The following steps are done iteratively to pinpoint the source of disagreement. The protocol determines timeout periods within which Solvers and Challengers must respond. Failure to respond within time bounds results in immediate loss for the non-responding party.

1. The Solver selects c configurations equally spaced apart in time across the current range of the dispute. In the initial iteration, for example, the Solver selects configurations across the entire tableau at the following time steps:

$$\frac{t}{c}, \frac{2t}{c}, \frac{3t}{c}, \dots, \frac{ct}{c}.$$

He then produces c Merkle trees, each with s leaves, where the leaves constitute the complete machine state at each of these times, and broadcasts each of the roots of these Merkle trees to the blockchain.

2. The Challenger responds with a number $i \leq c$, indicating the first time step in this list that differs from her own, and broadcasts this number i to the blockchain.
3. The Judges check that the Solver indeed posted c Merkle roots to the blockchain and that the Challenger's value i satisfies $1 \leq i \leq c$. If either of these checks fails, the Solver, or respectively the Challenger, loses immediately.
4. The next iteration of the game continues similarly, but restricted to configurations between the $i - 1$ -st and i -th indexed configurations. Here we interpret a 0 as the computation's initial configuration.

Final stage. After $\log t / \log c$ rounds, the loop above converges to the first, disputed computational step, and the Judges explicitly check this step of the computation as follows. Suppose that the first disagreement occurs at time e . The Solver provides paths from the Merkle root for time e to its leaves containing:

- the location of the machine head at time e ,
- the content of the tape square under the machine head, the tape contents of its left and right neighbors, and
- the machine state at time e .

The Solver also provides the analogous paths for time $e - 1$. The Solver loses the dispute if the Judges find that any of these paths fail to be valid. Finally, using this information and the original task code (which existed on the blockchain prior to the verification game), the Judges check the computational step at time e and rule whether the challenge was justified.

3.4 Runtime and security analysis

We now show that the work done by the Judges is small compared to the work required to perform the task, and that the number of interaction rounds is modest. Note that the majority of effort from the Judges simply involves

recording data onto the blockchain. The Judges’ work is split over many transactions and therefore avoids any chance of a Verifier’s Dilemma.

We fix σ as the size of each hash in the Merkle trees, and p as the space required to store a single machine state for the computation. For computational tasks that run in time t and space s with c configuration roots posted per round, the protocol requires $\log t / \log c$ rounds of interaction on the main loop, writes $c\sigma + \log c$ bits to the blockchain during each iteration of the main loop, and writes $2 \cdot ((2\sigma + 1) \log s + 5 + p)$ bits in the final round (assuming that tape squares can contain either 0, 1, or be blank, so we need $3 \cdot \log 3 / \log 2 \approx 5$ bits to describe 3 of them). Summing these parts, the total time required for community-wide verification from the Judges is then

$$O \left[\frac{\log t}{\log c} \cdot (c\sigma + \log c) + 2 \cdot ((2\sigma + 1) \log s + 5 + p) \right]$$

where the hidden constant depends on the implementation of this protocol on the Judges’ local machines. This estimate does not include the space and time required to store the instructions for the initial computation task. We assume that this data already exists on the blockchain before the verification game begins.

In comparison with other proposed blockchain-based computation systems, or even PCP-based verifiable computing systems [55], the work imposed on the Judges by the verification game is small [11], as is the computational overhead required to join the system [54]. Moreover, the economic costs to the Task Giver, who pays only a Solver and a Challenger and not all Judges to process the full task, is modest, as are the computational overheads for the Solver and Challenger.

Note that larger values for σ reduce the chances of hash collisions. Hash collisions might allow a dishonest Solver to substitute configurations in his computation tableau with fake ones sharing the same Merkle roots. Two consecutive configurations in the tableau whose Merkle roots happen to agree with the Merkle roots of two consecutive configurations that the dishonest Solver wishes to apply as substitute suffice to disrupt the integrity of the verification game protocol. Indeed, the Challenger will agree with both Merkle roots, and the Judges will confirm that the transition between them is valid and therefore incorrectly reject the challenge. For a fixed size σ , the chance that such collisions occurs by accident even becomes likely for sufficiently enormous tasks.

One could entirely eliminate the security risk discussed in the previous paragraph by posting complete machine states on the blockchain rather than just Merkle roots, but this makes the protocol more expensive. Alternatively,

one could either increase the parameter σ or the number of states checked by the Judges in the final step. In this sense, the choice of σ bounds the maximum complexity of secure computations in TrueBit. We will also see in Section 4.1 that the capital held by the TrueBit contract as a “jackpot” for finding a forced error poses an additional constraint on computational capacity, but this value can scale as well. Finally, the effectiveness of TrueBit’s verification game may degrade for extremely complex tasks due to the computational limitations of the Challenger. If the expected execution time of a verification game exceeds a human life expectancy, for example, then potential Challengers may lack motivation to participate despite financial incentives. Note that even when a task itself is feasible, its corresponding verification game may not be as the verification game carries significant overhead. We anticipate that future versions of TrueBit may optimize the verification game for certain kinds of tasks so as to reduce this discrepancy (see Section 7.4).

Remark. The verification game provides a significant privacy-protecting effect in that only disputed parts of a given computation must touch the public eye [37].

4 Incentive layer

We now discuss the financial incentives which encourage anonymous parties to both contribute CPU cycles and perform requested computational tasks correctly. We use the dispute resolution layer from Section 3 as a subroutine and connect its role to the present construction. A *Verifier* is a potential Challenger, that is, someone who checks submitted solutions and calls for a challenge (see Section 3) if and when she detects an error. By assumption (II) in Section 2.2, Verifiers must receive adequate payment for verification tasks. As argued in the beginning of Section 2, however, Verifiers can only receive rewards when they actually find bugs. It follows that the reward for discovering errors must encompass amortized payment for all verification tasks checked, including those in which no bugs were found.

Any Ethereum user who offers a reward through TrueBit for performing a computational *task* is called a *Task Giver*. A party which offers a *solution* for performing these tasks in exchange for a reward is called a *Solver*, and, as indicated in the previous paragraph, Verifiers check that Solvers’ solutions are correct. Solvers and Verifiers will play the roles of “Solvers” and “Challengers” respectively from the dispute resolution layer (Section 3) whenever a Verifier initiates a verification game. *Referees* enforce the rules of

the TrueBit protocol in the incentive layer. While Ethereum miners enforce protocol rules both through the roles of “Judges” in Section 3 and the role of “Referees” in the present section, we use distinct labels for them because the layers and functions differ for these roles. The primary role of Judges is to interactively resolve a dispute, whereas Referees primarily enforce that Solvers and Verifiers timely submit appropriate data in the incentive layer.

In a nutshell. The main steps of a TrueBit contract are as follows. Items in quotes will be thoroughly explained in due course.

1. A Task Giver announces a task and offers a reward for its solution.
2. A Solver is elected by lottery, and he prepares both a “correct” and an “incorrect” solution to the task.
 - (a) If a “forced error” is in effect, the Solver reveals the incorrect solution on the blockchain.
 - (b) Otherwise the Solver reveals the correct one.
3. Verifiers check the Solver’s solution. They win a large “jackpot” payout if both:
 - (a) they correctly identify the solution as erroneous, and
 - (b) a forced error was in effect.
4. If no Verifier signals an error, then the system accepts the solution. Otherwise acceptance depends on the outcome of a verification game.

In the rest of this section, we discuss the security measures in TrueBit which ensure that participants comply with the skeletal procedure outlined above. A more complete description of the protocol appears in Section 4.6.

4.1 Jackpots

TrueBit periodically imposes *forced errors* in which a Solver must offer a wrong solution to a task (see Section 2). This ensures that Verifiers who diligently search for errors will eventually find them. Accordingly, Verifiers who correctly report forced errors receive substantial *jackpot* payouts. By design, Verifiers cannot predict when forced errors will occur and therefore have incentive to check all tasks carefully. Forced errors occur only rarely, and we expect Solvers to make only few, if any, other errors. Rewards for identifying unforced errors may be modest compared to jackpot payouts,

and we consider any such payments of these type incidental, rather than fundamental, to TrueBit’s secure incentive structure.

The jackpot payout effectively bounds the complexity of computation tasks that TrueBit can perform securely. The Verifier must, on average, receive a payment which fairly compensates her for the task at hand, which means that the jackpot payout should at least consist of fair compensation for the current task times the forced error rate. In this way, the jackpot cap bounds the Verifier’s expected per task compensation which, by assumption (II) in Section 2.2, restricts Verifiers’ available CPU cycles per task.

We fix a rate for forced errors among tasks. This fraction should not be so low so as to discourage Verifier participation through excessively infrequent rewards, but neither should it be so high so as to run the risk of Referees bias (see Section 5.1). We set forced errors to occur, on average, once every thousand tasks.

4.2 Taxes

Given that a jackpot *repository* must exist, we now describe the mechanism for funding it. We assume that a generous philanthropist deposits some initial funds into the repository, but thereafter TrueBit will be self-sustaining through *taxes*. Any Task Giver that calls a TrueBit contract must pay not only the cost of computational work done by the Solver but also for the work done by the Verifier(s) (excluding unforced errors and bogus challenges), as well as the work done by Referees and Judges. We refer to the latter two costs as the *verification tax*. We ensure that the jackpot repository never disappears entirely by placing a cap on the jackpot size. To this end, we set the maximum jackpot payout for a forced error to be one third of the total repository size.

While a single, attentive Verifier suffices to ensure correctness and achieves ideal tax efficiency, in practice the verification tax requires a substantial cushion. We estimate the necessary verification tax to be 500% – 5000% of the cost of performing the given task. As we shall see in Section 5.2, there is a quantitative tradeoff between tax efficiency, Solver deposits, and security of computations, so Solvers could potentially absorb some of this tax burden by contributing higher deposits. Our tax rate estimate incorporates Verifiers’ incentives for participation and the fact that both the Task Giver and the Solver for a given task each have incentive to perform verification. Participation from these parties may necessitate higher taxes because the total jackpot payoff decreases exponentially as the number of challenges increases (see Section 5.3). Expected jackpot payoffs must be sufficiently high to con-

sistently attract at least one Verifier per task. The required tax amounts may also depend on peculiarities of human behavior. Indeed, we may have to pay Verifiers more than Solvers per CPU cycle because Verifier rewards have the human-undesirable property of being sporadic whereas Solvers always receive immediate rewards. Thus the optimal tax rate must, at least in part, be determined experimentally.

4.3 Deposits

TrueBit requires deposits from Solvers and Verifiers in order to thwart Sybil attacks (see Section 5.1) and incentivize correct computations. We set these deposits to be more than the expected jackpot payout for a given task plus the cost of playing a verification game. In particular, the deposits must be large enough to:

1. pay for the (expensive) cost of a verification game, including all rewards and penalties for Solver and Challengers and work done by Judges,
2. discourage Solvers and Verifiers from sacrificing deposits in order to obtain jackpots without performing verification (see Section 5.1 and Section 5.3),
3. discourage Task Givers who collude with Solvers in effort to get bogus solutions onto the blockchain (see Section 5.2),
4. refund taxes to the Task Giver in case the Solver causes an unforced error,
5. deter Solvers from randomly guessing solutions to obtain task rewards instead actually performing computations (as might be profitable for binary decision tasks with very high task rewards), and
6. deter external, temporal pathologies.

Note that the currency used to pay Solvers and Verifiers need not be the same as the currency used to pay Judges and Referees, but for simplicity of presentation, we fix ether (ETH) as the unique underlying currency.

As an example of the second type, consider a situation where the Solver deposit is small (say 10 ETH) but the expected jackpot payout per task is high (say 1000 ETH). An individual playing both the role of the Solver and Verifier could offer a bogus solution and then challenge his own answer, hypothetically netting, on average, $1000 - 10 = 990$ ETH without providing

any useful service. Such an action would degrade other Verifiers' incentive to participate.

As an example of the last case, if the outcome of a time-sensitive TrueBit contract controlled a 1000 ETH payout but only required a 10 ETH Solver deposit, then the party who stands to lose 1000 ETH from the TrueBit contract could attempt to cause a delay by posing as a Solver to the TrueBit contract and giving a bogus answer. We leave it to the Task Giver to determine appropriate minimum deposit values for such specific situations, as such contextual determinations lie outside of the scope of TrueBit itself.

4.4 Generating forced errors

In order to motivate verification of all tasks and to guard the jackpot repository against swindle, forced errors must appear unpredictably. TrueBit uses strings of random bits to determine whether or not a forced error occurs for a given task. The system derives its unpredictability via the following properties.

1. The Task Giver does not know the random bits at the time that she announces a task.
2. The Solver does not know the random bits until after he has committed his solution.
3. Verifiers do not know the random bits until after they decide whether to challenge.

The first property makes it difficult for a Task Giver to create a task designed to swindle the jackpot repository, and the second discourages Solvers from being lazy and only volunteering to solve tasks which have forced errors. In order to satisfy these three properties, TrueBit combines random bits from the following two sources:

- (a) "private" random bits from the Solver, and
- (b) the hash of the block mined immediately after the block containing the Solver's solution.

By *hash of the block*, or *block hash*, we more precisely mean the hash of the block's *header*, a roughly 200-byte piece of data that contains the timestamp, nonce, previous block hash, and Merkle root hash of the transactions occurring in the block [25]. Source (b) achieves properties 1. and 2. above, and source (a) achieves property 3.

TrueBit generates Solvers’ “private” random bits using a method reminiscent of RANDAO’s random generator mechanism [19]. Before participating in a TrueBit contract, a Solver must privately generate a string of random bits r and publish its hash on the blockchain. This action commits the Solver to using r in the protocol without revealing the value r to others.

The Solver establishes property 2. above by committing both a hash of a “correct” solution and a hash of an “incorrect” solution prior to the broadcast of the block hash from item (b). At the time that this block hash is revealed, only the Solver, who has access to the value r , knows whether or not a forced error is in effect. He publicly designates one of his two hashed solutions for evaluation, however potential Verifiers do not know *a priori* whether the Solver intended his solution to be “correct” or “incorrect.” Only after the timeout for challenges has elapsed do the Verifiers witness the Solver’s private random bits in the clear and learn whether a forced error was in effect. In case the Solver’s solution is challenged, the Solver must reveal r as well as his second “correct” solution in order to prove that the error was forced and avoid a penalty (see Section 5.4). Thus, in the case of a forced error, the Task Giver still obtains a correct solution. In case of an unforced error, i.e. when the Solver evidently fails to provide a correct solution, the Task Giver receives a full refund of her task reward and taxes (see Section 4.3 and Section 4.6).

Although in theory one could securely run this part of the protocol without hashing the Solver’s solutions, hashing makes it easier for the Solver to provide a convincing “incorrect” solution which appears, upon casual inspection, indistinguishable from a “correct” one. The Solver can effectively use any randomly selected “incorrect” solution during a forced error because he never has to reveal its preimage.

4.5 Solver and Verifier election

The system effectively chooses Solvers by lottery. When a task is announced, Solvers broadcast their interest in solving it to the Referees in the form of Ethereum transactions. Referees, or more specifically miners, choose one of these transactions to include the next block, thereby electing the Solver for that given task. In case the miner who chooses the transactions for the current block happens to also hold a lottery ticket, he may bias his chances of winning the lottery. This bias does not affect the security of TrueBit, however, since the Solver must still provide a correct solution.

In contrast to the Solver selection, TrueBit does not impose a limit on the number of Verifiers, and we describe in Section 5.3 how multiple Verifiers

who participate in a task must split their rewards after a successful challenge. Due to the computational work involved in verification, rational Verifiers will only verify tasks in case they expect to receive adequate compensation for their work (assumption (II), Section 2.2). Thus the number of Verifiers verifying each task will remain low due to the balance between computation costs and incentives (see Section 5.2 and Section 5.3).

4.6 Protocol overview

In this section we present an overview of the TrueBit protocol. We will discuss how our chosen parameters enable system security in Section 5. Throughout the description below, Solvers and Verifiers who deviate from protocol by broadcasting malformed data or failing to respond within time-out bounds forfeit their security deposits to the jackpot repository.

Preprocessing steps. The following must be done prior to commencing TrueBit operation:

1. A substantial jackpot repository must be established for the TrueBit contract prior to commencement of any interactions (See Section 4.1 for details).
2. Solvers and Verifiers that wish to participate must commit deposits to the TrueBit smart contract (See Section 4.3). The deposits may be placed at any time during the protocol so long as the underlying contract has sufficient time to confirm receipt of the funds.
3. Solvers must also generate private random bits and commit their respective hashes to the blockchain. We denote the private random bits of the unique Solver selected to perform the `task` below by r (see Section 4.4 for more details).
4. A universal tax rate T must be established (see Section 4.2).

Main algorithm. The protocol steps are as follows.

1. A Task Giver provides the following:
 - (a) a computational `task`,
 - (b) the `timeOut` for accepting bids, performing the computation, and waiting for a challenge. In the protocol below, we do not distinguish between these various timeouts with distinct notation, and we colloquially use “`timeOut`” to refer to both

events and lengths of time. In all cases, `timeOut` must be long enough to avoid microforks during which Referees temporarily disagree on current parameters.

- (c) a `reward` for a correct output, which must be at least the cash equivalent of the task difficulty d based on `timeOut` (see Section 5.5), plus a total tax of $T \cdot d$. The `reward` is held in escrow by the TrueBit contract while the taxes are immediately deposited into the jackpot repository.
 - (d) the minimum deposit, `minDeposit`, needed to participate as a Solver or Verifier (see Section 4.3, Section 5.1, Section 5.2, and Section 5.3).
2. Solvers who have the requisite `minDeposit` and random bits can bid to take on the task until the bidding `timeOut`. At most one Solver is selected (Section 4.5), and if no Solver takes on the task in the allotted time, the task request is canceled and the Task Giver receives a full refund.
 3. The Solver privately computes `task`. In case of a `timeOut`, Solver forfeits his deposit to the jackpot repository and the protocol terminates.
 - (a) Solver commits two distinct hashes to the blockchain, thereby committing both a “correct” and an “incorrect” solution.
 - (b) The hash of the next mined block is revealed, and then Solver knows whether or not there is a forced error (see Section 4.4).
 - (c) Solver designates one of the two hashes as the hash of his `solution`.
 4. Verifiers who have posted `minDeposit` can challenge (the hash of) `solution` until `timeOut`. Prior to `timeOut`, the Verifier must broadcast the hash of an even integer to the blockchain in order to commit to a challenge. Hashing an odd number in case of no challenge is optional and may be used to camouflage real challenges from other Verifiers (see Section 5.3). After `timeOut`, the Verifier broadcasts to the blockchain this hashed number in the clear to reveal her action.
 - (a) If no Verifier challenges `solution`, then
 - i. Solver reveals r to prove that there was no forced error (i.e. the criteria in Step 4(b)i below fails),
 - ii. Solver reveals `solution` in the clear on the blockchain,
 - iii. Solver receives the task `reward`, and then

- iv. the protocol terminates.
- (b) Otherwise, some Verifier challenges `solution`.
 - i. Solver reveals his private random string r on the blockchain, and Referees check it against his commitment from preprocessing step 3. If the hash of the concatenation of r and the block hash following the `solution` announcement from 3(b) is small (as determined by the forced error rate, see Section 4.1), then a forced error is in effect (see Section 4.4).
 - ii. If the value r reveals that Solver's error was forced, then Solver must reveal his secondary solution in the clear (see Section 5.4).
 - A. If no Verifier challenges Solver's secondary solution before `timeOut`, then Verifier wins a fraction the maximum jackpot amount J , scaled for task difficulty. In case of challenges from multiple Verifiers, the jackpot is split among them. In more detail, if there are k distinct challenges, then each participating Verifier receives $J/2^{k-1}$ times the ratio of the task's difficulty to the system's maximum task difficulty (See Sections 4.1, 5.3, and 5.4 for further discussion).
 - B. Otherwise the Solver must play the verification game with the challenging Verifier(s). Verifier penalties, Verifier rewards, Solver penalties, and refunds to the Task Giver here are the same as described in the next step. In case the Solver successfully defends himself against all challenges, however, then the jackpot payouts proceed as in Step 4(b)i above.
 - iii. Otherwise the error was not forced. Solver reveals `solution` in the clear, and then Solver and Verifier must play a verification game (Section 3.2). In case of challenges from multiple Verifiers, the steps below are repeated until either some Verifier wins a challenge or Solver defeats all Verifier challenges. The Solver collects `reward` only once he wins all of the challenges, and if this does not happen, then the Task Giver receives a refund on his `reward`, and tax payments are reimbursed through Solver's deposit as described below.
 - A. If Solver wins, then Verifier forfeits half of her deposit to

the jackpot repository and the other half to the Solver, where the latter suffices to compensate the Solver for playing the verification game (see Section 4.3).

- B. Otherwise Verifier wins. Solver pays at most half of his deposit to the Verifier(s), according to the distribution scheme in Step 4(b)ii above, pays back the Task Giver’s taxes out of this amount, and forfeits the remaining funds to the jackpot repository (see Section 5.1 and Section 5.2).

End of protocol.

Note that Solver does not receive `reward` when a forced error occurs in Step 4(b)ii. In such situations, we expect the Solver to challenge his own “mistake” and receive a jackpot payout much greater than `reward`. TrueBit makes this payout automatic since the Solver would lose the verification game anyway if he and Verifier were to play it out.

4.7 Sanity check

The table on the next page recaps the parameters used in the TrueBit protocol and hints at how to estimate their exact values. Each parameter is either a fixed constant, an input from a participant or something already on the blockchain, or can be computed from the elements which appear above it in the table. By inspecting the list below from top to bottom, one can confirm that all notions are well-defined.

5 Defenses

We now analyze the security of the incentive layer from Section 4. TrueBit’s security relies on the presence of at least one Verifier to check each task performed. In this section we show that available incentives suffice to guarantee the existence of such Verifier(s) according to the network assumptions in Section 2.2 and the attacker model from Section 2.3. TrueBit defends against shortcuts which could divert verification incentives to non-verifying parties. In particular, we show that TrueBit resists Sybil attacks, collusion pools, opportunistic attacks related to jackpot payoffs, and certain external threats.

parameter	dependency
dispute layer parameters p , c , and σ (Section 3)	- fixed constants.
tax rate (Section 4.2, Section 5.2)	- fixed constant (500% – 5000%).
forced error rate (Section 4.1)	- fixed constant (1/1000).
maximum jackpot payout (Section 4.2)	- 1/3 of current jackpot repository.
cash equivalent of CPU cycles (Section 5.5)	- based on external markets.
maximum task difficulty (Section 4.1)	- maximum jackpot payout divided by cash equivalent of a CPU cycle.
task	Parameters in this box are chosen by the Task Giver with minimums as described below:
timeouts (Section 4.6)	- long enough to avoid microforks.
task difficulty	- determined by timeouts.
Solver reward (Section 5.5)	- no less than the cash equivalent of the task difficulty.
expected jackpot payout (Section 4.1, Section 5.2, Section 5.4)	- cash equivalent of task difficulty, and number of active Verifier deposits.
Solver & Verifier deposits (Section 4.3, Section 5.1, Section 5.2, Section 5.3)	- more than the cost of verification game plus the expected jackpot payout. Also depends on tax rate.
actual number of Verifiers (Section 4.5)	- as many join (incentives limit over-participation).
jackpot payout for challenging a forced error (Section 5.3)	- based on maximum jackpot payout, actual number of verifiers, and ratio of task difficulty to maximum task difficulty.
Payout for detecting an unforced error (Section 5.1, Section 5.2)	- at most half of Solver deposit is split among all Verifiers (rest goes to jackpot repository).

Table: Relations between parameters used in TrueBit.

5.1 Pairwise Sybil attacks

In a Sybil attack, an adversary assumes multiple identities on the network in order to execute an exploit. Identities on TrueBit include Task Givers, Solvers, Verifiers, Judges, and Referees. By our assumption (Section 2.3), Judges and Referees always function as intended, and we provide additional justification for this axiom here. In this subsection, we consider all sets of pairs among the remaining identity types and show that pairwise cooperation does not harm operation of the system. While parties can freely join and leave TrueBit, each identity must make a deposit in order to participate. This deposit alone is a general deterrent against Sybil attacks, however as we shall see, multiple identities do not provide much cheating advantage.

Judges and Referees. Recall that Ethereum miners play the roles of Judges and Referees in TrueBit. Our analyses in Section 3.4 and Section 4.6 show that the work done by these parties is small, and hence they are not vulnerable to a Verifier’s Dilemma (Section 3.1). Nakamoto consensus [47] therefore ensures that miners will not post bogus transactions, i.e. enforce rules incorrectly, lest other miners reject their block and cost them a block reward. Therefore the only threat from miners is block withholding, specifically with regard to random number generator bias (see Section 4.4).

While in theory miners could discard blocks, along with their mining reward and associated random bits for TrueBit, in practice this may never happen. A miner who drops a block must expect in exchange some income greater than the usual mining reward in order to make this action worthwhile. If such income were to come as a bribe from a TrueBit participant, it would have to come from a Solver since only the Solver, who has unique access to his own private random bits, knows how to bias miners’ random bits towards or away from a forced error. In short, miners cannot disturb randomness in TrueBit without a tip-off from a Solver. The Solver, however, has no reason to bias against a forced error because this would prevent him from challenging his own answer and winning the jackpot, or at the very least would invite others to share it with him, thereby decreasing the total jackpot payout (see Section 5.3). Moreover, the Solver is unlikely to succeed in biasing towards a forced error since miners have little control over their own block hashes. This “one-sided” effect of block withholding, which can lock away jackpot funds but (almost) never release them, makes block hashes a safe source of randomness in TrueBit. Hence the Solver’s potential reward for challenging his own solution under a forced error is not merely an artifact of TrueBit’s incentive structure — it guarantees unbiased Referees.

Task Giver and Solver. A Task Giver who also acts as a Solver does not benefit from solving her own task. One idea for an attack vector would be to create a task such that the Solver’s private random bits force an error for the given task. Since the Task Giver cannot predict the random bits from the block hash at the time the task is created (Section 4.4, part (b)), the Task Giver cannot create such a task. Moreover, the Task Giver cannot afford to flood TrueBit with useless tasks and solve them herself in the hopes of eventually, by chance, encountering a forced error. The Task Giver would pay more taxes to the jackpot repository than she would expect to win from the jackpot payout (Section 4.2), even taking into consideration any rewards she might win by correctly solving her own tasks. Her losses are amplified through payments to other Solvers, Verifiers, and Judges who choose to participate in the protocol.

Solver and Verifier. The Solver’s burned deposit always exceeds the Verifier’s income for successfully challenging an unforced error (see Section 4.3). Hence the Solver has no incentive to challenge himself as a Verifier in such cases. Similarly, a Verifier can only lose her deposit by posing bogus challenges. In certain situations it is conceivable that a Solver–Verifier pair could benefit from submitting a false solution and then challenging it due to temporal constraints external to TrueBit itself (as mentioned in Section 4.3), and in such cases the Task Giver must determine the necessary deposits to deter such actions.

In the case of an forced error, we expect the Solver will challenge himself in order to win the jackpot. Nevertheless, TrueBit’s tax rate (Section 4.2), and hence its jackpot payout (Section 4.1), suffices to incentivize an independent Verifier to also check the solution. As we shall see in Section 5.3, the Solver lacks incentive to claim the same jackpot more than once, hence the Solver’s self-verification does not spoil other Verifiers’ motivation to participate.

Task Giver and Verifier. A Task giver can certainly verify the solution she receives. This checking can only improve the reliability of the system! We cannot, however, assume that the Task Giver will always perform such checks due to the Task Giver’s possible computational and financial resource constraints or lack of motivation to involve herself in the details of a TrueBit contract.

5.2 The trifecta

Up until this point, we have implicitly assumed that the attacker’s goal was to extract money from the jackpot without performing verification. However there is another possibility: the attacker wishes to get a bogus solution accepted onto the blockchain. No penalty scheme can deter all such attacks since TrueBit, as a closed system, has no way to estimate the true economic impact of a bogus solution on Ethereum’s blockchain. Hence we now consider scenarios in which a party is willing to sacrifice a jackpot payout (or more) in order to skew computational results.

Scaring off Verifiers. Suppose that a Task Giver wishes to obtain an incorrect solution for a task. Posing as a Solver, she gives a wrong solution to her own task. Now the only thing that could stand in the way of her success would be a pesky Verifier. In order to dodge this obstacle, the attacker poses as a regular Verifier prior to posting her malicious task. She checks every task that comes along, until eventually he encounters a forced error. At this point, the attacker challenges not once but a hundred times, which according to Step 4(b)ii in Section 4.6 essentially wipes out the jackpot payout for herself and any other Verifiers that also challenged this solution. The upshot is that legitimate Verifiers no longer wish to participate because TrueBit failed to deliver their well-deserved jackpot payout, and more to the point, they have no way to estimate their income from future jackpots. So when the attacker finally broadcasts her malicious task, legitimate Verifiers simply ignore it.

As Verifier participation, and hence security, depends entirely on the expected value of jackpot payouts (assumption (ii) in Section 2.2), TrueBit must ensure sufficient expected value for these payouts. We therefore offer the following default strategy for all Verifiers. New tasks are submitted to TrueBit at various and overlapping times. A prospective Verifier who wishes to earn rewards for verification should count the number of active Verifier deposits and compare this to the average number of simultaneous tasks occurring on the network. If the tax rate and task flow profitably support an additional Verifier, then the prospective Verifier submits a deposit, randomly chooses an active task, and begins verifying. As soon as she finishes verifying this task, she randomly selects another active task and begins verifying that next one. Thus a Verifier with either more powerful hardware or better algorithms can gain a market advantage in TrueBit. Such a Verifier can check more solutions per unit time and therefore increase her expected jackpot rewards per unit time. Verifiers are free to leave the network at

any time, for any reason. If the current number of active Verifiers exceeds the number supported by the tax rate and task flow, then the probabilistic default strategy tells each Verifier to exit the system with the minimal probability that one would expect to restore equilibrium if all Verifiers were following the default strategy.

Verifiers who choose tasks to verify at random according to the default strategy, as opposed to changing their behavior based on guesses as to who solved-what before, can reasonably estimate their compensation for participation. When a significant number of Verifiers (not necessarily a majority) follow the default strategy, the adversary wishing to establish a bogus solution can neither predict which tasks will be verified, nor can he influence which tasks default strategy Verifiers will verify. According to assumption (ii) in Section 2.2, there exists some bound on the adversary’s financial means. If not, he could simply flood the Verifier pool with deposits and begin verifying, which in turn would incentivize legitimate Verifiers to leave the system. Therefore the adversary can only post a bogus solution and hope that, by chance, it gets ignored. In this way, the default Verifier strategy spoils the scare-off strategy above.

Let us now estimate the chances that the adversary succeeds to post a bogus solution which all legitimate Verifiers fail to notice. We make the worst-case assumption that the adversary Solver always gets elected to solve his own task. Suppose that the tax rate (see Section 4.2) supports at least 6 Verifiers per task, and that the adversary holds at most 1/6 of all Verifier deposits. Then effectively there are an average of 5 verifiers per task. Suppose n is the average number of simultaneous tasks on the network. Then the number of Verifiers is approximately $5n$, and so the chance that none of them check a particular solution is the chance that each of them independently and randomly choose other tasks, or $[(n-1)/n]^{5n}$. Therefore the probability that the adversary gets caught, regardless of n , is

$$1 - \left(1 - \frac{1}{n}\right)^{5n} > 1 - e^{-5} > 99\%.$$

The leftmost inequality follows from the standard limit definition for the Calculus constant e . We conclude that the adversary will most likely lose his Solver deposit by attempting such an attack.

By raising taxes further, we can make such an attack even more costly. In this sense the tax rate, together with the minimums for Solver deposits, bounds the mandatory “honesty” of the network. If bogus solutions never occur, then there would be no need for taxes. On the other hand, higher taxes make it more expensive to cheat. There is a tradeoff between overhead

tax expenses for Task Givers, minimum deposits for Solvers, and security of computations.

High-stake tasks. For some high-stakes applications it is possible to entirely fall back on the dispute resolution layer of Section 3. In a high-stakes situation, it is reasonable to expect that other parties who have a significant financial stake in the outcome of a given task will make an effort to verify the solution regardless of expected jackpot payouts. This motivation acts as a secondary deterrent against the “Scaring off Verifiers” attack above, as does the fact that the Verifier receives a fraction of the Solver’s deposit in case of positive error detection (TrueBit burns part of the Solver’s deposit in order to avoid collusion between Solvers and Verifiers).

5.3 Collusion pools

In this section, we analyze the potential effects of pooling computational, informational, and financial resources among Verifiers and Solvers. Participants in TrueBit may voluntarily form a *Cartel* which may in turn use external smart contracts to enforce cooperation among its (mutually distrusting) members, however, Solver and Verifier deposits in TrueBit alone prevent many potentially harmful financial collusions designed to extract jackpot funds without exerting computational effort. Secondly, carefully chosen jackpot payout amounts, together with Solver and Verifier incentives and private commitment channels, prevent financial gain from unintended signaling of private random bits or challenges.

Rich but powerless. First, consider a Cartel of Solvers with limited CPU bandwidth but deep pockets which volunteers to solve each and every task that comes along. The Cartel makes no attempt to provide correct solutions to any of these tasks, but instead intends to absorb lost Solver deposits until a forced error comes along, at which point it splits the jackpot among its members. By construction (see Section 4.3), Solver deposits exceed the expected jackpot payout per task. Therefore, in the long run, such a Cartel strategy loses money. If members of the Cartel instead decide to start producing correct solutions, this does not harm security because that’s what they are supposed to do anyway. We remark that, also by construction (Section 4.4), a Solver cannot simply decide to accept only tasks which have forced errors because the Solver does not know whether or not a forced error is in effect until after he has committed his solution to the blockchain.

Similarly, a Cartel which challenges every single task in the hopes of eventually obtaining a jackpot sustains losses in the long run due to lost Verifier deposits.

A flood of Challengers. Timely information about forced errors has economic value. While Solvers can earn jackpots by challenging their own forced errors, they risk sharing jackpot payouts by divulging information about such errors. It follows that Solvers have incentive to keep their knowledge of forced errors private and their fake solutions convincing. Information about active challenges also has economic value because a challenge posted to the blockchain could be a tip-off about a forced error. For this reason, Verifiers have incentive to announce “fake” challenges rather than remaining silent on tasks in which they don’t detect errors (See Step 5. in Section 4.6). “Fake” challenges serve as noise to mask real challenges and protect Verifiers’ jackpot revenue from potential copycats who might challenge just because they see someone else doing it.

If the jackpot for a given task were a fixed amount and equally divided among challenging participants, then one could potentially flood the set of participants with aliases to obtain more than one’s share of the jackpot. A Cartel in which members inform each other about forced errors could therefore have the economic means to flood the network with challenges and monopolize jackpot payouts. This in turn would dilute jackpots for legitimate Verifiers which would, in turn, degrade Verifier incentives for participation. While such a Cartel might successfully detect forced errors, it might also cause tasks with unforced errors to go unverified.

In order to dissuade a Cartel from flooding the network with harmful challenges, we must reduce the total jackpot payout based on the number of challenges that occur. The total jackpot paid out on a forced error with k challengers must not exceed

$$\frac{J}{2^{k-1}}. \tag{5.1}$$

The reason is as follows. We want a Challenger’s net reward to decrease if she superfluously clones herself via the Cartel. Assume that there are n legitimate challenges plus a Cartel which initially contributes k challenges to the pool and is considering adding another alias. We want the per challenge jackpot distribution among $n + k$ challenges to be less than it would be among $n + k + 1$ challenges regardless of k , and n . Let J_i denote the total reward paid out when there are exactly i challenges. Then we want

$$\frac{k}{n} \cdot J_{n+k} > \frac{k+1}{n+1} \cdot J_{n+k+1},$$

or

$$J_{n+k+1} < J_{n+k} \cdot \underbrace{\frac{n+1}{n}}_{>1 \text{ for all } n \geq 0} \cdot \underbrace{\frac{k}{k+1}}_{\geq 1/2 \text{ for all } k \geq 1}.$$

Thus it suffices to set $J_{n+k+1} = J_{n+k}/2$, or by induction $J_k \leq J_1/2^{k-1}$. In fact we cannot do better than this because the case $k = 1$ asymptotically forces a matching bound.

The upper bound in (5.1) does not take into consideration the Verifier’s cost of playing the verification game. In the case where the Solver’s solution has been found to have an error through one of the challenges, there is no need to repeat the verification game with the other challenges. The Verifier who actually plays the verification game with the Solver receives a bonus from the Solver’s deposit to compensate her for the CPU cycles spent during the challenge.

5.4 On low-hanging fruit

We analyze the security implications of Verifiers and Solvers who opportunistically disappear for hard tasks and reappear for easy ones, or who take advantage of solution discrepancies.

Easy winners. If the jackpot for all tasks were equal, rational Verifiers might choose to verify only simple tasks and ignore complex ones. For this reason, the jackpot for each task scales proportionally with the task’s complexity (Step 5. in Section 4.6). Scaling ensures that a Verifier’s expected jackpot payout per CPU cycle remains constant across all tasks, and it equally incentivizes Verifiers to inspect simple and complex tasks. Solvers always appear to perform tasks since the minimum task reward suffices to compensate them for their work (see Section 4.6 and assumption (II) in Section 2.2,).

Multiple solvers. For security reasons, TrueBit explicitly does not allow Task Givers to hire redundant Solvers for a single task (Main Algorithm, Step 2., Section 4.6). Suppose that two Solvers provided solutions to a single task and exactly one of them receives a forced error. Any Observer who notices a difference in the two solutions on the blockchain could, without verifying anything, challenge both solutions. By playing two verifications and sacrificing one deposit, such an Observer could potentially win a jackpot at negligible cost, thereby degrading the overall incentives for verification.

Forced error in disguise. What happens if a Solver has a forced error but gives a correct solution anyway? The Solver could then challenge his own solution while other Verifiers ignore it, resulting in both the Solver receiving a bigger share of the jackpot as well as damage to the system’s verification mechanism. Therefore, when the Solver reveals that a forced error was in effect, the protocol also forces him to reveal his committed “correct” solution which must be distinct from the “incorrect” solution that he showed at first. If the second solution revealed by the Solver is correct and distinct from his first solution, then by uniqueness of deterministic processes, the first solution must have been incorrect (as desired). Verifiers have an opportunity to challenge this second solution. If an error is detected in it via a verification game, the Verifier(s) split the Solver’s deposit rather than the jackpot payout according to Step 4(b)ii in Section 4.6. Since the Verifier(s) presumably already performed the task themselves when checking the first solution, no additional CPU cycles are required to do this second check. As the Solver loses both a jackpot payout and a deposit by following this course, by assumption (II) in Section 2.3, the situation described in this paragraph should never arise.

5.5 A cash equivalence problem

Consider the following scenario. A Task Giver wishes to get a bogus solution onto the blockchain. He offers a minimal reward for a difficult task so as to ensure that no legitimate Solvers or Verifiers volunteer to inspect it. Acting as a Solver, she then provides a bogus solution, and the system accepts her solution because no one bothers to check it. It follows that for security purposes, the system must require that Task Givers compensate fairly based on the difficulty of the task. But how can TrueBit estimate the true cost of executing a given task in its native currency?

While one could potentially establish a long-term lower bound on the cost of a CPU cycle relative to a stable currency like the US dollar or Euro, calculating a lower bound relative to the value of a cryptocurrency token is another matter. Cryptocurrencies are extremely volatile. Moreover, TrueBit lives on the blockchain and does not have access to a newspaper with current exchange rates.

In the first iteration of TrueBit, we will manually update the internal cash equivalent of a CPU cycle based on a live feed (e.g. [3, 17, 20]). Ultimately, however, we would like to input these prices in a decentralized way without relying on a third-party. Later versions of the protocol may make use of Augur [1], a decentralized prediction market which ports outside information

sources onto the Ethereum blockchain. As of this writing, Augur is currently in beta testing. Alternatively, we may build an independent blockchain for reporting prices whose “transactions” consist of exchange rate updates.

6 Implementation

Formally, TrueBit is a smart contract in Ethereum which uses Ethereum’s existing smart contract framework to bootstrap the creation of computationally-intensive TrueBit contracts. Tasks in TrueBit take the form of C, C++, or Rust code, but the user must pass this code as input to the Google Lanai architecture [14] (discussed below) prior to submission to TrueBit. This latter step guarantees consistency of simulated architecture and allows Judges to adjudicate fairly.

Google’s Lanai interpreter. The theoretical framework for the verification game requires a fixed computer architecture to be used for all verification tasks. In [33], the authors used the regular Intel X86 architecture for this purpose, but we believe that this architecture is far too complicated to be used in TrueBit. All participants of the verification game, including the Judges, have to simulate the behavior of the entire architecture. Even a slight difference in these implementations could result in catastrophic loss of funds for one of the parties. Moreover, simpler architecture costs less to simulate on the blockchain, and because of its simplicity and the fact that there now exists an efficient and trusted LLVM compiler writer, we have chosen to use Google’s Lanai architecture [14]. The full Google Lanai interpreter will be available as a smart contract in Ethereum so that it can be used by Judges. In all TrueBit contracts, the Lanai bytecode will be authoritative.

TrueBit’s on-chain interpreter runs in Solidity. For efficiency reasons, tasks will not be run using the interpreters except in case of dispute. In general, users will run tasks written in native programming languages and running on regular hardware. In the majority of cases where tasks do not require the Lanai interpreter, Solvers and Verifiers can privately optimize implementation of task executions and potentially gain a market advantage over other participants.

7 Applications

TrueBit is more than just an outsourced computation system. It is designed for use in trustless smart contracts. We present some examples of possible

use cases.

7.1 Practical decentralized pooled mining

Mining rewards are extremely competitive. A typical desktop computer might only mine on average one Bitcoin block every thousand years. To reduce income variance, miners often join *mining pools* which share both computing resources and mining rewards. Each pool has an *operator* who distributes computing tasks and rewards to pool *members*. This centralization degrades the security of the system by giving the operator undue influence to censor transactions [43]. In extreme cases where an operator controls more than half of the network’s hash rate, as has been the case with DwarfPool [8] in Ethereum, GHash.io [24] in Bitcoin, and could happen again with Bitmain’s proposed gigantic mining center [2], the operator can even withdraw cleared transactions and double-spend money by way of a *51% attack* [40].

SmartPool [21] introduces mining pools whose operators are Ethereum smart contracts. As decentralized pool operators, smart contracts have many virtues. In addition to counteracting the censorship issues described in the previous paragraph, they can operate at low cost relative to centralized pools and do not rely on a social contract to ensure fairness. Unlike other decentralized mining pools, which either have higher payout variance [18] that negate incentives for joining the pool or require a change in the proof-of-work protocol [45], SmartPool offers low variance payouts and retrofits existing cryptocurrencies, and it can handle a huge number of participants with a wide range of computing power.

SmartPool’s proposed Ethereum mining pool minimizes verification work by precomputing the 1 GB data sets needed to check Ethereum’s proof-of-work [9, 43]. While this shortcut may help in checking Ethereum’s proof-of-work, not all cryptocurrencies have this lucky feature. Fortunately TrueBit contracts can check *any* proof-of-work, which means that with TrueBit, we can build a smart contract-based mining pool for *any* Nakamoto consensus-based cryptocurrency. At the time of this writing, for example, the effort required to check a Zcash proof-of-work [26] appears to exceed Ethereum’s `gasLimit` capacity by a factor of 50. TrueBit is an option for bringing this task within reach of Ethereum smart contracts. Project Alchemy [16], which aims to bring smart contract functionality to Zcash in the spirit of the Dogecoin–Ethereum bridge below, may also benefit from TrueBit’s ability to check Zcash’s proof-of-work.

7.2 Dogecoin–Ethereum bridge

We can use TrueBit to build a two-way peg between Dogecoin [5] and Ethereum, enabling Dogecoin users to move dogecoins between Dogecoin’s blockchain and Ethereum’s blockchain without transmitting currency through a third-party exchange and effectively adding smart contract functionality to Dogecoin. The Dogecoin community maintains an active interest in such a bridge [7], and current offers a prize of more than 6000 ETH for its construction [4, 6].

TrueBit contracts can check Dogecoin’s Scrypt-based proof-of-work whereas traditional Ethereum smart contracts cannot. If Dogecoin were to enhance its scripting language with an operation indicating currency transfer to Ethereum addresses, a TrueBit contract could then confirm such transfers. The newly created dogecoin token could then be passed around Ethereum, and a final signed transaction in Ethereum could finally send the dogecoin back onto Dogecoin’s blockchain, assuming that Dogecoin miners are willing to treat the Ethereum blockchain as authoritative for such transfers.

7.3 Scalable transaction throughput

Building a system that can securely meet even the modest transaction volume demands of current Bitcoin users remains a formidable challenge [34]. Proposed Byzantine “sharding” [29, 35, 39, 41, 44, 46, 48], and miner-based “serializing” [36, 51] protocols exist which aim to distribute verification work, but here we take a simple and fundamentally different approach to decouple the two tasks which miners perform, namely

1. selecting which transactions to include in the blockchain, and
2. verifying that blockchain transactions are valid.

Using TrueBit, one can construct a verification protocol whose incentives guarantee that task 2 is correctly performed by off-chain Solvers and Verifiers (with some help from Judges and Referees), while miners continue to perform task 1. In this way, complex transactions can securely reach the blockchain without overburdening miners.

7.4 Towards a big data system

In order to perform as a truly scalable cloud computer, TrueBit must have access to a scalable data storage system. Ethereum’s blockchain alone does not suffice as storing even moderate amounts of data directly on Ethereum’s

blockchain is prohibitively expensive. TrueBit can securely access and use portions of massive data sets so long as the data is stored somewhere publicly and permanently, for example in Swarm [22] or on another blockchain. Parties who wish to rely on such data in a TrueBit contract must be sure that Verifiers have access to the full data set.

To use TrueBit on external data, one need only store a Merkle root of the massive data set on the blockchain and add non-deterministic steps in the verification game in which the Solver can “guess” the original data set represented by the Merkle root. While Solvers and Verifiers must have access to the full data, Judges and Referees do not. Indeed, if we modify the verification game so as to permit tasks for nondeterministic Turing machines, then the Solver can nondeterministically guess the certificate data as a step in the TrueBit contract. Only in cases of disputes would the Solver have to reveal external certificate data to the Judges via the blockchain. In some applications, the Solver might even be able to reveal to the Judges a privacy-preserving zkSNARK rather than the data itself. zkSNARKs have the potential to enable privacy for many different kinds of systems on Ethereum [50].

While in theory TrueBit’s scalable protocol can process arbitrarily complex tasks, in practice the verification game is inefficient and therefore security of TrueBit computations degrades as tasks reach the level of big data. For big data applications, TrueBit may not be able to rely on a one-size-fits-all verification game. Therefore we anticipate optimizing the verification game for certain classes of tasks. For example, Section 3.2 gives an example of an efficient, specialized verification game for matrix multiplication. In future versions of TrueBit, Task Givers might broadcast not only tasks but also an indication of the corresponding verification game to be played in case of disputes.

Remark. We conclude with a caveat: TrueBit may expose latent security vulnerabilities in the underlying Ethereum network as a result of new kinds of interactions between smart contracts and miners. By allowing smart contracts to check proof-of-works, for example, TrueBit may facilitate 38.2% attacks [53].

Acknowledgments. We thank Vitalik Buterin and Vlad Zamfir for suggesting the use of forced errors in the TrueBit protocol and Eli Bendersky for introducing us to Google Lanai. We also thank Loi Luu and Julia Koch for useful discussions.

References

- [1] Augur. <https://www.augur.net/>.
- [2] Bitmain responds to controversy surrounding its upcoming 140,000 kw mining center. <http://www.newsbtc.com/2016/11/04/bitmain-response-new-mining-center/>.
- [3] Crypto-currency market capitalizations. <http://coinmarketcap.com/>.
- [4] The Doge connection Bounty Dao is live and working. https://www.reddit.com/r/ethereum/comments/41ohhr/the_doge_connection_bounty_dao_is_live_and_working/.
- [5] Dogecoin. <http://dogecoin.com/>.
- [6] Dogecoin–Ethereum bounty smart contract. <https://etherscan.io/address/0xdbf03b407c01e7cd3cbea99509d93f8dddc8c6fb>.
- [7] Dogethereum. <https://www.reddit.com/r/dogethereum/>.
- [8] Dwarfpool is now 50.5%. <http://forum.ethereum.org/discussion/5244/dwarfpool-is-now-50-5>.
- [9] Ethash: defining the seed hash. <https://github.com/ethereum/wiki/wiki/Ethash#defining-the-seed-hash>.
- [10] Ethereum. <http://ethereum.org/>.
- [11] Ethereum Computation Market. <http://www.ethereum-computation-market.com/>.
- [12] Golem. <http://github.com/imapp-pl/golem/wiki/FAQ>.
- [13] The Golem Project: crowdfunding whitepaper. <https://golem.network/doc/Golemwhitepaper.pdf>.
- [14] Google Lanai. <http://llvm.org/docs/CompilerWriterInfo.html#lanai>.
- [15] I thikn the attacker is this miner—today he made over \$50k. https://www.reddit.com/r/ethereum/comments/55xh2w/i_thikn_the_attacker_is_this_miner_today_he_made/.

- [16] Introducing Project Alchemy. <https://z.cash/blog/project-alchemy.html>.
- [17] Oraclize. <http://www.oraclize.it/>.
- [18] P2Pool. <http://p2pool.org/>.
- [19] RANDAO. <https://github.com/randao/randao>.
- [20] Reality Keys. <https://www.realitykeys.com/>.
- [21] Smartpool. <http://smartpool.io>.
- [22] Swarm. <http://swarm-gateways.net/>.
- [23] The DAO. <http://daohub.org/>.
- [24] Warning: Ghash.io is nearing 51% – leave the pool. <http://www.cryptocoinsnews.com/warning-ghash-io-nearing-51-leave-pool/>.
- [25] White paper. <https://github.com/ethereum/wiki/wiki/White-Paper>.
- [26] Why equihash? <https://z.cash/blog/why-equihash.html>.
- [27] Zcash. <https://z.cash/>.
- [28] Some miners generating invalid blocks. <https://bitcoin.org/en/alert/2015-07-04-spv-mining>, July 2015.
- [29] Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Alexander Spiegelman. Solidus: An incentive-compatible cryptocurrency based on permissionless Byzantine consensus. <https://arxiv.org/abs/1612.02916>, 2016.
- [30] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013: 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, pages 90–108, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

- [31] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. Indistinguishability obfuscation vs. auxiliary-input extractable functions: One must fall. <https://eprint.iacr.org/2013/468.pdf>.
- [32] Ran Canetti, Ben Riva, and Guy N. Rothblum. Practical delegation of computation using multiple servers. In *Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS '11*, pages 445–454, New York, NY, USA, 2011. ACM.
- [33] Ran Canetti, Ben Riva, and Guy N. Rothblum. Refereed delegation of computation. *Information and Computation*, 226:16 – 36, 2013. Special Issue: Information Security as a Resource.
- [34] Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Wattenhofer. On scaling decentralized blockchains (a position paper). In *Financial Cryptography and Data Security 2016 BITCOIN Workshop*, volume 9604 of *Lecture Notes in Computer Science*, pages 106–125. Springer Berlin Heidelberg, February 2016.
- [35] Christian Decker, Jochen Seidel, and Roger Wattenhofer. Bitcoin meets strong consistency. In *Proceedings of the 17th International Conference on Distributed Computing and Networking, ICDCN '16*, pages 13:1–13:10, New York, NY, USA, 2016. ACM.
- [36] Ittay Eyal, Adem Efe Gencer, Emin Gun Sirer, and Robbert Van Renesse. Bitcoin-NG: A scalable blockchain protocol. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pages 45–59, Santa Clara, CA, March 2016. USENIX Association.
- [37] Tim Goddard. AdversariallyVerifiableMachine. https://www.reddit.com/r/ethereum/comments/51qjz6/interactive_verification_of_c_programs/d7ey41n/.
- [38] Sanjay Jain, Prateek Saxena, Frank Stephan, and Jason Teutsch. How to verify computation with a rational network. <https://arxiv.org/abs/1606.05917>, June 2016.
- [39] Eleftherios Kokoris Kogias, Philipp Jovanovic, Nicolas Gailly, Ismail Khoffi, Linus Gasser, and Bryan Ford. Enhancing Bitcoin security and performance with strong consistency via collective signing. In *25th*

USENIX Security Symposium (USENIX Security 16), pages 279–296, Austin, TX, 2016. USENIX Association.

- [40] Joshua A. Kroll, Ian C. Davey, and Edward W. Felten. The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. <http://www.econinfosec.org/archive/weis2013/papers/KrollDaveyFeltenWEIS2013.pdf>, June 2013.
- [41] Loi Luu, Viswesh Narayanan, Chaodong Zheng, Kunal Baweja, Seth Gilbert, and Prateek Saxena. A secure sharding protocol for open blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 17–30, New York, NY, USA, 2016. ACM.
- [42] Loi Luu, Jason Teutsch, Raghav Kulkarni, and Prateek Saxena. Demystifying incentives in the consensus computer. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS 2015)*, pages 706–719, New York, NY, USA, 2015. ACM.
- [43] Loi Luu, Yaron Welner, Jason Teutsch, and Prateek Saxena. SmartPool: Practical decentralized pooled mining. <http://smartpool.io/docs/smartpool.pdf>.
- [44] Silvio Micali. ALGORAND: the efficient and democratic ledger. <http://arxiv.org/abs/1607.01341>, 2016.
- [45] Andrew Miller, Ahmed Kosba, Jonathan Katz, and Elaine Shi. Nonout-sourceable scratch-off puzzles to discourage Bitcoin mining coalitions. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, pages 680–691, New York, NY, USA, 2015. ACM.
- [46] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of BFT protocols. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 31–42, New York, NY, USA, 2016. ACM.
- [47] Satoshi Nakamoto. Bitcoin P2P e-cash paper. <http://www.mail-archive.com/cryptography@metzdowd.com/msg09959.html>, November 2008.
- [48] Rafael Pass and Elaine Shi. Hybrid consensus: Efficient consensus in the permissionless model. <https://eprint.iacr.org/2016/917.pdf>, 2016.

- [49] Christian Reitwiessner. From smart contracts to courts with not so smart judges. <https://blog.ethereum.org/2016/02/17/smart-contracts-courts-not-smart-judges/>.
- [50] Christian Reitwiessner. zkSNARKs in a nutshell. <https://blog.ethereum.org/2016/12/05/zksnarks-in-a-nutshell/>, December 2016.
- [51] Yonatan Sompolinsky, Yoad Lewenberg, and Aviv Zohar. SPECTRE: A fast and scalable cryptocurrency protocol. <https://eprint.iacr.org/2016/1159.pdf>, 2016.
- [52] Nick Szabo. The idea of smart contracts. http://szabo.best.vwh.net/smart_contracts_idea.html, 1997.
- [53] Jason Teutsch, Sanjay Jain, and Prateek Saxena. When cryptocurrencies mine their own business. (to appear in *Financial Cryptography and Data Security (FC 2016)*).
- [54] Jelle van den Hooff, M. Frans Kaashoek, and Nikolai Zeldovich. Ver-sum: Verifiable computations over large public logs. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, pages 1304–1316, New York, NY, USA, 2014. ACM.
- [55] Michael Walfish and Andrew J. Blumberg. Verifying computations without reexecuting them. *Communications of the ACM*, 58(2):74–84, January 2015.