Proof of Luck: an Efficient Blockchain Consensus Protocol

Mitar Milutinovic
UC Berkeley
mitar@cs.berkeley.edu

Howard Wu
UC Berkeley
howardwu@berkeley.edu

Warren He
UC Berkeley
-w@berkeley.edu

Maxinder Kanwal
UC Berkeley
mkanwal@berkeley.edu

ABSTRACT

In the paper, we present designs for multiple blockchain consensus primitives and a novel blockchain system, all based on the use of trusted execution environments (TEEs), such as Intel SGX-enabled CPUs. First, we show how using TEEs for existing proof of work schemes can make mining equitably distributed by preventing the use of ASICs. Next, we extend the design with proof of time and proof of ownership consensus primitives to make mining energy- and time-efficient. Further improving on these designs, we present a blockchain using a *proof of luck* consensus protocol. Our proof of luck blockchain uses a TEE platform's random number generation to choose a consensus leader, which offers low-latency transaction validation, deterministic confirmation time, negligible energy consumption, and equitably distributed mining. Lastly, we discuss a potential protection against up to a constant number of compromised TEEs.

CCS Concepts

•Security and privacy → Hardware-based security protocols; •Computer systems organization → Peer-to-peer architectures;

Keywords

Blockchain, Trusted Execution Environments, Consensus Protocol, Intel SGX

1. INTRODUCTION

Bitcoin [15], a widely-used blockchain system, as well as other popular cryptocurrencies [21], have demonstrated that it is practical to use a blockchain as a distributed ledger. Maintained by a peer-to-peer network of mutually distrusting participants, these systems use proof of work [1] to solve the key challenge of reaching consensus among participants.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SysTEX '16 December 12–16, 2016, Trento, Italy

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4670-2/16/12...\$15.00

DOI: <http://dx.doi.org/10.1145/3007788.3007790>

While proof of work is robust against misbehaving and malicious participants, these algorithms require participants to dedicate computation time, energy, and silicon towards contrived “work.” Moreover, to reduce the number of intermediary forks, which decrease the effective power of the network, Bitcoin’s proof of work is designed to produce a new block on average every 10 minutes. It is recommended to wait for 6 blocks before accepting a transaction, [19] which makes it impractical for many applications (e.g., point of sale transactions).

To address these shortcomings, we can use modern trusted execution environments (TEEs) to build new consensus primitives for use in decentralized electronic currency designs. The capabilities of TEEs can enforce correct processing of critical operations and can also limit the effect of Sybils running under single units of hardware.

In this paper, we propose a novel consensus algorithm, *proof of luck*, along with a proof-of-concept blockchain design using it, which achieves low-latency transaction validation while using minimal energy and computing power under rational attackers and benign participants. Our design utilizes the capabilities of Intel’s SGX platform, although any TEE platform with similar features will work. Our design scales to large numbers of participants and sidesteps the issue of ASICs in mining, allowing consumer-grade hardware to participate equally in the network.

Our contributions are as follows:

- We present three basic consensus primitives that use a TEE—proof of work, proof of time, and proof of ownership. These are energy efficient drop-in replacements for existing consensus primitives.
- We present a novel fourth consensus primitive—proof of luck—and a blockchain design that uses it to achieve low-latency transaction validation, deterministic confirmation time, negligible energy consumption, and equitably distributed mining.
- We discuss a potential protection against an attacker with a small number of compromised TEEs.

Our paper is structured as follows: we present related work in Section 2; formally define the problem we address in Section 3; present three drop-in TEE-enabled consensus primitives in Section 4; present our proof of luck blockchain in Section 5; summarize our analysis in Section 6; describe future directions in handling TEE compromise in Section 7; and finally conclude in Section 8.

2. RELATED WORK

To address the energy consumption of Bitcoin [16], various alternative consensus mechanisms have been proposed, e.g., proof of stake [10] and proof of burn [17]. It remains unclear whether these suitably maintain the security properties and incentives of current proof of work schemes [18]. Alternative approaches include replacing “useless” proof of work puzzles with meaningful problems [6, 9], making energy consumption less wasteful, but still requiring such computing resources to be available, which limits mobile and IoT devices from participating.

When there is a known set of participants, the practical Byzantine fault tolerance algorithm [3] allows a system to commit transactions and reach consensus within a few network communications in the presence of up to one-third of the participants exhibiting arbitrary, malicious behavior. However, this algorithm does not scale well in the number of participants, relies critically on network timing assumptions, and only guarantees liveness when all participants behave as expected. While several improvements have been proposed [11, 12, 14], these still require maintaining a list of participants or trust relations between subgroups of participants.

To improve transaction confirmation times in Bitcoin, several approaches have been proposed, e.g., using uncles [20] and dividing blocks into micro and macro blocks [5]. These approaches are orthogonal to our proposed advancements and can be used to further improve our blockchain system.

Intel is concurrently and independently developing an SGX-based distributed ledger as part of the Sawtooth Lake project [7]. While it simulates Bitcoin mining in an energy-efficient manner similar to proof of time, our work provides additional consensus schemes and benefits, most notably proof of luck, addresses the issue of possible compromised TEEs, and does not require estimation of the number of participants.

3. PROBLEM DEFINITION

The problem of consensus is that the distributed system must agree on a single shared state. Current designs of blockchains are slow, using significant time and energy as part of the consensus mechanism. Validation prevents arbitrary changes to state, but still allows the blockchain to “fork” into two or more valid continuations. A participant may be incentivized to prefer their own state over other participants’ states, e.g. to gain a block mining reward, or perform a double-spend attack [8]. The system must be able to determine the state of the participants, uncontrollable by a minority of colluding, malicious participants.

Our goal is to design a consensus algorithm with:

1. Quick, deterministic transaction confirmations.
2. Energy and network communication efficient protocol.
3. Resistant to custom hardware not commonly available.
4. An attacker controlling under a threshold of TEEs cannot control the blockchain.
5. An attacker cannot control the blockchain without controlling a majority of CPUs and without breaking the TEE platform.
6. No requirement for a synchronized clock between participants.

We assume the following about the setting of the problem:

1. The participants have CPUs that implement a suitable TEE, such as Intel SGX.
2. TEE programs can produce unbiased uniform random numbers, which an attacker cannot influence. (In Intel SGX, the `RDRAND` instruction is available.)
3. TEE programs can detect concurrent invocations. (In Intel SGX, this can currently be done by attempting to create the maximum allowed number of monotonic counters for each enclave identity, 256 at the time of writing.)

3.1 Principals

Participants, the main principals in our protocol, are required to use a TEE, and perform routines to maintain the blockchain and help others read and write to the blockchain. The *trusted platform vendor* controls the correct execution of the algorithm inside each participant’s TEE. *Clients* rely on the blockchain and communicate with participants to read and write to the blockchain. As clients do not authenticate participants, they must protect the content of the blockchain end-to-end, i.e., by signing any data to be added.

3.2 Threat Model

We analyze our protocol under attack from an adversary controlling some fraction (less than half) of the participants’ machines. We assume that the adversary cannot break any cryptographic primitives with non-negligible probability. It follows that, while the adversary can run the protocol in the TEE for each controlled machine, they cannot generate a valid attestation proof if they deviate from the protocol.

In Section 7, we consider an adversary that can generate arbitrary valid attestation quotes on a limited number of CPUs. We base our economic security model of TEE platforms on the fact that compromising many TEEs is prohibitively expensive.

Furthermore, the adversary can read and send its own network messages, but cannot modify messages sent by honest participants or cause a network split. Messages in our protocol do not identify the sender, although some implementations may choose to do so (e.g., using TCP/IP); in this case, it is assumed that the adversary can spoof the sender’s information in its own messages and tamper with the sender’s information from honest participants.

Attacks that compromise an entire trusted platform vendor are out of this paper’s scope, as such an attack would have consequences much broader than our blockchain.

3.3 SGX Background

Our work is based on Intel’s SGX platform [13]. In addition to the isolation and remote attestation features of a typical TEE, we utilize the SGX platform’s trusted services that provide relative timestamps and monotonic counters. SGX uses Enhanced privacy ID (EPID) [2] signatures, allowing anonymous and pseudonymous attestation. When initializing a TEE, the platform records a digest of the code and data (called the *measurement*) which serves as the cryptographic identity of the TEE. In the remote attestation protocol, the TEE certifies a computation result with a signed *report* and remotely verifiable *quote*, which serves as the proof that the results arose from appropriate computations on proper hardware.

4. BUILDING BLOCKS

Existing blockchains and other systems utilizing proof of work can benefit from using these TEE-enabled consensus building blocks directly. In this section, we progress through a series of designs that inform and motivate our proof of luck consensus protocol, which we present in Section 5.

4.1 Proof of Work

One major discussion in the cryptocurrency community is whether proof of work algorithms should be ASIC-resistant, preventing custom ASIC hardware from speeding up the computations. Using a TEE-enabled proof of work sidesteps this issue, as only supported platforms can be used for mining (e.g., Intel SGX-enabled CPUs). As a result, mining power is decentralized, mining rewards are equity distributed among participants, and mining pools bring no advantages.

Algorithm 1 TEE-enabled proof of work (inside TEE)

```

1: function POW(nonce, difficulty)
2:   result  $\leftarrow$  ORIGINALPOW(nonce, difficulty)
3:   assert ORIGINALPOWSUCCESS(result)
4:   return TEE.ATTESTATION( $\langle$ nonce, difficulty $\rangle$ , null)
5: end function

```

We propose Algorithm 1, which wraps any existing proof of work algorithm into a TEE-enabled one. Inside a TEE, we call the ORIGINALPOW algorithm with a given *nonce* and *difficulty*. In Bitcoin, *nonce* is the newly-mined block header, *difficulty* is the target hash, and ORIGINALPOWSUCCESS checks that ORIGINALPOW succeeded. TEE.ATTESTATION returns a *proof* attesting that the algorithm ran inside the TEE, and that the code was unmodified (based on the measurement). The second argument in our pseudocode, **null** here, calls for an anonymous *random base* EPID signature.

4.2 Proof of Time

Proof of work effectively enforces that a sufficient amount of time has passed by requiring the participant to do work; a TEE can enforce this directly by requiring the participant to wait for the desired time, saving CPU cycles and energy for other meaningful work. We now propose Algorithm 2.

Algorithm 2 Proof of time (inside TEE)

```

1: counter  $\leftarrow$  INCREMENTMONOTONICCOUNTER()
2: function POT(nonce, duration)
3:   SLEEP(duration)
4:   assert counter = READMONOTONICCOUNTER()
5:   return TEE.ATTESTATION( $\langle$ nonce, duration $\rangle$ , null)
6: end function

```

TEE platforms, like Intel SGX, provide relative timestamps to the TEE. The SLEEP function can busy-wait for *duration*, or it can yield control from the TEE for other outside processes, returning after *duration* and verifying that the platform’s relative timestamp has increased sufficiently.

A malicious participant may try to game proof of time by running multiple instances in parallel on the same CPU. By maintaining a monotonic counter, provided as a TEE platform service, we can prevent such malicious behavior by incrementing the counter (INCREMENTMONOTONICCOUNTER) every time the TEE starts. After control is transferred out of

the TEE and back, we check (READMONOTONICCOUNTER) that the counter has not changed.

4.3 Proof of Ownership

Proof of work can be seen as a countermeasure against Sybils, where an attacker acts as multiple participants—such an attacker would have to do many times the amount of work. When participants are required to use a TEE, instead of consuming resources, one can make it expensive to maintain such virtual participants by physically limiting each participant to owning a unique CPU. In SGX, we can operate the EPID signature (part of the remote attestation protocol) in *name base* mode, which produces a “pseudonym” that reveals if multiple proofs came from the same CPU. This works even if a malicious user resets the owner epoch register on their CPU.

Algorithm 3 Proof of ownership (inside TEE)

```

1: function POO(nonce)
2:   return TEE.ATTESTATION( $\langle$ nonce $\rangle$ , nonce)
3: end function

```

In Algorithm 3, we generate a *proof* using *nonce* (a block header) as the name base (second argument on Line 2). Proofs generated by a single TEE for the same name base bear the same pseudonym. A blockchain can then reach consensus by selecting the block that has the most proofs with unique pseudonyms as the leader.

5. PROOF OF LUCK

Our proof of luck blockchain is built upon the TEE-based proof of luck primitive proposed in Algorithm 4.

The algorithm consist of two functions, POLROUND and POLMINE. At the start of every round, the participant prepares the TEE to mine on a particular chain by calling POLROUND and passing the currently known latest *block*. After *ROUND_TIME* passes, the participant calls POLMINE to mine a new block. The participant passes the *header* of the new block and the block that it will extend (as *previousBlock*). *previousBlock* may be different from the *roundBlock* that was passed to POLROUND, but we require that *roundBlock* and *previousBlock* have the same *parent*. This ensures that a participant waits the mandatory *ROUND_TIME* between mining blocks, while allowing them to switch to a luckier, alternative block should they receive one while waiting. We verify that *ROUND_TIME* indeed has passed.

The POLMINE function generates a random value $l \in [0, 1)$ from a uniform distribution, which is used to determine the winning block from all mined blocks of participants in this round. This will tie in with the protocol in Section 5.1.

To optimize protocol communication, the algorithm delays releasing the *proof* depending on $f(l)$, which prescribes a shorter delay period for luckier (larger) numbers and a longer wait period for unluckier numbers. If a participant receives a luckier block before their own mining completes, they will not need to broadcast their own block.

We use a monotonic counter as in Section 4.2 to prohibit concurrent invocations of the TEE.

5.1 Protocol

Every participant in the protocol receives transactions from clients and other participants and maintains a current

Algorithm 4 Proof of luck primitive (inside TEE)

```
1: counter  $\leftarrow$  INCREMENTMONOTONICCOUNTER()
2: roundBlock  $\leftarrow$  null
3: roundTime  $\leftarrow$  null

4: function POLROUND(block)
5:   roundBlock  $\leftarrow$  block
6:   roundTime  $\leftarrow$  TEE.GETTRUSTEDTIME()
7: end function

8: function POLMINE(header, previousBlock)
  // Validating link between header and previousBlock.
9: assert header.parent = HASH(previousBlock)
  // Validating previousBlock matches roundBlock.
10: assert previousBlock.parent = roundBlock.parent
  // Validating the required time for a round passed.
11: now  $\leftarrow$  TEE.GETTRUSTEDTIME()
12: assert now  $\geq$  roundTime + ROUND_TIME

13: roundBlock  $\leftarrow$  null
14: roundTime  $\leftarrow$  null
15: l  $\leftarrow$  GETRANDOM()
16: SLEEP(f(l))

  // Validating that only one TEE is running.
17: newCounter  $\leftarrow$  READMONOTONICCOUNTER()
18: assert counter = newCounter

19: nonce  $\leftarrow$  HASH(header)
20: return TEE.ATTESTATION( $\langle$ nonce,l $\rangle$ , null)
21: end function
```

blockchain representing their view of the system state. In every round, participants execute an algorithm to commit pending transactions into a new block, containing a proof of luck generated from inside a TEE, and extending their current chain. Participants broadcast their new chain to other participants should the chain be luckier than any alternative chain received in the meantime.

Algorithm 5 Extending a blockchain with a new block

```
1: function COMMIT(newTransactions, chain)
2:   previousBlock  $\leftarrow$  LATESTBLOCK(chain)
3:   parent  $\leftarrow$  HASH(previousBlock)
4:   header  $\leftarrow$   $\langle$ parent, newTransactions $\rangle$ 
5:   proof  $\leftarrow$  POLMINE(header, previousBlock)
6:   newBlock  $\leftarrow$   $\langle$ parent, newTransactions, proof $\rangle$ 
7:   return APPEND(chain, newBlock)
8: end function
```

Participants use COMMIT in Algorithm 5 to return a new chain with a *newBlock* made from the set of *newTransactions*. The *newBlock* consists of a hash *parent* of the previous block, the data *newTransactions*, and a proof of luck *proof*.

The LUCK algorithm described in Algorithm 6 computes a numeric score (luck) of a given blockchain by summing the *l* values of each block.

Our blockchain prefers the chain with the highest luck, which results in desirable behaviors. First, when appending a block, the new chain is preferred. Second, when a network split heals, the larger half's chain is likely to have greater luck. Third, our analysis in Section 6.1 shows that, with high probability, a minority attacker's chain will not overtake the majority's.

Algorithm 6 Computing a luck of a valid blockchain

```
1: function LUCK(chain)
2:   luck  $\leftarrow$  0
3:   for block in chain do
4:     luck  $\leftarrow$  luck + TEE.PROOFDATA(block.proof).l
5:   end for
6:   return luck
7: end function
```

Algorithm 7 Validating a blockchain

```
1: function VALID(chain)
2:   previousBlock  $\leftarrow$  null
3:   while chain  $\neq$   $\epsilon$  do
4:     block  $\leftarrow$  EARLIESTBLOCK(chain)
5:      $\langle$ parent, transactions, proof $\rangle$   $\leftarrow$  block
6:     if parent  $\neq$  HASH(previousBlock)  $\vee$ 
       not VALIDTRANSACTIONS(transactions)  $\vee$ 
       not TEE.VALIDATTESTATION(proof) then
7:       return false
8:     end if
9:      $\langle$ nonce, l $\rangle$   $\leftarrow$  TEE.PROOFDATA(proof)
10:    if nonce  $\neq$  HASH( $\langle$ parent, transactions $\rangle$ ) then
11:      return false
12:    end if
13:    previousBlock  $\leftarrow$  block
14:    chain  $\leftarrow$  WITHOUTEARLIESTBLOCK(chain)
15:  end while
16:  return true
17: end function
```

The VALID algorithm in Algorithm 7 traverses the *chain* from the earliest (genesis) block (returned from EARLIESTBLOCK) to the latest block, ensuring each block has valid transactions, a valid proof of luck, and a matching previous block hash. TEE.VALIDATTESTATION is provided by the TEE platform to validate attestations, and TEE.PROOFDATA exposes data used when creating the *proof*.

We now introduce the proof of luck blockchain protocol in Algorithm 8. Every participant starts with an empty blockchain *currentChain* = ϵ , a set of pending *transactions* = ϵ , and an initial *roundBlock* = null.

After initializing state, participants listen for network messages. Upon receiving a *transaction* network message, participants add the *transaction* to their *transactions* if it has not been already included, and broadcasts it to their peers. If the message is a new *chain*, they verify that the chain is VALID (Algorithm 7) and has higher LUCK (Algorithm 6). If so, the participant switches to the new chain and broadcasts it using NETWORK.BROADCAST.

Before broadcasting, the participant may start a new round of mining. This happens if it is the first round, or if the new chain has a different *parent* block from its latest block. During one round of mining, the participant continues to receive messages about luckier chains and switches should it have higher luck. In this instance, the *parent* will not change. However, if a participant was part of a network split and just reconciled, the *parent* will differ and the participant will have to restart the mining process on the new chain.

Every time a new round is started (by calling NEWROUND), we call POLROUND to bind mining to the new chain, clear any pending callbacks, and schedule a new *callback* to start

Algorithm 8 Proof of luck blockchain protocol

```
1: currentChain  $\leftarrow \varepsilon$ 
2: transactions  $\leftarrow \varepsilon$ 
3: roundBlock  $\leftarrow \text{null}$ 

4: function NEWROUND(chain)
5:   roundBlock  $\leftarrow \text{LATESTBLOCK}(\text{chain})$ 
6:   POLROUND(roundBlock)
7:   RESETCALLBACK(callback, ROUND_TIME)
8: end function

9: on transaction from NETWORK
10:  if transaction  $\notin$  transactions then
11:    transactions  $\leftarrow \text{INSERT}(\text{transactions}, \text{transaction})$ 
12:    NETWORK.BROADCAST(transaction)
13:  end if
14: end on

15: on chain from NETWORK
16:  if VALID(chain)  $\wedge$  LUCK(newChain) >
    LUCK(oldChain) then
17:    currentChain  $\leftarrow$  chain
18:    if roundBlock = null then
19:      NEWROUND(chain)
20:    else
21:      latestBlock  $\leftarrow \text{LATESTBLOCK}(\text{chain})$ 
22:      if latestBlock.parent  $\neq$  roundBlock.Parent then
23:        NEWROUND(chain)
24:      end if
25:    end if
26:    NETWORK.BROADCAST(chain)
27:  end if
28: end on

29: on callback
30:  newTransactions  $\leftarrow$  transactions
31:  transactions  $\leftarrow \varepsilon$ 
32:  chain  $\leftarrow \text{COMMIT}(\text{newTransactions}, \text{currentChain})$ 
33:  NETWORK.SENDTOSELF(chain)
34: end on
```

after *ROUND_TIME*. During normal operation, this means that every participant mines a new block approximately once every *ROUND_TIME* interval. Participants do not have to have a synchronized clock and their rounds do not have to be synchronized, but the protocol tends to synchronize rounds.

Inside a *callback*, pending transactions are appended to a new *chain* using COMMIT (Algorithm 5) and the new chain is sent back for processing using NETWORK.SENDTOSELF.

Because the proof of luck TEE releases the proof for luckier (larger) numbers first, participants will receive new chains with numbers luckier than their own before receiving their own from the TEE. If the participant receives a chain from its TEE that is worse than the one that they have already received, they do not broadcast its chain this round.

6. ANALYSIS

We analyze the behavior of the proof of luck blockchain protocol.

6.1 Persistence against Minority Attacker

We show that several blocks after a fork, it is exponentially improbable for a minority attacker to produce a chain that

is preferred to a majority of honest participants.

Let M = majority population size, m = minority population size. For block t , we have that the population luck is distributed according to the maximum of uniform random variables:

$$l_M(t) \stackrel{\text{id}}{\sim} \max\{\{\text{Uniform}(0, 1)\}^M\}$$
$$l_m(t) \stackrel{\text{id}}{\sim} \max\{\{\text{Uniform}(0, 1)\}^m\}$$

After h blocks from a fork, we define the relative total luck:

$$L^{(h)} := \sum_{t=1}^h l_M(t) - l_m(t)$$

Using a Chernoff bound and independence assumptions, we now show that the probability of the event that a minority wins is exponentially small in h :

$$\begin{aligned} \Pr(L^{(h)} \leq 0) &\leq \min_{s>0} \mathbb{E} \left[e^{-sL^{(h)}} \right] \\ &= \min_{s>0} \prod_{t=1}^h \mathbb{E} \left[e^{-sl_M(t)} \right] \mathbb{E} \left[e^{sl_m(t)} \right] \\ &= \min_{s>0} \left(\mathbb{E} \left[e^{-sl_M(t)} \right] \mathbb{E} \left[e^{sl_m(t)} \right] \right)^h \end{aligned}$$

Because $M > m$, there exists an $s > 0$ such that the product of the inner expectations is less than 1.

Hence we have that the probability of the event that the total minority luck exceeds the total majority luck decreases exponentially in the number of blocks, h , after a fork.

6.2 Proportional Control of Blocks

Consider a group of participants $A \subseteq G$, from all participants G . Among honest participants that append to the longest chain, the new chain with the largest luck value l in the newly added block is preferred. At each round, every participant has an equal probability of generating the largest random number, because the participants sample independently from identical distributions. Thus over time, the expected number of blocks mined by group A is proportional to the number of participants in A .

6.3 Round Time and Confirmation Time

We propose a *ROUND_TIME* of 15 seconds, which creates block confirmation times slightly larger than 15 seconds, comparable with Ethereum and much faster than Bitcoin's 10 minutes. We have chosen this value based on an evaluation of information in the Bitcoin network [4], where median block propagation time has been observed to be around 6.5 seconds. The initial selection of the winning block with maximal luck value l can be implemented without transmitting whole blocks, only block headers (i.e., without the full transaction data), which should take just one round-trip based on Bitcoin's transaction propagation time, and only after a winner is determined should the whole block be propagated.

7. COMPROMISED TEE

Our consensus protocol assumes that the security requirements of the TEE are met for all participants. Although it would be expensive to violate these requirements, a motivated attacker may compromise a limited number of TEEs. We discuss a mitigation extension to our protocol.

7.1 Luckiest m

To strengthen our design against high-cost attacks on individual TEEs, we discuss a possible extension to our protocol, which constructs a blockchain consisting of *super-blocks*: super-blocks are made by merging m normal blocks and their proofs of luck. Participants continue to create individual blocks based on pending transactions and attach a proof of luck with a luck value l to the block. However, participants now select the m luckiest blocks (m blocks with highest l) and merge, in a deterministic way, those m blocks into a super-block.

Each super-block has m proofs of luck, with values $l_1, \dots, l_m$, where $l_1 > \dots > l_m$ (i.e., l_1 is the luckiest). Under this extension, the luck of the super-block as a whole is l_m , the least luckiest value. This ensures that even if $m - 1$ CPUs have been compromised, they still cannot fully control the l value of the super-block. Future work may analyze the persistence property of this extension.

This approach is similar to our proof of ownership primitive from Section 4.3. However, instead of recording proofs from all participants, we record only the luckiest m , making this approach more scalable. In our SGX-based prototype, we ensure that the m proofs come from different CPUs by having the participants generate linkable quotes during remote attestation, using the parent block's hash as a name base in the EPID signature scheme. EPID signatures produce a pseudonym based on the name base and the CPU's private key, and proofs with the same pseudonym would reveal that they came from the same physical CPU. Thus, even if a CPU is compromised, it is still limited to supplying one proof of luck per super-block.

7.2 Merging

After the honest participants produce their own blocks, they broadcast their block to other participants that chose the same blockchain to extend (they may choose different blockchains due to network latency). Participants then merge the m luckiest blocks into a super-block. Blocks from honest participants have nearly the same transactions, so a super-block can often be compressed efficiently. Having participants merge blocks into super-blocks is similar to Ethereum's concept of block uncles, in which a block may have multiple parents, and which improves the throughput and liveness of the protocol.

8. CONCLUSION

We have proposed three TEE-enabled building block designs for systems using existing proof of work schemes: TEE-enabled proof of work, proof of time, and proof of ownership. We have combined ideas from these primitives into a new blockchain system based on the proof of luck consensus protocol. Our analysis shows that our blockchain ensures liveness and persistence, while providing energy efficient mining, low-latency transaction validation with deterministic confirmation time, and decentralized mining power. Lastly, we discuss a potential protection against a constant number of compromised TEEs participating in the blockchain.

9. ACKNOWLEDGEMENTS

We thank Dawn Song for helpful suggestions. This work was supported by IC3, NSF (TWC-1518899), DARPA (N66001-15-C-4066). Any opinions, findings, and conclusions or rec-

ommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSF and/or DARPA.

10. REFERENCES

- [1] A. Back et al. Hashcash: a denial of service counter-measure, 2002.
- [2] E. Brickell and J. Li. Enhanced privacy ID: A direct anonymous attestation scheme with enhanced revocation capabilities. In *Proceedings of the 2007 ACM workshop on Privacy in electronic society*, pages 21–30. ACM, 2007.
- [3] M. Castro, B. Liskov, et al. Practical Byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.
- [4] C. Decker and R. Wattenhofer. Information propagation in the bitcoin network. In *Peer-to-Peer Computing (P2P), 2013 IEEE Thirteenth International Conference on*, pages 1–10. IEEE, 2013.
- [5] I. Eyal et al. Bitcoin-ng: A scalable blockchain protocol. *arXiv preprint arXiv:1510.02037*, 2015.
- [6] R. Halford. Gridcoin, 2013.
- [7] Intel Corporation. Sawtooth lake, 2016.
- [8] G. O. Karame et al. Double-spending fast payments in bitcoin. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security, CCS '12*, pages 906–917, New York, NY, USA, 2012. ACM.
- [9] S. King. Primecoin: Cryptocurrency with prime number proof-of-work. *July 7th*, 2013.
- [10] S. King and S. Nadal. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake, 2012.
- [11] J. Kwon. Tendermint: Consensus without mining, 2014.
- [12] D. Mazières. The Stellar consensus protocol: A federated model for Internet-level consensus, 2015.
- [13] F. McKeen et al. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*. ACM, 2013.
- [14] A. Miller et al. The honey badger of BFT protocols. In *CCS*, 2016. To appear.
- [15] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.
- [16] K. J. O'Dwyer and D. Malone. Bitcoin mining and its energy footprint. In *ISSC 2014/CHICT 2014, 25th IET*, pages 280–285. IET, 2013.
- [17] P4Titan. Slimcoin: A peer-to-peer crypto-currency with proof-of-burn, 2014.
- [18] A. Poelstra et al. Distributed consensus from proof of stake is impossible, 2014.
- [19] M. Rosenfeld. Analysis of hashrate-based double spending. *arXiv preprint arXiv:1402.2009*, 2014.
- [20] Y. Sompolinsky and A. Zohar. Accelerating Bitcoin's transaction processing. fast money grows on trees, not chains. *IACR Cryptology ePrint Archive*, 2013:881, 2013.
- [21] G. Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper*, 2014.