



Limits on Revocable Proof Systems, With Implications for Stateless Blockchains

Miranda Christ^{1,3} and Joseph Bonneau^{2,3}

¹ Columbia University, New York, USA

mchrist@cs.columbia.edu

² New York University, New York, USA

jcs@cs.nyu.edu

³ a16z Crypto Research, San Francisco, USA

Abstract. Motivated by the goal of building a cryptocurrency with succinct global state, we introduce the abstract notion of a revocable proof system. We prove an information-theoretic result on the relation between global state size and the required number of local proof updates as statements are revoked (e.g., coins are spent). We apply our result to conclude that there is no useful trade-off point when building a stateless cryptocurrency: the system must either have a linear-sized global state (in the number of accounts in the system) or require a near-linear rate of local proof updates. The notion of a revocable proof system is quite general and also provides new lower bounds for set commitments, vector commitments and authenticated dictionaries.

Keywords: Stateless Blockchains · Authenticated Data Structures

1 Introduction

Modern cryptocurrencies prevent double-spending attacks using a public, append-only log called a blockchain. Classically, a blockchain records *all* transactions, and validating a new transaction requires checking that it doesn't conflict with any prior transaction. This approach was first successfully deployed by Bitcoin [22] though it was proposed earlier [14].

A challenge of the blockchain paradigm is that each validator traditionally must store the entire *state* of the system. In Bitcoin, this consists of a set of unspent transaction outputs (UTXOs), which has consistently grown and now contains 80 million elements, requiring several GB to store. Ethereum's state is even larger [31], requiring roughly 35 GB to represent 200 million accounts.

The requirement that validators store this large (and growing) state raises concerns about centralization if the state grows so large that only well-funded organizations can afford to store it. As a result, most blockchain systems impose strict limits on state growth, which in turn limit transaction throughput. Famously, Bitcoin originally imposed a maximum size of 1 MB per block, limiting throughput to about three transactions per second.

The tension between throughput and state growth leads to a natural question: can we achieve high throughput with a small (perhaps even constant-sized) global state? This led to the proposal of *stateless* blockchain designs [27], although this term is a misnomer: they typically assume validators store a small commitment to the global state of the system (e.g., a Merkle root committing to the set of unspent coins). Users wishing to make a transaction must publish a *witness* that their transaction is valid given the current state commitment (e.g., a Merkle proof that a coin is included in the valid set). Validators can then accept transactions without knowing the full state of the system. Since Todd’s original proposal using Merkle trees, several other designs have been proposed using Merkle-tree-based accumulators [6, 16], RSA accumulators [4], and vector commitments [12, 18, 26, 28, 30]. Stateless blockchains are distinct from *succinct blockchains* [1, 11, 20] such as the Mina protocol [5]. Succinct blockchains use verifiable computation to achieve $O(1)$ storage and verification costs for light clients, but still require validators to store the entire system state in order to process new transactions and build (and prove correct) new blocks.

Unfortunately, all known stateless blockchain designs introduce a new problem: users’ witnesses can become invalid as other (unrelated) transactions update the global state, requiring users to monitor the network and periodically refresh their witnesses. This is a departure from the traditional blockchain model, in which users can stay offline for long periods of time and then successfully create and broadcast a transaction. This is not simply a matter of convenience; there are important security benefits of supporting offline participation, as private keys can be kept in air-gapped machines such as hardware wallets.

In this work we show that, regrettably, the trade-off between a large global state and requiring frequent witness changes is fundamental. More specifically, in Theorem 1 we show a lower bound on the global state size as a function of the number of revoked statements and the desired maximum number of witness changes. In Corollary 1, we show that there is no trade-off which does not require either an (asymptotically) linear-sized global state or an (asymptotically) near-linear number of witness updates as a constant fraction of coins are spent.

Model. To analyze the efficiency of stateless blockchains and similar authenticated data structures, we introduce a new cryptographic notion: a *revocable proof system* (RPS, Sect. 2). An RPS is a simple abstraction capturing a class of schemes that involve a *global state* V encapsulating a set S of *valid statements*. Correctness ensures that each valid statement $s_i \in S$ has a corresponding *proof* π_i which can be efficiently verified given s_i , the public parameters, and the global state. A subset $T \subseteq S$ of the initial set of valid statements may later be *revoked*, yielding an updated global state V' . Security requires that these revoked statements’ proofs no longer verify. This functionality is quite natural and captures a wide range of useful cryptographic notions, including accumulators and vector commitments. We discuss these connections in Sect. 4.

Contributions. Using our revocable proof system definition we prove a trade-off between the size of the global state and the frequency with which proofs must

be updated (Theorem 1). We do so using a compression argument: if the global state is small and with constant probability there are few ($\leq k$) proof updates, an adversary can use the global state and a small amount of additional information to encode the revoked set. We apply Shannon’s Coding Theorem to show a lower bound on the size of the global state given the number of proof updates k .

As a corollary, we observe that there is no useful asymptotic trade-off between the size of the global state and the number of proof changes: either the state size is linear, or a (nearly) linear number of proofs must change (Sect. 3.1). As a second corollary, we show that a useful notion of *persistence* (proofs of certain statements never change) requires linear storage for these persistent statements. In Sect. 4, we show that accumulators, vector commitments, and authenticated dictionaries fit the framework of a revocable proof system and thus our lower bound applies to them, giving results of independent interest.

Implications. Finally, we discuss the implications of our results on stateless blockchain proposals (Sect. 5). We plot the minimum number of required witness changes per day for blockchains with practical transaction rates and global state sizes. For a blockchain with a transaction rate on the scale of Visa, the number of witness changes is infeasibly large for any meaningfully compressed global state. We discuss three ideas for mitigating the witness update issue and analyze them in light of our impossibility result. The first idea is a versioning system, which stores explicitly all transactions occurring during the current epoch and consequently requires no witness updates within this epoch. The second is a scheme where users lock up their coins for a period of time in exchange for a guarantee that their witnesses will remain unchanged during that time. The third and most promising is the introduction of new state-storing third parties, which are neither users nor validators, called proof-serving nodes.

2 Model

Notation. We use λ to denote the security parameter. We use \lg to denote a logarithm with the base 2. Right \rightarrow and left \leftarrow arrows denote the output of a (possibly randomized) algorithm.

A *revocable proof system* (RPS) maintains a *global state* V , a *valid set* S , and a set of proofs π_i for each element $s_i \in S$ which we’ll also call a *statement*. The global state commits to the valid set, such that proofs of elements s_i in S can be verified. More formally, a *revocable proof system* is a tuple of algorithms $(\text{Setup}, \text{ComputeState}, \text{Revoke}, \text{Verify})$ where:

$\text{Setup}(1^\lambda) \rightarrow \text{pp}$ is a randomized algorithm that takes as input 1^λ , where λ is the security parameter, and outputs public parameters pp .

$\text{ComputeState}(\text{pp}, S) \rightarrow V, (\pi_1, \pi_2, \dots, \pi_n)$ is a deterministic algorithm that takes as input the public parameters pp and a valid set S of size n . It outputs the corresponding global state V and a list of proofs $(\pi_1, \pi_2, \dots, \pi_n)$ where π_i is the proof for $s_i \in S$.

$\text{Revoke}(\text{pp}, S, T, V, (\pi_1, \pi_2, \dots, \pi_n)) \rightarrow V', (\pi'_1, \pi'_2, \dots, \pi'_n)$

is a deterministic algorithm that takes as input the public parameters pp , the initial valid set S , a revoked set $T \subseteq S$, an initial global state V , and a list of proofs for elements in S . It outputs an updated global state V' and updated proofs.¹

$\text{Verify}(\text{pp}, V, s_i, \pi_i) \rightarrow \{\text{true}, \text{false}\}$ is a deterministic algorithm that takes as input the public parameters pp , a global state V , a statement s_i , and a proof π_i . It outputs true or false.

A revocable proof system must be correct and secure. By correct, we mean that genuine proofs for valid elements should verify against the corresponding global state. By secure, we mean that it should be difficult to find a proof for a revoked element; that is, it should be computationally hard for an adversary to produce a revoked set such that a proof for a revoked element still verifies against the updated global state. More formally, correctness and security are defined as follows:

Definition 1 (Correctness²). A revocable proof system is *correct* if for every set S , every set $T \subseteq S$, and every $s_i \in S \setminus T$,

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda) \\ V, (\pi_1, \pi_2, \dots, \pi_n) \leftarrow \text{ComputeState}(\text{pp}, S) \\ V', (\pi'_1, \pi'_2, \dots, \pi'_n) \leftarrow \text{Revoke}(\text{pp}, S, T, V, (\pi_1, \pi_2, \dots, \pi_n)) \\ \text{Verify}(\text{pp}, V, s_i, \pi_i) = \text{true} \\ \text{Verify}(\text{pp}, V', s_i, \pi'_i) = \text{true} \end{array} \right] = 1$$

Definition 2 (Security). A revocable proof system is *secure* if for every p.p.t. adversary A ,

$$\Pr \left[\begin{array}{l} \text{pp} \leftarrow \text{Setup}(1^\lambda) \\ S, T, s^*, \pi^* \leftarrow A(1^\lambda, \text{pp}) \\ V, (\pi_1, \pi_2, \dots, \pi_n) \leftarrow \text{ComputeState}(\text{pp}, S) \\ V', (\pi'_1, \pi'_2, \dots, \pi'_n) \leftarrow \text{Revoke}(\text{pp}, S, T, V, (\pi_1, \pi_2, \dots, \pi_n)) \\ s^* \in T \\ \text{Verify}(\text{pp}, V', s^*, \pi^*) = \text{true} \end{array} \right] \leq \text{negl}$$

3 Main Result

Our main result is an inequality describing the relationship between the size of the global state and the number of proofs of valid statements which must be

¹ Although for ease of notation the output includes n proofs, security dictates that proofs for elements in T should not verify.

² While correctness with probability 1 is standard for the schemes we consider, our main result (Theorem 1) still holds for a relaxed notion of correctness which holds with overwhelming probability. In fact, it does not rely on correctness at all and rather on the prevalence of a notion of a k -good revoked set. A revocable proof system that is secure, correct with overwhelming probability, and has few witness changes should have many k -good revoked sets.

updated. We first introduce the notion of a k -good revoked set; that is, a set of statements such that when these statements are revoked, none of their proofs still verify, and at most k still-valid statements' proofs must be changed.

Definition 3 (k -good revoked set). We say a revoked set T is k -good given a revocable proof system RPS , an initial valid set S , size parameter $n = |S|$, and public parameters pp if, for $V, (\pi_1, \pi_2, \dots, \pi_n) \leftarrow \text{ComputeState}(\text{pp}, S)$ and $V', (\pi'_1, \pi'_2, \dots, \pi'_n) \leftarrow \text{Revoke}(\text{pp}, V, T)$ if:

1. For every revoked statement $s_i \in T$, $\text{Verify}(\text{pp}, V', s_i, \pi_i) = \text{false}$
2. There are at most k non-revoked statements $s_j \in (S \setminus T)$ such that $\text{Verify}(\text{pp}, V', s_j, \pi_j) = \text{false}$

Condition (1), that the original proof of a revoked statement no longer verifies, is a consequence of the security requirement that should hold for most revoked sets. Condition (2) is that few ($\leq k$) proofs of non-revoked statements no longer verify (i.e., need to be updated). Suppose that we want to have a secure revocable proof system such that most of the time, at most k non-revoked statements change when a set of size m is revoked. This is equivalent to having many k -good revoked sets of size m . By security (regardless of $|V|$), an overwhelming fraction of sets of size m must satisfy condition (1); otherwise an adversary could choose a set of size m at random, revoke it, and succeed in finding a proof for a revoked element. Condition (2) is exactly the other property we want: that at most k non-revoked statements change when our set is revoked. Therefore, our desired revocable proof system must have many k -good revoked sets of size m .

We now show that if any public parameters yield many k -good revoked sets, the size of the global state must be large. In other words, if the size of the global state is small, many proofs of non-revoked statements must change when an average set is revoked.

Theorem 1. *Let $\text{RPS} = (\text{Setup}, \text{ComputeState}, \text{Revoke}, \text{Verify})$ be a revocable proof system satisfying correctness, and let pp be any public parameters occurring with nonzero probability over $\text{Setup}(1^\lambda)$. Let S be a set of size n , and let \mathcal{X}_k^* denote the set of subsets $T \subseteq S$ that are k -good given RPS, S , and public parameters pp . Then $|V|$, the size of the global state in bits, satisfies*

$$|V| \geq \lg |\mathcal{X}_k^*| - \lceil k \lg n \rceil$$

Proof. We show that if $|V|$ is any smaller, there exists an efficient encoding of \mathcal{X}_k^* using fewer than $\lg |\mathcal{X}_k^*|$ bits, contradicting Shannon's Coding Theorem [25]. Note that \mathcal{X}_k^* can be computed by trying every revoked set $T \subseteq S$ and determining whether it fits the conditions of a k -good revoked set. While this algorithm is not efficient, it does not need to be as the contradiction we derive is compression beyond information theoretic limits, which poses no computational bounds on the communicating parties.

Consider two parties A and B interacting with a challenger in a game given as input S and pp . The goal is for A to succinctly encode a uniformly chosen revoked

set $T \subseteq S$ for B to decode. The challenger computes the initial global state and proofs $V_0, (\pi_1, \pi_2, \dots, \pi_n) \leftarrow \text{ComputeState}(\text{pp}, S)$. The challenger passes $V_0, (\pi_1, \pi_2, \dots, \pi_n)$ to both A and B. A chooses T uniformly at random from \mathcal{X}_k^* and computes the updated global state $V \leftarrow \text{Revoke}(\text{pp}, V_0, T)$. Then, for each $s_i \in (S \setminus T)$, A checks whether its proof verifies; i.e., whether $\text{Verify}(\text{pp}, V, s_i, \pi_i) = \text{true}$. If *not*, A adds s_i to a list L of still-valid statements with changed proofs. A sends V and L to B.

We now show that given V and L , B can decode T exactly. Let B's decoding be the set T' consisting of all statements s_i such that $\text{Verify}(\text{pp}, V, s_i, \pi_i) = \text{false}$ and $s_i \notin L$. First, any statement $s_i \in T$ must be in T' , since by definition of a k -good revoked set³, no proofs for revoked statements verify. Therefore, $T \subseteq T'$. Next, B's decoding algorithm ensures any statement $s_i \in T'$ is not in L , and $\text{Verify}(\text{pp}, V, s_i, \pi_i) = \text{false}$. All elements s_j that were not revoked (i.e., not in T) and whose proofs no longer verify ($\text{Verify}(\text{pp}, V, s_j, \pi_j) = \text{false}$) are included in L . Since s_i is not in L , it must in fact have been revoked, so $s_i \in T$. Therefore, $T' \subseteq T$, which implies that $T' = T$.

Finally, we observe that A can encode L by listing a $(\lg n)$ -bit representation of each of its elements. Since $|L| \leq k$ by definition of \mathcal{X}_k^* , this encoding takes at most $\lceil k \lg n \rceil$ bits, and A sends $|V| + \lceil k \lg n \rceil$ bits in total after choosing T . Since T was chosen uniformly from \mathcal{X}_k^* , and the entropy of the uniform distribution over \mathcal{X}_k^* is $\lg |\mathcal{X}_k^*|$, we have by Shannon's Coding Theorem that $|V| + \lceil k \lg n \rceil \geq \lg |\mathcal{X}_k^*|$. \square

3.1 No Useful Trade-Offs for Sublinear State Size

We show that under certain regimes (when $|\mathcal{X}_k^*|$ includes at least a constant fraction of subsets of size m for $\lg n \leq m \leq \frac{n}{2}$), there is no useful trade-off between the global state size and the frequency of proof changes when m elements are deleted. That is, the global state size is either linear in the size of the stored set, or $\Omega\left(\frac{m}{\lg n}\right)$ proofs must be updated.

Corollary 1 (No useful trade-offs). *Let n be the size of the initial valid set S and $m \leq \frac{n}{2}$ be the number of deleted elements.⁴ If $|\mathcal{X}_k^*|$ includes at least a constant fraction of subsets $T \subseteq S$ of size m , and the global state size is $|V| = o(\lg \binom{n}{m})$, then $k = \Omega\left(\frac{m}{\lg n}\right)$.*

Proof. This holds by a straightforward application of Theorem 1. First, observe that the number of possible $T \subseteq S$ of size m is $\binom{n}{m} \geq \frac{2^m}{m^m} \geq 2^m$. \mathcal{X}_k^* includes at least a constant fraction of these subsets T of size m , so $\lg |\mathcal{X}_k^*| = \Omega(\lg \binom{n}{m}) =$

³ It is tempting to instead cite security of a revocable proof system here, but security guarantees only that for *most* revoked sets T , proofs of revoked statements do not verify. Our definition of k -good gives us exactly what we need.

⁴ If more than $\frac{n}{2}$ elements are deleted in sequence, as in stateless blockchains, we can set $m = \frac{n}{2}$ since there must be an intermediate point where $\frac{n}{2}$ elements were deleted, and this bound still applies.

$\Omega(m)$. In order for the inequality from Theorem 1 to hold, we must have $k \lg n \geq \Omega(\lg \binom{n}{m}) - o(\lg \binom{n}{m})$, or $k = \frac{\Omega(\lg \binom{n}{m})}{\lg n} = \Omega\left(\frac{m}{\lg n}\right)$.

This bound on k holds for *any* global state size $|V|$ that is sublinear in $\lg \binom{n}{m}$. Once the global state size becomes $\Omega(\lg \binom{n}{m})$, we can (asymptotically) store the full list of deleted elements and require no witness updates. One especially interesting regime for this bound is when $m = \Theta(n)$ and $|V| = o(n)$. Then Corollary 1 implies that $k = \Omega\left(\frac{n}{\lg n}\right)$. In other words, if we want to avoid a near-constant fraction of proof updates, we need a linear global state size, at which point we can (asymptotically) store the full state naively and require no witness updates. This suggests that there is no asymptotically useful trade-off between global state size and number of proof changes in this regime; at least one of the two must be (nearly) linear.

3.2 Persistence Requires Linear Storage

We now show that another desirable property, which we call *persistence*, is also not possible without linear global state. Suppose that we want proofs of certain statements to always verify as long as those statements remain true. This guarantee would be very useful in cryptocurrencies, allowing a user to stay offline until she is ready to make a transaction, without fear of her proof becoming stale. We call this notion *persistence* and formalize it below.

Definition 4 (Persistence). A statement $s_i \in S$ is *persistent* given initial valid set S of size n and public parameters pp if for all $T \subseteq S$ such that $s_i \notin T$,

- $V, (\pi_1, \pi_2, \dots, \pi_n) \leftarrow \text{ComputeState}(\text{pp}, S)$
- $V', (\pi'_1, \pi'_2, \dots, \pi'_n) \leftarrow \text{Revoke}(\text{pp}, S, T, V, (\pi_1, \pi_2, \dots, \pi_n))$
- $\text{Verify}(\text{pp}, V', s_i, \pi_i) = \text{true}$

A corollary of Theorem 1 shows that there can be very few persistent statements:

Corollary 2 (Persistence requires linear storage). *Let RPS be a secure and correct revocable proof system such that there exists an initial set S and a set $S^* \subseteq S$ such that*

$$\Pr_{\text{pp} \leftarrow \text{Setup}(1^\lambda)} [\text{every } s \in S^* \text{ is persistent}] > \frac{1}{2}$$

Then the the global state of RPS has size at least $|S^| - 1$.*

Proof. Let S be any initial set and S^* be any subset of S . We wish to show that with high probability, \mathcal{X}_0^* is large, where \mathcal{X}_0^* is the family of revoked sets that require no witness changes and for which proofs of revoked statements do not verify. We first argue that by security, few revoked sets $T \subseteq S^*$ yield proofs of revoked statements that still verify. Then it follows that \mathcal{X}_0^* contains all other

revoked subsets of S^* , since by definition of persistence they require no witness changes.

Suppose for the sake of contradiction that for *all* parameters $\mathbf{pp} \leftarrow \text{Setup}(1^\lambda)$ that occur with nonzero probability and for which every $s \in S^*$ is persistent, more than half of the revoked sets $T \subseteq S^*$ yield a global state such that the proof of a revoked statement verifies. Then the following adversary A forges a proof with non-negligible probability, breaking security. Let A compute $\mathbf{pp} \leftarrow \text{Setup}(1^\lambda)$ and $V, (\pi_1, \pi_2, \dots, \pi_n) \leftarrow \text{ComputeState}(\mathbf{pp}, S)$. A then chooses $T \subseteq S^*$ uniformly at random, computes $V', (\pi'_1, \pi'_2, \dots, \pi'_n) \leftarrow \text{Revoke}(\mathbf{pp}, V, T)$, and checks whether $\text{Verify}(\mathbf{pp}, V', s_i, \pi_i)$ for each $s_i \in T$. If A finds such an s_i , it outputs S, T, s_i, π_i . Independently, A chooses \mathbf{pp} such that all of S^* is persistent with probability at least $\frac{1}{2}$ and T such that the proof of a revoked statement verifies with probability at least $\frac{1}{2}$. Thus, A is efficient and succeeds with probability $\frac{1}{4}$, contradicting security. Therefore, there must be some parameters \mathbf{pp} occurring with nonzero probability such that all of S^* is persistent and at least half of the revoked sets $T \subseteq S^*$ have no revoked statements whose original proofs verify. The family of these sets is exactly \mathcal{X}_0^* , whose size is at least $\frac{1}{2} \cdot 2^{|S^*|}$.

Thus, there exist public parameters \mathbf{pp} occurring with nonzero probability such that $|\mathcal{X}_0^*| \geq \frac{1}{2} \cdot 2^{|S^*|}$. Applying Theorem 1 for $k = 0$, we have that the size of the global state is at least $\lg |\mathcal{X}_k^*| = |S^*| - 1$.

4 Implications for Authenticated Data Structures

We show that cryptographic accumulators, vector commitments, and authenticated dictionaries, are instances of revocable proof systems. Thus, our lower bound from Theorem 1 applies. This result is of interest since these data structures are frequently used in distributed settings, in which users maintain proofs of portions of the committed data that are verified against a global state. Our bound dictates that these users must update their proofs frequently as the global state changes.

4.1 Cryptographic Accumulators

A cryptographic *accumulator* [3, 9], also called a *set commitment*, commits to an accumulated set X via a succinct digest A . Different accumulator schemes support efficiently proving various properties about the accumulated set X , such as membership or non-membership of elements. Some schemes may also allow X to be modified and the corresponding proofs updated. A typical accumulator supports additions, deletions, and membership proofs. That is, given a set X there is a function computing a digest A and a membership proof (also called a *witness*) w_i for each $x_i \in X$, corresponding to the ComputeState function of a revocable proof system. When a new element x is added to X , a new global state A' can be computed using x and A . Furthermore, each membership proof w_i can be updated given x and A . When an element $x_i \in X$ is deleted, a new global state A' can be computed using x_i , A , and the membership proof w_i for

x_i . The membership proofs of the other elements of X can be updated using the same information. Some accumulator schemes also allow batch updates, where multiple elements can be efficiently added and/or deleted at once [4].

Constructing a Revocable Proof System Using an Accumulator. We show how an RPS can be constructed using an accumulator scheme Acc supporting addition, deletion, and membership proofs. Addition is only necessary for the initial set S . The Setup function for our RPS calls the Setup function for Acc . The ComputeState function for our RPS, given public parameters pp and a valid set S , adds S to our accumulator given pp to obtain a digest A and a membership witness w_i for each $s_i \in S$. We let the proof π_i for s_i be this membership witness w_i . We implement Revoke for our RPS by, given a set $T \subseteq S$, removing T from the accumulated set and updating all witnesses according to the accumulator scheme. The resulting global state is the resulting accumulator value A' , and the resulting proofs π'_i are the updated witnesses w'_i . We let the Verify function for our RPS be the same as the Verify function for the accumulator scheme.

Accumulator schemes have correctness and security definitions that are analogous to those of a revocable proof system; full definitions can be found in [8]. By correctness of the accumulator, membership witnesses for elements of the accumulated set (equivalently, valid statements) verify. By security of the accumulator, it is hard for an adversary to find verifying membership witnesses for elements not in the accumulated set (equivalently, revoked or invalid statements). Thus, this construction is indeed a revocable proof system, and our lower bound from Theorem 1 applies.

We note that we can also construct a revocable proof system using an accumulator that supports only addition and non-membership witnesses (but not deletion). Given a finite data universe U and a set $X \subseteq U$, a delete/membership accumulator storing X can be implemented using an add/non-membership accumulator storing $U \setminus X$.

Camacho-Hevia Result. Our accumulator lower bound is reminiscent of a lower bound proved by Camacho and Hevia [8]. They consider a dynamic accumulator supporting addition, deletion, and membership proofs. Their model allows batch updates: if w_1, \dots, w_n are witnesses for an initial accumulated set X , after deletion of a set T the state-update function outputs a string $Upd_{X, X \setminus T}$ that can be used to update all witnesses to w'_1, \dots, w'_n to reflect the updated state. They show that if there are $|T| = m$ deletions, $Upd_{X, X \setminus T}$ must have length $\Omega(m)$. Baldimtsi et al. show an analogous result for a universal accumulator supporting addition, deletion, and *non-membership* proofs, using the same proof style [2]. While these results are similar in spirit to ours, they do not address how this string $Upd_{X, X \setminus T}$ is incorporated into the new witnesses or how many witnesses must change. It is possible in this model that some elements require very long witness changes, while nearly all other witnesses can remain the same. Our result addresses the separate question of how many witness changes are required.

We note a small gap in the Camacho-Hevia proof (and similarly in the Baldimtsi et al. proof) in the appendix of the full version of this paper. In our proof, we address this issue by defining the notion of a k -good revoked set.

4.2 Vector Commitments

A *vector commitment (VC)* [10] stores a vector $\mathbf{v} = [v_1, \dots, v_k]$ in the form of a succinct digest C . For each index i and corresponding component v_i , the scheme produces a proof π_i that can be used alongside C to verify that $\mathbf{v}_i = v_i$. When a component is changed, the digest and proofs of some or all components may change. Correctness dictates that properly generated proofs of true components verify with their corresponding digests. Security dictates that it is hard to find a proof for an incorrect component. Recently several vector commitment schemes have been constructed with cryptocurrency applications in mind; see [12, 18, 26, 28, 30].

Constructing a Revocable Proof System Using a Vector Commitment Scheme. Our construction commits to a vector storing valid statements. In describing our construction, we use the syntax for a VC scheme from [26]. Let q be an upper bound on the total number of valid statements. Let \perp be some special value used to denote that there is no statement stored at that vector position. The `Setup` function for our RPS calls the `KeyGen` function of the VC scheme with security parameter λ and vector length n (the size of our initial valid set) to obtain public parameters `pp`. The `ComputeState` function, given `pp` and an initial valid set S of size n , first calls the commitment function of the VC scheme on the vector $[s_1, s_2, \dots, s_n, \perp, \dots, \perp]$ (using some arbitrary ordering of S). This outputs a commitment C that is the global state, along with auxiliary information `aux`. To generate the proof w_i for each s_i , `ComputeState` then calls the `Open` function of the VC given i, s_i , and `aux`. It outputs the commitment C and a proof w_i for each s_i . The `Revoke` function of our RPS, given $T \subseteq S$, revokes each statement $s_i \in T$ by setting the corresponding position of the committed vector to \perp . We describe how to do so assuming no batch updates, updating the state and all proofs for each revocation before moving onto the next. More precisely, for each $s_i \in T$, it calls `VC.Update`(C, s_i, \perp, i) to obtain an updated state C' and update information U . It then updates the proof w_j for each other component s_j using `VC.ProofUpdate` given C, w_j, s_i, i, U . After all updates have been made, it outputs all proofs and the resulting commitment. Finally, the `Verify` function of our RPS, given C, s_i, w_i , calls `VC.Ver`(C, s_i, j, w_i) for each vector component j . `Verify` outputs `true` if and only if there exists a j such that `VC.Ver` outputs `true`.

We give an overview of how correctness and security for a VC scheme relate to the corresponding definitions for a revocable proof system; full definitions of correctness and security for a VC scheme are given in [10]. VC schemes offer correctness with overwhelming probability, ensuring that properly generated proofs for committed components verify. See footnote 2 for a discussion of how this compares to correctness with probability 1 for a revocable proof system. The

security definition for a VC guarantees that it's hard for an adversary to find two valid proofs for different values s_i and s'_i of the i^{th} component. This implies security of our constructed revocable proof system: if an adversary finds a proof w^* of a statement s^* that is *not* in the valid set, it has succeeded in finding a proof that the value of the vector at some index i is s^* . Since s^* is not in the valid set, the actual value at i must be \perp or some other s' . The proof w of this other value yields a pair of proofs that verify for different values at index i . Thus, our constructed scheme is an RPS.

4.3 Authenticated Dictionary

A related notion is an *authenticated dictionary* [17, 23], which produces a commitment to a set of key-value pairs, such that proofs of these stored pairs can be generated and verified against the commitment. Throughout time, more key-value pairs can be added, and existing pairs can be modified. When the dictionary is updated, a new shared commitment is generated, potentially invalidating old proofs. The existence of these proofs both for the original data and the updated data corresponds to correctness for a revocable proof system. Security of an authenticated dictionary guarantees that it is difficult to generate proofs of a key-value pair that is not in the stored set. The argument that we can construct a revocable proof system from an authenticated dictionary is along the same lines as the arguments from vector commitments and accumulators. One way to see this is to observe that we can construct a vector commitment scheme using an authenticated dictionary, by storing the vector index-value pairs as key-value pairs in the dictionary. Authenticated dictionaries therefore fit the framework of a revocable proof system, and thus our lower bound holds, implying that proofs must be updated often.

Aardvark [21], a recently proposed distributed authenticated dictionary with applications to stateless blockchains, proposes an interesting versioning scheme to overcome the need to change witnesses enough to accommodate many users making transactions concurrently. We discuss this idea further in Sect. 5.2.

5 Implications for Blockchains

Blockchains typically operate in one of two models: the *unspent transaction output (UTXO)* model or the *account-based* model. A stateless blockchain functions slightly differently in each of these models. We describe the models below and argue that each requires the functionality of a revocable proof system, meaning that our lower bound from Theorem 1 holds.

UTXO Model. In the UTXO model, the global state stores the set of unspent coins. When a user wants to make a transaction, they must specify the coin(s) (UTXOs) they wish to spend and submit a proof that these coins are unspent. A stateless blockchain needs to satisfy correctness: a proof for an unspent coin

should verify against the corresponding global state. If the transaction is successful, the global state is updated, and the spent coins' proofs should no longer verify. In order to prevent users from double spending, it should be computationally hard to produce a proof for a spent coin—this is equivalent to the definition of security for a revocable proof system. A stateless blockchain in the UTXO model is commonly constructed using a dynamic accumulator, where the accumulated set is the set of valid UTXOs. Such accumulators include RSA accumulators [4], Merkle-tree-based accumulators [7, 16], and Verkle trees [6].

Account-Based Model. In the account-based model, the global state stores a list of account-balance pairs. Each account owner, or user, maintains a proof of their account balance. When a user u wants to make a transaction, they submit a proof π that their account-balance pair is included in the global state. The validator verifies the user's account balance using this proof, and they check that the balance is high enough to make the desired transaction. The amount spent is then deducted from the user's balance, and the global state is updated accordingly.

In the context of a revocable proof system, the valid set is the set of account-balance pairs. An account-balance pair is revoked when the corresponding user makes a transaction, changing their account balance. Security ensures that it is hard to generate a proof for an incorrect account-balance pair. Correctness ensures that every user can prove that their true account balance is valid. An account-based blockchain is often constructed using a vector commitment or authenticated dictionary, where each index of the vector represents an account and the value is that account's balance (e.g., [26]).

5.1 Interpreting Our Bound in Practice

An interesting question is exactly what implications Theorem 1 has for practical stateless blockchains. Toward answering this, we graph the number of witness (or proof) changes for various parameter values.

We first apply Theorem 1 to obtain a lower bound on the number of witness changes required after some number m deletions, given an initial valid set of size n . The number of possible deleted sets of size m is $\binom{n}{m} \geq \frac{n^m}{m^m}$ (by, e.g., [13]). Ideally, we would like at least half of these sets to (1) require few ($\leq k$ for some k) witness changes, and (2) allow no deleted elements to be double spent. These are exactly the conditions for a k -good revoked set; thus, in our application of Theorem 1 we can set $|\mathcal{X}_k^*| = \frac{1}{2} \binom{n}{m} \geq \frac{n^m}{2m^m}$. Our next step is to obtain a lower bound for k , showing that many witnesses must change.

Rearranging, we have $\lceil k \lg n \rceil \geq \lg \frac{n^m}{2m^m} - |V|$. Simplifying further,

$$\begin{aligned} k \lg n &\geq m \lg n - m \lg m - |V| - 2 \\ k &\geq m - \frac{|V| + m \lg m}{\lg n} - \frac{2}{\lg n} \end{aligned} \tag{1}$$

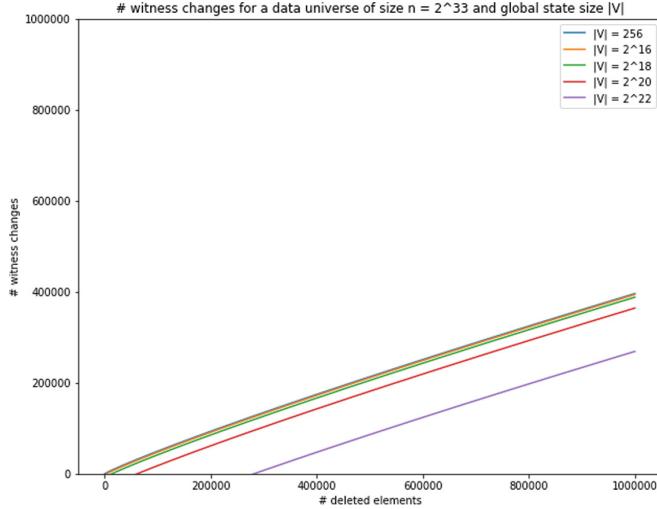


Fig. 1. Number of witness changes $f(m, n, |V|)$ given $0 \leq m \leq 10^6$ deleted elements, a data universe of size $n = 2^{33}$, and varying global state size $|V|$.

Let $f(m, n, |V|)$ denote the right hand side of Eq. 1. We graph f , showing that if at least half of the sets of size m are k -good revoked sets, k must be at least $f(m, n, |V|)$. In our graphs, we use two natural values of n . The first is 2^{33} , approximately the world’s current human population. The second is 2^{26} , approximately the current number of UTXOs in Bitcoin [15]; these graphs are included in the appendix of the full version of this paper.

In Fig. 1, we can see that the relationship between f and m is approximately linear, with the $\frac{m \lg m}{\lg n}$ term having little impact since m is small relative to n in our ranges. Furthermore, increasing the size of the global state V results in a horizontal shift of the curve and has little benefit until it becomes very large.

Like Fig. 1 and Fig. 2 shows that there is not a useful trade-off between the global state size and the number of witness changes per day. The global state size must become very large, at least 2^{22} for most throughput values, before there is much impact on the number of witness changes. This concrete effect mirrors the asymptotic relation of by Corollary 1.

While the number of witness changes may seem small in comparison to the number of UTXOs or accounts in the system, without some additional recovery mechanism, the consequences of a user missing their witness update are severe as they will no longer be able to make transactions. Furthermore, if the system has enough throughput to adequately serve the data universe, there will be many more witness changes: for 24,000 transactions per second (the maximum throughput supported by Visa [29]) the number of witness changes per day for $n = 2^{33}$ becomes roughly 1.25×10^8 . Our graphs show that if most users are not willing to refresh their witnesses continually, hundreds of thousands of these users will lose their coins per day. As a result, most stateless blockchain proposals

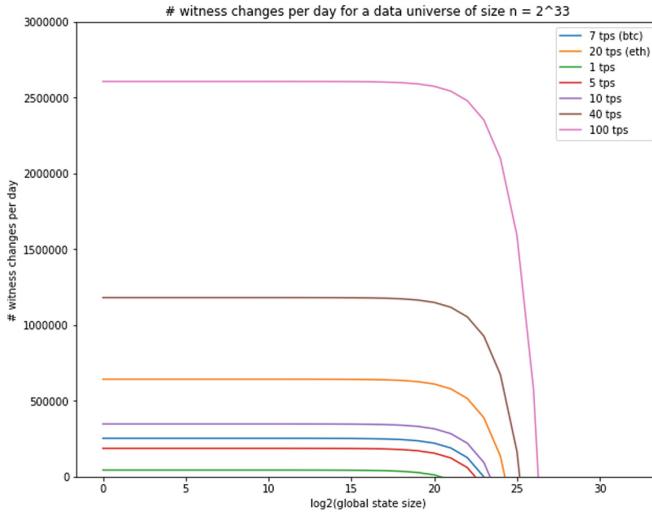


Fig. 2. Number of witness changes per day for a data universe of size $n = 2^{33}$ and varying global state size, for blockchains with various throughput. In particular, Bitcoin and Ethereum support roughly 7 and 20 transactions per second respectively.

have included a way for lazy users to obtain updated proofs, at the cost of more storage for certain parties; the most prominent such solution uses proof-serving nodes (PSNs). Below, we discuss two more limited solutions (a versioning model and a partially persistent model), then conclude with a discussion of PSNs and potential future work relating to them.

5.2 Versioning Model

An issue arises when at some time t , many users simultaneously provide a proof of their account balance (an element in the authenticated dictionary) and a transaction that they wish to make (an update of their element in the dictionary). If the transactions are executed in sequence, each user's transaction requires updating the dictionary, invalidating the other users' proofs. One solution is to store this set of transactions temporarily, so we can verify each user's proof against the global state at time t and then check manually that none of the subsequent transactions changed that user's account balance. We call this a versioning system.

Aardvark [21], an authenticated dictionary designed with cryptocurrency applications in mind, does essentially this: it stores all transactions that happen in the next τ time, for some tunable time parameter τ . The current state commitment at time t is also stored. At a future time up to $t + \tau$, any proof at least as recent as time t can still be verified by checking it against the state commitment at time t , then naively checking that it does not conflict with the cached transactions. This approach essentially ensures that proofs do not need to

change for k transactions by storing k additional state, where k is the number of transactions happening in time τ . This matches our lower bound from Corollary 2 (up to constants), which when translated to this setting says that if we want no proof changes when deleting k elements, we must store at least k state. Thus, this versioning scheme is essentially the best one can hope to achieve without introducing parties such as PSNs storing more state (see, e.g., [24, 26, 28]).

5.3 Partially Persistent Model

A desirable feature of a stateless blockchain is that users know in advance when their proofs will need to change, so they can go online only at that time. Perhaps users could pay a fee for the guarantee that their proofs will remain valid for some number of transactions in the future. A natural question is how much additional state is necessary to accommodate these special requests.

This property is exactly our notion of *persistence*: the persistent set S^* corresponds to the set of proofs that are guaranteed to remain valid. Unfortunately, Corollary 2 says that any secure and correct revocable proof system with a persistent set S^* must have global state size at least $|S^*| - 1$. If any significant portion of the user base wants persistent proofs, the stateless blockchain model does essentially no better than storing the full state.

We can achieve persistence if users are willing to lock up their coins for a set period of time. That is, a user wanting their proof to remain valid for at least a day would sacrifice their ability to spend their coin during that day. We could then separate the blockchain into two state commitments: one state S_1 storing the set of locked coins and another state S_2 storing all other (liquid) coins. Since locked coins can only be spent at the end of the day, S_1 remains the same and no proofs of locked coins change throughout the day. At the end of the day, users may unlock their coins and move them from S_1 to S_2 . We could extend this scheme to support other time ranges, incurring the cost of extra storage as more time ranges are supported.

While potentially helpful in limited settings, this model has serious drawbacks for general use. The most obvious is the fact that users cannot spend their locked coins. Furthermore, the benefits are all-or-nothing in the following way: If a user wants to maintain *any* liquid coins, they must continually update these liquid coins' witnesses, at which point updating their locked coins' witnesses would require minimal additional effort.

5.4 Proof-Serving Node Model

Prior work proposes offloading witness updates to a *proof-serving node* (PSN) [24, 26, 28]. Instead of maintaining its proof itself, a user can delegate this task to a PSN and come online only when it wishes to make a transaction. In any revocable proof system, the PSN can update a user's proof simply by using the *Revoke* algorithm. The storage required for this simple approach scales with the number of users: the PSN can serve k users by storing only these users' proofs and constantly checking for updates. This property that PSNs can use storage

proportional to the number of proofs they maintain somewhat mitigates the centralization issues posed by requiring storing a large state, allowing anyone to operate a small PSN. PSNs also interact nicely with *hybrid nodes*, a newly introduced [19] type of node that stores much less state than full nodes yet can perform nearly all full node functionalities. The PSN model is especially promising in light of our result that there is no holy grail revocable proof system achieving few witness updates on its own.

However, centralization is still a major concern with PSNs, and the PSN model raises interesting questions regarding incentives. PSNs must be incentivized in some way to do this work. Hyperproofs [26] suggests a PSN model where users pay PSNs to maintain their proofs for them. This payment model seems to have an interesting relationship with batch updates, which hyperproofs also allow. That is, while it takes a user time t to update a single proof, a PSN can update the proofs of all n users in the system in time $t \cdot f(n)$ (for some sublinear function f). PSNs that serve enough users to take advantage of batch updates can offer much cheaper prices than small PSNs. There can only be a few PSNs that serve this many users. The resulting system will have a few PSNs storing the full state, and the users they serve will store nothing. This is a significant risk: an adversary that attacks these PSNs can compromise the entire blockchain, preventing many users from spending their coins.

Acknowledgments. This research was conducted primarily at a16z crypto research. Miranda Christ was also supported in part by NSF Award CCF-2107187, by JPMorgan Chase & Co, by LexisNexis Risk Solutions, and by the Algorand Centres of Excellence programme managed by Algorand Foundation. Joseph Bonneau was also supported by NSF Award CNS-1940679 and DARPA Award HR00112020022, and served as a technical advisor to Mina. Any opinions, findings, and conclusions or recommendations expressed in this material are solely those of the authors.

References

1. Abusalah, H., Fuchsbauer, G., Gaži, P., Klein, K.: SNACKs: leveraging proofs of sequential work for blockchain light clients. Cryptology ePrint Archive, Paper 2022/240 (2022)
2. Baldimtsi, F., et al.: Accumulators with applications to anonymity-preserving revocation. In: IEEE Euro S&P (2017)
3. Benaloh, J., de Mare, M.: One-way accumulators: a decentralized alternative to digital signatures. In: Hellseeth, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 274–285. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48285-7_24
4. Boneh, D., Bünz, B., Fisch, B.: Batching techniques for accumulators with applications to IOPs and stateless blockchains. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11692, pp. 561–586. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_20
5. Bonneau, J., Meckler, I., Rao, V., Shapiro, E.: Mina: decentralized cryptocurrency at scale (2020). <https://docs.minaprotocol.com/static/pdf/technicalWhitepaper.pdf>. Accessed 9 Aug 2022
6. Buterin, V.: A state expiry and statelessness roadmap. https://notes.ethereum.org/@vbuterin/verkle_and_state_expiry_proposal

7. Buterin, V.: The stateless client concept (2017). <https://ethresear.ch/t/the-stateless-client-concept/172>
8. Camacho, P., Hevia, A.: On the impossibility of batch update for cryptographic accumulators. In: Abdalla, M., Barreto, P.S.L.M. (eds.) LATINCRYPT 2010. LNCS, vol. 6212, pp. 178–188. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14712-8_11
9. Camenisch, J., Lysyanskaya, A.: Dynamic accumulators and application to efficient revocation of anonymous credentials. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 61–76. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45708-9_5
10. Catalano, D., Fiore, D.: Vector commitments and their applications. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 55–72. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36362-7_5
11. Chen, W., Chiesa, A., Dauterman, E., Ward, N.P.: Reducing participation costs via incremental verification for ledger systems. Cryptology ePrint Archive, Paper 2020/1522 (2020). <https://eprint.iacr.org/2020/1522>
12. Chepurnoy, A., Papamanthou, C., Srinivasan, S., Zhang, Y.: EDRAx: a cryptocurrency with stateless transaction validation. Cryptology ePrint Archive, Paper 2018/968 (2018)
13. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to Algorithms. MIT Press, Cambridge (2022)
14. Dai, W.: b-money (1998). <https://www.weidai.com/bmoney.txt>
15. Delgado-Segura, S., Pérez-Solà, C., Navarro-Arribas, G., Herrera-Joancomartí, J.: Analysis of the Bitcoin UTXO set. In: Zohar, A., et al. (eds.) FC 2018. LNCS, vol. 10958, pp. 78–91. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-662-58820-8_6
16. Dryja, T.: Utreexo: a dynamic hash-based accumulator optimized for the Bitcoin UTXO set. Cryptology ePrint Archive, Paper 2019/611 (2019)
17. Goodrich, M.T., Shin, M., Tamassia, R., Winsborough, W.H.: Authenticated dictionaries for fresh attribute credentials. In: Nixon, P., Terzis, S. (eds.) iTrust 2003. LNCS, vol. 2692, pp. 332–347. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-44875-6_24
18. Gorbunov, S., Reyzin, L., Wee, H., Zhang, Z.: Pointproofs: aggregating proofs for multiple vector commitments. In: ACCM CCS (2020)
19. Hegde, P., Streit, R., Georghiades, Y., Ganesh, C., Vishwanath, S.: Achieving almost all blockchain functionalities with polylogarithmic storage. In: Eyal, I., Garay, J. (eds.) Financial Cryptography and Data Security: 26th International Conference, FC 2022, Grenada, 2–6 May 2022, Revised Selected Papers, pp. 642–660. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-18283-9_32
20. Kassis, A., Bonneau, J.: Proof of necessary work: succinct state verification with fairness guarantees. In: Financial Crypto (2023). <https://eprint.iacr.org/2020/190.pdf>
21. Leung, D., Gilad, Y., Gorbunov, S., Reyzin, L., Zeldovich, N.: Aardvark: an asynchronous authenticated dictionary with applications to account-based cryptocurrencies. In: USENIX Security (2022)
22. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2008)
23. Papamanthou, C., Tamassia, R.: Time and space efficient algorithms for two-party authenticated data structures. In: Qing, S., Imai, H., Wang, G. (eds.) ICICS 2007. LNCS, vol. 4861, pp. 1–15. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-77048-0_1

24. Reyzin, L., Meshkov, D., Chepurnoy, A., Ivanov, S.: Improving authenticated dynamic dictionaries, with applications to cryptocurrencies. Cryptology ePrint Archive, Paper 2016/994 (2016)
25. Shannon, C.E.: A mathematical theory of communication. Bell Syst. Tech. J. **27**(3), 379–423 (1948)
26. Srinivasan, S., Chepurnoy, A., Papamanthou, C., Tomescu, A., Zhang, Y.: Hyper-proofs: aggregating and maintaining proofs in vector commitments. IACR Cryptol. ePrint Arch. **2021**, 599 (2021)
27. Todd, P.: Making UTXO Set Growth Irrelevant with Low-Latency Delayed TXO Commitments (2016). <https://petertodd.org/2016/delayed-txo-commitments>
28. Tomescu, A., Abraham, I., Buterin, V., Drake, J., Feist, D., Khovratovich, D.: Aggregatable subvector commitments for stateless cryptocurrencies. Cryptology ePrint Archive, Paper 2020/527 (2020)
29. Visa: Visa acceptance for retailers. <https://usa.visa.com/run-your-business/small-business-tools/retail.html>
30. Wang, W., Ulichney, A., Papamanthou, C.: BalanceProofs: maintainable vector commitments with fast aggregation. Cryptology ePrint Archive, Paper 2022/864 (2022)
31. Wood, G., et al.: Ethereum: a secure decentralised generalised transaction ledger (2014)