

# Mutator Sets and their Application to Scalable Privacy

Alan Szepieniec      Thorkil Værge  
alan@neptune.cash   thor@neptune.cash

Neptune

**Abstract.** A mutator set is a cryptographic data structure for authenticating operations on a changing set of data elements called *items*. Informally:

- There is a short commitment to the set.
- There are succinct membership proofs for elements of the set.
- It is possible to update the commitment as well as the membership proofs with minimal effort as new items are added to the set or as existing items are removed from it.
- Items cannot be removed before they were added.
- It is difficult to link an item’s addition to the set to its removal from the set, except when using information available only to the party that generated it.

This paper formally defines the notion, motivates its existence with an application to scalable privacy in the context of cryptocurrencies, and proposes an instantiation inspired by Merkle mountain ranges and Bloom filters.

## 1 Introduction

Blockchain-based money has the potential to transform and disrupt the financial industry and even the global monetary system. However, its wider adoption and success are obstructed by two critical challenges: *privacy* and *scalability*. Privacy is the quality of a blockchain protocol that thwarts attempts to trace the source and destination of payments based on information present in the public ledger. Scalability is the quality of a blockchain protocol that resists feature degradation as a result of increased use or age. Accordingly, much research has been devoted to cryptographic techniques that promise to achieve one or both of these features.

*Cryptographic accumulation schemes* compress a large set of items into a compact commitment in order to dramatically reduce the complexity of elementary set operations at the modest cost of sacrificing perfect security for computational security. They have applications in timestamping, blacklists, public key infrastructure, anonymous credentials, group signatures, to name examples without being exhaustive. All these applications have one thing in common: they enable a lightweight agent to verify with cryptographic certainty the correct computation of elementary operations on relatively large sets.

Accumulation schemes have seen renewed interest in recent years in the context of blockchains, due to their potential to improve scalability by reducing the cost of participation. The lower barrier to entry and lower running costs incentivize more participation, benefiting the protocol's decentralization metrics and the features that are downstream from there such as resilience against legal attacks.

While accumulation schemes are extremely well suited to reduce the workload of verifying traceable ledgers, they are incapable of the same task when it comes to untraceable ledgers. This limitation comes from the static nature of the items accumulated by an accumulation scheme. Whenever the same item is added and later removed, or added and later modified, or modified once and later modified again, it is apparent to all observers that the items in question are one and the same. A basic requirement towards addressing privacy would stipulate items be hidden behind cryptographic commitments, and that these commitments *mutate* between operations such that external observers cannot link them.

However, up until now blockchains have relied on entirely different techniques from accumulation schemes to achieve privacy, namely decoys or mixnets. In either case, the transaction inputs and outputs are hidden behind cryptographic commitments that hide the amount, and every transaction comes with a zero-knowledge proof that establishes that no inflation is taking place and that all amounts are positive.

*Decoys* are plausible transaction inputs listed alongside the true origin inputs, which is where the funds that are being spent actually come from. The external observer cannot distinguish between decoys and true origins. In order to catch double-spending attempts, every transaction input that is being spent unlocks a short string of data called the *nullifier*, and a zero-knowledge proof hides which of the inputs to the transaction it came from. The blockchain maintains a set of nullifiers and whenever a transaction is confirmed, the nullifiers announced by it are added to the set. A transaction is only valid if its nullifiers are not already in the set.

The downside of this decoy-and-nullifier approach is its poor scalability: this nullifier set grows with the number of transactions in the history of the network. In a naïve implementation, the verifying protocol participant stores the entire set and incurs the associated storage cost. A more sophisticated implementation uses an accumulation scheme that admits proofs of non-membership. It shifts the burden of work but does not eliminate it: the *verifying protocol participant* only needs to store a compact set commitment and verify succinct non-membership proofs, but the *transaction initiator* must produce a non-membership proof relative to this set, which integrates knowledge of all nullifiers. An even more sophisticated implementation fixes a snapshot of the set commitment at the time the coin was generated. To spend it, the transaction initiator must prove that the nullifier was not added to the nullifier set since it was created. The user must therefore remain online to continually update his non-membership proofs, or else synchronize them upon rejoining. The complexity of this latter synchronization task scales linearly with the intervening time because the updated proofs must

depend on all the nullifiers announced in the time spent offline. To adapt this approach for better scalability, either the nullifier set must be prunable somehow or else its *non-membership proofs should be practically independent of members*.

*Mixnets* are an alternative to the decoy-and-nullifier approach. In this construction, unspent transaction outputs (UTXOs) are hidden behind *publicly re-randomizable* commitments, which enable *anyone* to transform an old commitment into a new one that *a)* binds to the same information, and *b)* can be opened only by the same parties that could open the old one, but *c)* is unlinkable to the old commitment in the eyes of everyone else. Depending on the blockchain architecture, the entire UTXO-commitment set or a subset thereof is mixed at regular or variable intervals by miners, transaction initiators, or designated operators that are neither.

The first downside of this approach is that the owner of a UTXO must learn his new commitment by iterating over all mix outputs and either trying to unlock them or running some other ownership test on them. The workload of this task is linear in the size of the batch and thus proportional to the size of the anonymity set. Even if the UTXO-commitment passes through multiple smaller mixes instead, the workload merely shifts from tracing the UTXO across one large mix to tracing it across multiple smaller ones. While the anonymity set size does expand exponentially in the number of mixes, the user's workload scales linearly in this number, and thus once again either requires users to remain online or face a costly catch-up task upon rejoining.

The second downside of this mixnet approach is the fact that the operator selects the secret permutation and can be coerced or bribed to release it, or to select one that is not uniformly random. A construction for eliminating the operator's randomness would obviate this trust assumption and moreover make the mix outputs deterministic, thus simplifying the UTXO owner's tracing task.

In summary, the existing constructions for scalability and for privacy seem to induce tradeoffs that preclude achieving both *scalability and privacy* simultaneously. Specifically:

- accumulation schemes reduce the cost of participation but are inherently traceable;
- decoys and nullifiers generate untraceable transactions but require operations on an ever-growing nullifier set that cannot be pruned or truncated;
- mixnets generate untraceable transactions but require users to trace their coins across mixes and moreover require trust in the mixnet operator.

Mutator sets are cryptographically authenticated data structures that fix these deficiencies. Informally, mutator sets enable users to add items to a set and remove them later. The essential differences relative to accumulation schemes are that mutator sets accumulate *commitments* to data items rather than items directly; and that the commitment when an item is added is unlinkable to its commitment when it is removed. Phrased crudely, *mutator sets are accumulation schemes with addition-removal unlinkability*.

Mutator sets are not just idly reminiscent of nullifier sets and decoy commitments. What the full list of all historical commitments to transaction outputs and

the full nullifier set jointly represent, is a set of *unspent* transaction outputs. A transaction is valid only if all its inputs are members of this UTXO set: the nullifier set keeps track of which items are removed from this set. As the commitment and nullifier are unlinkable, this construction provides addition-removal unlinkability. The relevant difference relative to mutator sets is the workload. Phrased crudely, *mutator sets are succinct decoy-and-nullifier sets*.

As for mixnets, the similarity is readily apparent: the items are removed in a (possibly) different order from the one in which they were added, and transformed untraceably in the process. The main difference is that there is no operator. Consequently, there is no leakable permutation or re-randomization randomness. Phrased crudely, *mutator sets are operator-free mixnets*.

CONSTRUCTION OVERVIEW. Our construction uses Merkle mountain ranges (MMRs) [38,16], which are sequences of Merkle trees that support adding elements to a list, in addition to inheriting the basic features of Merkle trees such as compact membership proofs and support for modifying leaves. MMRs are used to compactify two data structures, the *append-only commitment list (AOCL)* and the *sliding-window Bloom filter (SWBF)*.

The AOCL tracks additions to the mutator set by recording hiding commitments to data items. These commitments are known as *addition records*.

The SWBF tracks removals from the mutator set by recording indices of bits that are announced when items are removed from the mutator set. These indices are uniquely determined by the data item, but *which* indices is impossible for the computationally bounded external observer to predict. A zero-knowledge proof certifies that a given set of bits corresponds to *some* addition record that lives in the AOCL. A *removal record* is a set of bit indices along with this zero-knowledge proof and along with the bits' membership proofs in the SWBF MMR. The removal record is valid only if some of the listed bits are not set yet.

The SWBF has infinite size in principle, but every addition record can only sample bits from a finite window. This window slides periodically.

Figure 1 captures key elements visually:

- the commitment to the set consists of the Merkle roots of both MMRs and the active window;
- an addition record consists of a commitment to the data item;
- the zero-knowledge proof establishes that for some addition record, *a*) the addition record lives in the AOCL MMR, and *b*) the given bit indices were derived correctly;
- a removal record consists of a set of announced bits, the zero-knowledge proof, and membership proofs in the SWBF MMR for all bits that no longer live in the active window.

The effect of the sliding window is to reduce the dependency of membership proofs and removal records for items that were added early on, on the removal records of items added later. Specifically, the MMR membership proofs for old chunks of the SWBF require only a logarithmic number of updates as a function of the number of items that were added (and maybe removed) after the item in

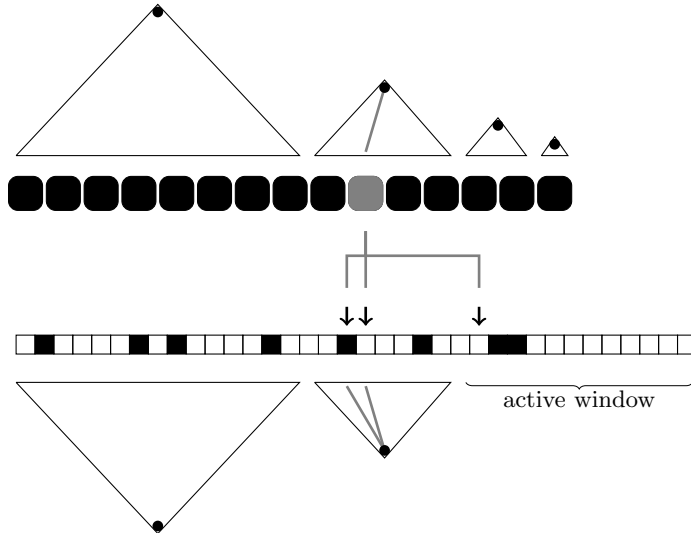


Fig. 1: Schematic overview of the present construction.

question was. Under the realistic assumption that at any point in time freshly added items are more likely to be removed than old ones, this property translates to a workload for updating membership proofs that decreases over time.

**RELATED WORK.** *Accumulators.* Arguably Merkle trees [28] are the first accumulation schemes and when their leaves are sorted, they even admit non-membership proofs. Benaloh and de Mare [6] introduce the term “accumulator” and describe a construction where the order of accumulated items does not matter. Nyberg [33] proposed a similar function based on symmetric primitives. Barić and Pfitzmann [4] argue that accumulators should be presented as tuples of algorithms, extending the term to accumulator *schemes*. Since then research has focused on finding new building blocks for instantiating accumulator schemes [32,26,27,1], extending their features and properties [10,13], and applying them to various protocols and tasks [34,23].

The combined use of Bloom filters and Merkle trees to construct cryptographic accumulator schemes was first proposed by Rambaja and Avdullahu [35]. Compared to just Merkle trees, the advantage is non-membership proofs without sacrificing updatability; and compared to just Bloom filters, the advantage is compactness. This construction does not provide addition-removal unlinkability.

*Accumulators with Privacy.* After finishing the bulk of this paper, the authors became aware of two papers extending accumulator schemes in such a way as to achieve a notion very similar to the one presented here. The first paper [2] presents a generic framework for building accumulator schemes with fancy properties from accumulator schemes without. It then proceeds to construct one from two additive RSA-based accumulators, one with membership proofs and one with

non-membership proofs. The new construction features *join-revoke unlinkability* — which is essentially the same property as addition-removal unlinkability of the present paper.

The second paper [3] generically constructs an *oblivious* accumulator scheme, which is one with add-delete unlinkability (also essentially the same property) and add-delete indistinguishability (additionally hiding the set size). It builds this from a key-value commitment scheme, and the key-value commitment scheme from a universal accumulator and an extendable-length vector commitment scheme. What is missing relative to mutator sets is a polylogarithmic bound on the running times of *all* algorithms.

*Light Clients.* Starting with informal discussions [30,18], there have been a host of proposals for reducing the cost of participating in a blockchain network by shrinking the state of a client to a short cryptographic commitment [16,8]. A related objective is to shrink the blockchain itself by having miners produce succinct proofs of history verification, in addition to proof-of-work [24,25,22].

*Privacy on Blockchains.* Zerocoin [29] first proposed the use of accumulators to ensure privacy in cryptocurrencies. ZCash [5] was the first to use zk-SNARKs with the same objective. Monero uses ring signatures, a related tool to accumulators and zero-knowledge proofs, to achieve privacy [36]. All three fall in the decoy-and-nullifier category.

Mimblewimble [21,20,19] first constructed private UTXOs without nullifier sets. Unfortunately, this approach limits privacy to *amount confidentiality* (or *state confidentiality*) as UTXOs remain linkable across transactions.<sup>1</sup> Quisquis [17] first proposed to use mixnets to effect unlinkability without nullifier sets.

*Mixnets.* Mixnets were originally proposed by Chaum [11] for anonymous messaging. Since then they have been used as a building block for a host of privacy-enhancing technologies, including anonymous remailers [12], onion routing [15], electronic voting [37], etc.

ROADMAP. **Section 2** covers some basic preliminaries, and after that **Section 3** presents formal definitions of both accumulation schemes and mutator sets. **Section 4** details how a mutator set can be used in the context of a UTXO-based blockchain to achieve scalable privacy. **Section 5** presents our construction in a sequence of incremental steps, and proves its security properties. Finally, **Section 6** finishes with some concluding thoughts.

## 2 Preliminaries

**Hash Function Security.** We reduce the security of our construction to security of some concrete hash function  $H : \Sigma^* \rightarrow D$  (the particular choice is irrelevant). To capture this notion we use two insecurity functions: 1)  $\text{InSec}_{\text{coll}}^H$ , the insecurity function for collision resistance for  $H$ ; and 2)  $\text{Adv}_{\text{PRF}}^H$ , an adversary’s

<sup>1</sup> Mimblewimble blockchains notably do not record individual transactions but only aggregates of them. Nevertheless, the individual transactions can be collected by the online but otherwise passive observer as they are broadcasted.

advantage over a random guess at distinguishing  $H$  from a random function. In order to justify conflating a concrete function  $H$  with a family of functions (as the PRF game requires) we implicitly designate part of the input as key material and pretend as though the adversary cannot read it. These quantities are defined as

$$\text{InSec}_{\text{coll}}^H \triangleq \max_{\mathbf{A}} \Pr [\mathbf{A}() \Rightarrow a, b \in \Sigma^* \wedge a \neq b \wedge H(a) = H(b)] \quad (1)$$

$$\text{Adv}_{\text{PRF}}^H \triangleq \max_{\mathbf{A}} \left| \Pr [\mathbf{A}^{f_b}() \Rightarrow b] - \frac{1}{2} \right|, \quad (2)$$

where  $f_0 = H, f_1 \stackrel{\$}{\leftarrow} \{f \mid f : \Sigma^* \rightarrow \mathbb{D}\}, b \stackrel{\$}{\leftarrow} \{0, 1\}$ .

The omitted arguments of  $\text{InSec}$  and  $\text{Adv}$  can be the security parameter  $\lambda$  (whose unary representation is  $\lambda$ ), the allowance for adversaries with respect to running time, memory, or (in a black box model) number of queries. The maximum is taken over all adversaries  $\mathbf{A}$  satisfying those constraints.

**Commitment Scheme.** A commitment scheme binds a party to a data item whose value can be revealed in a way that prevents equivocation. Specifically, a commitment scheme consists of two algorithms, presented here abstractly:

- **commit** :  $T \times \mathbb{R} \rightarrow \mathbb{C}$  takes a data item and possibly some randomness and outputs a *commitment*.
- **open** :  $T \times \mathbb{C} \times \mathbb{O} \rightarrow \{\text{True}, \text{False}\}$  takes a data item, the commitment, possibly some auxiliary opening information, and outputs **True** if the item was used to produce the commitment and **False** otherwise.

A commitment scheme generally satisfies two security properties, *hiding* and *binding*, although the former requirement is sometimes dropped. Hiding is equivalent to semantic security: no resource-bounded adversary should be able to compute any function of the committed message, and therefore hiding commitment functions require a randomness parameter. Binding captures non-equivocation: no adversary should be able to produce a commitment that it can open to distinct messages.

A commitment scheme (**commit**, **open**) is *hiding* if for all the polynomial-time adversaries  $\mathbf{A} = (\mathbf{A}_0, \mathbf{A}_1)$  the advantage over a random guess in the semantic security game is negligible:

$$\text{Adv}_{\text{ss}}^{\mathbf{A}} \triangleq \left| \Pr [\mathbf{A}_1(st, c) \Rightarrow b] - \frac{1}{2} \right| \leq \epsilon(\lambda), \quad (3)$$

where  $m_0, m_1 \leftarrow \mathbf{A}_0(\lambda); b \stackrel{\$}{\leftarrow} \{0, 1\}; r \stackrel{\$}{\leftarrow} \mathbb{R}; c \leftarrow \text{commit}(m_b; r)$ .

A commitment scheme (**commit**, **open**) is *binding* if for all polynomial-time adversaries  $\mathbf{A}$  the binding insecurity  $\text{InSec}_{\text{binding}}^{\mathbf{A}}$  is negligible:

$$\text{InSec}_{\text{binding}}^{\mathbf{A}} \triangleq \Pr [m_0 \neq m_1 \wedge \text{open}(m_0, c, o_0) = \text{open}(m_1, c, o_1)] , \quad (4)$$

where  $(m_0, o_0, m_1, o_1, c) \leftarrow \mathbf{A}(\lambda)$ .

A common commitment scheme is defined relative to a hash function  $H : \Sigma^* \rightarrow \mathbf{D}$  and known as the *canonical* commitment scheme. The commitment algorithm is defined as  $\text{commit} : T \times \mathbf{R} \rightarrow \mathbf{D}, (m, r) \mapsto H(m\|r)$ . In order to be hiding, the space of randomness  $\mathbf{R}$  must be at least as large as the space of digests  $\mathbf{D}$ . The binding property of the canonical commitment scheme reduces to the collision-resistance of the hash function.

The canonical opening algorithm is a special case of the *standard* opening algorithm, which is defined relative to a generic commitment algorithm  $\text{commit} : T \times \mathbf{R} \rightarrow \mathbf{C}$  as  $\text{open} : T \times \mathbf{C} \times \mathbf{R} \rightarrow \{\text{True}, \text{False}\}, (t, c, r) \mapsto \text{commit}(t, r) \stackrel{?}{=} c$ .

**Merkle Tree.** A Merkle tree [28] is a cryptographic data structure that compactly commits to a vector of  $2^k$  items called *leaves*. This commitment, called the *root*, is derived by repeatedly hashing together two neighboring nodes in every layer of a balanced binary tree of *height*  $k$ , whose bottom layer consist of the leaves. The list of siblings of nodes on a path from a leaf to the root is known as that leaf’s *authentication path* and can be used to verify the claim that the leaf belongs to the tree. We present only the interface here.

- $\text{root} : [T] \rightarrow \mathbf{D}$  computes the root from the leaves.
- $\text{open} : [T] \times \mathbb{N} \rightarrow [\mathbf{D}]$  computes the authentication path for an indicated leaf, from the list of all leaves and a leaf index.
- $\text{verify} : \mathbf{D} \times \mathbb{N} \times [\mathbf{D}] \times T \rightarrow \{\text{True}, \text{False}\}$  verifies that a leaf belongs to a Merkle tree, given the tree’s root, the leaf’s index, its authentication path, and the leaf itself.

Additionally, we define two functions that are not traditionally associated with Merkle trees but nevertheless represent relatively trivial extensions.

- $\text{modify} : \mathbf{D} \times \mathbb{N} \times [\mathbf{D}] \times T \times T \rightarrow \mathbf{D} \cup \{\perp\}$  modifies an indicated leaf, assuming its membership proof is valid. Specifically, it sends  $(\text{old\_root}, \text{leaf\_index}, \text{authentication\_path}, \text{old\_item}, \text{new\_item})$  to  $\text{new\_root}$  if the authentication path is valid, or to  $\perp$  if it is not. Note that afterwards the same authentication path should be valid for  $\text{new\_item}$  but relative to  $\text{new\_root}$ .
- $\text{update} : T \times \mathbb{N} \times [\mathbf{D}] \times (\mathbf{D} \times \mathbb{N} \times [\mathbf{D}] \times T \times T) \rightarrow [\mathbf{D}] \cup \{\perp\}$  updates the authentication path for a given leaf, given all of the arguments for modifying another leaf.

**Bloom Filter.** A Bloom filter [7] is a probabilistic data structure for cheap probabilistic set membership tests. While the test can generate a false positive, it cannot produce false negatives. In other words, the result is either “the item might be in the set” or “the item is definitely not in the set”. The cost of running the test only to get an undefinitive positive answer back, is compensated for by the definite negative alternative result, which allows the programmer to short-circuit whatever procedure comes next.



The state of a Bloom filter is defined by an array of  $w$  bits, which is initially all zero. Whenever an item  $t \in T$  is added to the set,  $k$  indices in  $[0, \dots, w - 1]$  are calculated from  $k$  independent hash functions  $[H_0, \dots, H_{k-1}]$  applied to  $t$ . The indicated bits are set.

To test an item  $t \in T$  for membership, the same  $k$  hash functions are evaluated. If all indicated bits are set, the item might be in the set. But if some indicated bit is not set, then the item cannot have been added to the set, and as a result the item is definitely not in the set.

### 3 Definition

#### 3.1 Accumulation Scheme

The following definition of *accumulation* schemes is different from formal definitions of *accumulator* schemes existing in the literature [9]. The difference is motivated by the need for scaffolding to exposit the main construction and prove its security. Regardless, the notions are the same in spirit.

**Definition 1 (accumulation scheme).** *An accumulation scheme is a tuple of polynomial time algorithms relative to a item data type  $T$ .*

- $\text{init} : [\{1\}] \rightarrow K$  generates a commitment to the empty set.
- $\text{prove} : K \times T \rightarrow M$  takes a set commitment and a new item, and outputs a membership proof that will be valid under the updated commitment if the item is added, but invalid if is removed thereafter.
- $\text{verify} : K \times T \times M \rightarrow \{\text{True}, \text{False}\}$  verifies a membership proof.
- $\text{add} : K \times T \rightarrow K \times A$  updates the set commitment so that the represented set includes the given item, and additionally outputs an addition record.
- $\text{remove} : K \times T \times M \rightarrow K \times R$  updates the set commitment so that the set no longer includes the given item, and additionally outputs a removal record.
- $\text{update} : K \times T \times M \times (A \vee R) \rightarrow M$  updates an item's membership proof given an addition or removal record.

The main difference relative to existing definitions of *accumulator* schemes comes from the explicit treatment of addition and removal records. Intuitively, records are the pieces of information that are associated with a change to the set commitment. As such, they are also needed to update a membership proof so that it is synchronized to the new set commitment. The addition and removal records expose a feature (or deficiency) of accumulation schemes: in all constructions the authors are aware of, the addition records and removal records pertaining to the same item are trivially linkable.

Another difference is that  $\text{prove}$  can only be invoked before the item in question was added. The scope of this paper does not cover generating proofs for items that were added in the past; it covers maintaining proofs and synchronizing them when they are out of date.

Some constructions enable modification of existing items. Strictly speaking, modification is implied by removing and adding. However, for practical use it is worthwhile to define the interface explicitly:

- **modify** :  $K \times M \times T \times T \rightarrow (K \times E) \cup \{\perp\}$  updates the set commitment so that the given item is replaced with the second given item, if the membership proof is valid, and if so outputs a *modification record* (“E” for edit).

When abstracting the type of record and intend only to capture some piece of data inducing an update to the set commitment, we use the term *update record* and reserve the symbol  $v \in A \cup R \cup E$ .

### 3.2 Security

An accumulation scheme needs to have the following security properties, defined in asymptotic terms relative to a security parameter  $\lambda$ . Additionally, the security games are defined with respect to a sequence  $O = (o_i)_{i=1}^{\#O}$  of valid set operations  $o_i : \{\text{add}, \text{remove}\} \times T$ , corresponding to an underlying set<sup>2</sup>  $S$  recursively defined as  $S(O \parallel (\text{add}, t)) = S(O) \cup \{t\}$ ,  $S(O \parallel (\text{remove}, t)) = S(O) \setminus \{t\}$ , and  $S(\emptyset) = \emptyset$ . Specifically, a sequence of operations  $O = (o_i)_{i=1}^{\#O}$  is *valid* iff  $\forall i \in \{1, \dots, \#O\}. o_i = (\text{remove}, t) \Rightarrow t \in S((o_j)_{j=1}^{i-1})$ .

**Definition 2 (completeness for accumulation schemes).** *Consider the completeness game  $\text{Game}_{\text{cmplt}}^{\text{AS}}$ .*

*An accumulation scheme is complete iff, for any valid sequence of operations  $O_1 \parallel (\text{add}, t) \parallel O_2$ , the completeness error  $\text{Err}_{\text{cmplt}}$  is negligible:*

$$\text{Err}_{\text{cmplt}}(\lambda) \triangleq 1 - \Pr [\text{Game}_{\text{cmplt}}(\lambda, O_1, O_2, t) \rightarrow \text{True}] \leq \epsilon(\lambda) . \quad (5)$$

It is tricky to define soundness for a proof system for set membership, when the set is determined by a compact commitment. Indeed: a single commitment may decommit to multiple sets, in which case soundness is a moot notion, and *knowledge-soundness* is preferred. We avoid definitional complexity by assuming instead that the commitment is computed honestly. The adversary can only provide the items to be added to the set and the randomnesses to be used when computing these additions. This definition of soundness fails in contexts where the adversary is capable of computing the commitment dishonestly. However, this dishonest calculation will be exposed at a later stage through the use of general-purpose zero-knowledge proofs.

**Definition 3 (honest-commitment soundness).** *Consider the honest-commitment soundness game  $\text{Game}_{\text{hcsnd}}$ .*

*An accumulation is honest-commitment sound iff the honest-commitment soundness error is negligible, where the maximum is taken over all polynomial-time adversaries  $A$ .*

$$\text{Err}_{\text{hcsnd}}(\lambda) \triangleq \max_A \Pr [\text{Game}_{\text{hcsnd}}^A(\lambda)] \leq \epsilon(\lambda) . \quad (6)$$

<sup>2</sup> Throughout this paper we use the word “set” to denote multiset, or unordered list.

<b>Game 1: Completeness</b>	<b>Game 2: Honest-Commitment Soundness</b>
<pre> 1 <b>define</b> Game<sub>cmplt</sub><sup>AS</sup>(<math>\lambda, O_1, O_2, t</math>)   <b>as:</b> 2   <math>\kappa \leftarrow \text{init}(\lambda)</math> 3   <b>for</b> <math>o \in O_1</math> : 4     <b>if</b> <math>o = (\text{add}, s)</math> : 5       <math>\kappa, \mu^*, - \leftarrow \text{add}(\kappa, s)</math> 6     <b>if</b> <math>o = (\text{remove}, s)</math> : 7       <math>\kappa, - \leftarrow</math> 8         <math>\text{remove}(\kappa, s, \mu^*)</math> 9   <math>\kappa, \mu, \alpha \leftarrow \text{add}(\kappa, t)</math> 10  <b>for</b> <math>o \in O_2</math> : 11    <b>if</b> <math>o = (\text{add}, s)</math> : 12      <math>\kappa, \mu^*, \alpha^* \leftarrow</math> 13        <math>\text{add}(\kappa, s)</math> 14      <math>\mu \leftarrow \text{update}(t, \mu, \alpha^*)</math> 15    <b>if</b> <math>o = (\text{remove}, s)</math> : 16      <math>\kappa, \rho^* \leftarrow</math> 17        <math>\text{remove}(\kappa, s, \mu^*)</math> 18      <math>\mu \leftarrow \text{update}(t, \mu, \rho^*)</math> 19  <b>return</b> <math>\text{verify}(\kappa, t, \mu)</math> </pre>	<pre> 1 <b>define</b> Game<sub>hcsnd</sub><sup>A</sup>(<math>\lambda</math>) <b>as:</b> 2   <math>\kappa \leftarrow \text{init}(\lambda)</math> 3   <math>\triangleright</math> <i>adversary supplies:</i> 4     -set operations <math>O</math>, 5     -randomnesses <math>\mathcal{R}</math>, 6     -item <math>t</math>, 7     -membership proof <math>\mu</math> 8   <math>O, \mathcal{R}, t, \mu \leftarrow \text{A}(\lambda)</math> ; 9   <b>for</b> <math>(o, r) \in \text{zip}(O, \mathcal{R})</math> : 10    <b>if</b> <math>o = (\text{add}, s)</math> : 11      <math>\kappa, -, - \leftarrow \text{add}(\kappa, s; r)</math> 12    <b>if</b> <math>o = (\text{remove}, s)</math> : 13      <math>\kappa, - \leftarrow</math> 14        <math>\text{remove}(\kappa, s; r)</math> 15  <b>return</b> 16    <math>\text{verify}(\kappa, t, \mu) \wedge t \notin S(O)</math> </pre>

### 3.3 Mutator Set

Informally, a mutator set is an accumulation scheme that additionally satisfies *addition-removal unlinkability*, which asserts that it is difficult to relate addition records to removal records. However, the set operations for accumulation schemes are deterministic, giving rise to a trivial attack against unlinkability whereby the adversary runs through all possible permutations of the given set operations to see which one matches with his view. In order to achieve unlinkability, the set operations must involve some randomness that the adversary does not have access to. Mutator sets provide this interface.

To define the accumulator set interface we start from accumulation schemes and modify this interface as follows:

- Users must *commit* to items  $t \in T$  using secret randomness  $r \xleftarrow{\$} \mathbb{R}$  to generate an addition record  $\alpha$  before they can be added. This function is achieved by the method  $\text{commit} : T \times \mathbb{R} \rightarrow A$ , where  $\mathbb{R}$  is a sufficiently large alphabet of random symbols. The addition record  $\alpha \in A$  is a commitment to the item  $t$  (not to be confused with the commitment  $\kappa$  to the set).
- The public input to the **add** method is the addition record  $\alpha \in A$ , instead of the randomness-free item  $t \in T$ . (Note that in contrast to **add** for accumulation schemes,  $\alpha$  is an input not an output.)

- To remove an item  $t \in T$  from the set, the user must first generate a removal record  $\rho \in R$ , which observers can then use to update the set commitment or membership proofs with. This separation of concerns is captured by to methods, **drop** and **remove**.
    - **drop** :  $K \times T \times M \rightarrow R$  generates a removal record from a set commitment, an item, and its membership proof.
    - **remove** :  $K \times R \rightarrow K \cup \{\perp\}$  updates the set commitment with a valid removal record.
- (Note: symmetrically to  $\alpha$ ,  $\rho$  has moved from the output of **remove** in an accumulation scheme to the input of **remove** in a mutator set.)

**Definition 4 (mutator set).** A mutator set is a tuple of polynomial time algorithms relative to a item data type  $T$ :

- **init** :  $[\{1\}] \rightarrow K$  generates a commitment to the empty set.
- **commit** :  $T \times R \rightarrow A$  generates an addition record by using the supplied randomness to commit to an item.
- **prove** :  $K \times T \times R \rightarrow M$  takes a set commitment, a new item, and randomness used to commit to it, and outputs a membership proof that will be valid under the updated commitment if the item is added, but invalid if is removed thereafter.
- **verify** :  $K \times T \times M \rightarrow \{\text{True}, \text{False}\}$  verifies a membership proof.
- **add** :  $K \times A \rightarrow K$  updates the set commitment so that the set includes the item committed to by the addition record.
- **drop** :  $K \times T \times M \rightarrow R$  generates a removal record from a set commitment, an item, and its membership proof.
- **remove** :  $K \times R \rightarrow K \cup \{\perp\}$  updates the set commitment so that the set no longer includes the item determined by the removal record, assuming it is a valid removal record.
- **update** :  $K \times T \times M \times (A \vee R) \rightarrow M$  updates an item's membership proof given an addition or removal record for another item, such that the updated membership proof is valid for the new set commitment resulting from applying the record.

### 3.4 Security

Completeness for mutator sets is defined analogously to **Definition 2**: any sequence of valid set operations should give rise to valid membership proofs. **Definition 5** adapts this intuition to the updated interface.

**Definition 5 (completeness for mutator sets).** Consider the completeness game  $\text{Game}_{\text{cmplt}}^{\text{MS}}$ . A mutator set is complete iff, for any valid sequence of operations  $O$ , and any  $t \in S(O)$ , the completeness error  $\text{Err}_{\text{cmplt}}$  is negligible:

$$\text{Err}_{\text{cmplt}}(\lambda) \triangleq 1 - \Pr \left[ \text{Game}_{\text{cmplt}}^{\text{MS}}(\lambda, O, t) \rightarrow \text{True} \right] \leq \epsilon(\lambda) . \quad (7)$$

Honest-commitment soundness is inadequate for mutator sets due to the changed interface. In particular, and in contrast to accumulation schemes, addition and removal records hide which items they pertain to. The set represented by the set commitment  $\kappa$ , cannot be defined *directly* in terms of the constituent items; but can only be defined *by proxy*, in terms of the items to which the addition and removal records are bound to. It follows that an adequate definition of soundness for mutator sets must imply that both addition records and removal records are binding commitments to particular items.

Moreover, in order for the value  $\kappa$  to compactly represent a set of items it must be binding to that set. However, this variant of binding is tricky to capture formally because the interface does not provide a functionality for opening the commitment and revealing the entire set. Instead, we opt for a simpler and stronger quality: the set commitment must be binding to the set of all addition and removal records that produced it, starting from the commitment to the empty set.

At this point what remains for security is a means for outlawing invalid set updates, such as removing non-members or removing the same item twice. Note that the interface delegates to the security definition the choice as to what happens when an item is added twice. We opt here to disallow removing an item more often than it was added. Technically speaking, this choice implicitly counts multiplicities and thus corresponds to multisets rather than sets, but we ignore this distinction and overload the latter term with the meaning of the former.

**Definition 6 (security for mutator sets).** *A mutator set is secure iff*

- (commit, open) where  $\text{open}(\alpha, t; r) \triangleq \left( \text{commit}(t, r) \stackrel{?}{=} \alpha \right)$  is a binding commitment scheme with negligible binding insecurity:  $\text{InSec}_{\text{bnd}}^{\text{commit}}(\lambda) \leq \epsilon(\lambda)$ ; and
- (drop, open) where  $\text{open}(\rho, t; \mu) \triangleq \left( \text{drop}(t, \mu) \stackrel{?}{=} \rho \right)$  is a binding commitment scheme with negligible binding insecurity:  $\text{InSec}_{\text{bnd}}^{\text{drop}}(\lambda) \leq \epsilon(\lambda)$ ; and
- (set\_commit, open) where  $\text{set\_commit} : \{A, R\}^* \times \{\emptyset\} \rightarrow K$ ,

$$(O, \emptyset) \mapsto \begin{cases} \text{init}(\lambda) & \Leftarrow O = \emptyset \\ \text{add}(\text{set\_commit}(O^*, \emptyset), \alpha) & \Leftarrow O = \alpha \| O^* \wedge \alpha : A \\ \text{remove}(\text{set\_commit}(O^*, \emptyset), \rho) & \Leftarrow O = \rho \| O^* \wedge \rho : R \end{cases}$$

and  $\text{open}(\kappa, O; \emptyset) \triangleq \left( \text{set\_commit}(O; \emptyset) \stackrel{?}{=} \kappa \right)$  is a binding commitment scheme with negligible binding insecurity:  $\text{InSec}_{\text{bnd}}^{\text{set\_commit}}(\lambda) \leq \epsilon(\lambda)$ ; and

- the non-negativity insecurity  $\text{InSec}_{\text{nonneg}}$  is negligible, where this quantity is defined as the maximum probability over all polynomial-time adversaries  $A$  of winning the non-negativity game of Game 4:

$$\text{InSec}_{\text{nonneg}}(\lambda) \triangleq \max_A \Pr \left[ \text{Game}_{\text{nonneg}}^A(\lambda) \right] \leq \epsilon(\lambda) .$$

<b>Game 3: Completeness</b>	<b>Game 4: Non-negativity</b>
<pre> 1 <b>define</b> Game<sub>empt</sub><sup>MS</sup>(<math>\lambda, O, t</math>) <b>as</b>: 2   <math>\kappa \leftarrow \text{init}(\lambda)</math> 3   <math>\mathcal{T} \leftarrow \emptyset</math> 4   <b>for</b> <math>o \in O</math> : 5     <b>if</b> <math>o = (\text{add}, s)</math> : 6       <math>r \xleftarrow{\\$} \mathbb{R}</math> 7       <math>\alpha \leftarrow \text{commit}(\kappa, t; r)</math> 8       <math>\mu \leftarrow \text{prove}(\kappa, t, r)</math> 9       <b>for</b> <math>(s^*, \mu^*) \in \mathcal{T}</math> : 10        <math>\mu^* \leftarrow</math> 11         <math>\text{update}(\kappa, s^*, \mu^*, \alpha)</math> 12        <math>\mathcal{T} \leftarrow \mathcal{T} \cup \{(s, \mu)\}</math> 13        <math>\kappa \leftarrow \text{add}(\kappa, s)</math> 14      <b>if</b> <math>o = (\text{remove}, s)</math> : 15        <b>find</b> <math>\mu</math> such that 16         <math>(s, \mu) \in \mathcal{T}</math> 17         <math>\rho \leftarrow \text{drop}(\kappa, s, \mu)</math> 18         <math>\mathcal{T} \leftarrow \mathcal{T} \setminus \{(s, \mu)\}</math> 19         <b>for</b> <math>(s^*, \mu^*) \in \mathcal{T}</math> : 20          <math>\mu^* \leftarrow</math> 21          <math>\text{update}(\kappa, s^*, \mu^*, \rho)</math> 22        <math>\kappa \leftarrow \text{remove}(\kappa, \rho)</math> 23      <b>find</b> <math>\mu</math> such that <math>(t, \mu) \in \mathcal{T}</math> 24      <b>return</b> <math>\text{verify}(\kappa, t, \mu)</math> </pre>	<pre> 1 <b>define</b> Game<sub>nonneg</sub><sup>A</sup>(<math>\lambda</math>) <b>as</b>: 2   <math>\kappa \leftarrow \text{init}(\lambda)</math> 3   <math>\text{ctr} \leftarrow 0</math> 4   <math>O, \mathcal{T}, \rho^* \leftarrow A(\lambda)</math> 5   <b>for</b> <math>o \in O</math> : 6     <b>if</b> <math>o = (\text{add}, \alpha)</math> : 7       <math>\kappa \leftarrow \text{add}(\kappa, \alpha)</math> 8       <math>\text{ctr} \leftarrow \text{ctr} + 1</math> 9     <b>if</b> <math>o = (\text{remove}, \rho)</math> : 10      <math>\kappa \leftarrow \text{remove}(\kappa, \rho)</math> 11      <math>\text{ctr} \leftarrow \text{ctr} - 1</math> 12   <b>return</b> <math>\text{ctr} =</math> 13     <math>\#\mathcal{T} \wedge \text{remove}(\kappa, \rho^*) \neq \perp \wedge</math> 14     <math>\forall (t, \mu) \in \mathcal{T}. \text{verify}(\kappa, t, \mu) \wedge</math> 15     <math>\rho^* \neq \text{drop}(\kappa, t, \mu)</math> </pre>

### 3.5 Unlinkability and Anonymity

Capturing unlinkability in cryptographic terms is tricky. As a first attempt, one might devise a game whereby the adversary sees two addition records and one removal record and nothing else, but with the guarantee that the removal record is linked to one of the addition records and not to the other, but the order in which they are supplied is random. The adversary wins if he guesses which one. No resource-bounded adversary should be able to win this game with more than a negligible probability.

However, this approach is rather moot because realistic adversaries have access to more information than just three records. They see the set commitment and the entire sequence of addition and removal records that generated it.

As a second attempt, one might devise a game whereby the adversary sees an entire mutator set history play out, and is then tasked with finding the addition record that matches with a given removal record. The baseline is the success probability of the optimal algorithm for this task that reads only the set operations and no cryptographic data. The concrete adversary's advantage when attacking a concrete scheme is the difference between its success probability and this baseline. A mutator set may be defined as *unlinkable* if for all bounded

resource adversaries, this *unlinkability advantage* is negligible in the security parameter.

While the above intuition captures the essence of an ideal mutator set in regards to linkability, it is still rather moot because in practice the produced privacy might be constrained by the sequence of set operations. In the extremal case where every addition except for the last is followed immediately by a removal, the adversary in this game can have a 100% success probability even without breaking any cryptography.

An alternative to capturing unlinkability in cryptographic terms, is to capture anonymity in entropic terms [14]. In the present context, the attacker observes one removal record and assigns probabilities to all possible addition records representing the likelihood, in the attacker’s estimation, that they are linked to the same item. The degree of anonymity is the entropy of this probability distribution. We refer to this notion as *removal entropy*.

### 3.6 Succinctness

In order for a mutator set to be practically useful, it must be *succinct*. Specifically, the commitments and records must be compact, and the operations on them efficient. Informally, *compactness* denotes a polylogarithmic size limit (as a function of the number of historical set operations) on the set commitments, membership proofs, and update records. Likewise, the operations are *efficient* if they have polylogarithmic running time.

**Definition 7 (succinctness).** *A mutator set is succinct relative to a number of set operations  $N$  iff both*

- *the set commitments  $\kappa \in K$ , addition records  $\alpha \in A$ , and removal records  $\rho \in R$ , are compact, i.e., all have a representation that grows at most polylogarithmically:  $\max(|\kappa|, |\alpha|, |\rho|) \leq \text{polylog}(N)$ ;*
- *all algorithms comprising the scheme are efficient: their running time is bounded from above by  $\text{polylog}(N)$ .*

**Succinct Synchronization** Suppose the set was updated  $N$  times in some period of time spent offline. Once online again, how much work does it take to update a membership proof accordingly? Ideally, the answer ought to be *much less than  $N$* .

We opted to omit this feature from the formal definition of succinctness and discuss it only informally, because capturing it formally in complexity theoretic terms is rather tricky. Surely the user wanting to update a membership proof needs to read all the updates, even if only to come to the conclusion that some or perhaps most of them are not necessary for updating *his* membership proofs. We leave the proper formalization of this intuition as an open question.

## 4 Scalable Privacy

In UTXO-based cryptocurrencies like Bitcoin [31], every transaction consumes existing coins and produces new ones. The consumed coins determine a public key, and the transaction is only valid if it is signed by the consumed coins. Synchronized nodes only need to keep track of the set of unspent transaction outputs (UTXO set) to verify new transactions. Once a coin is consumed, and after this expenditure confirmed by enough blocks, old coins can be discarded safely.

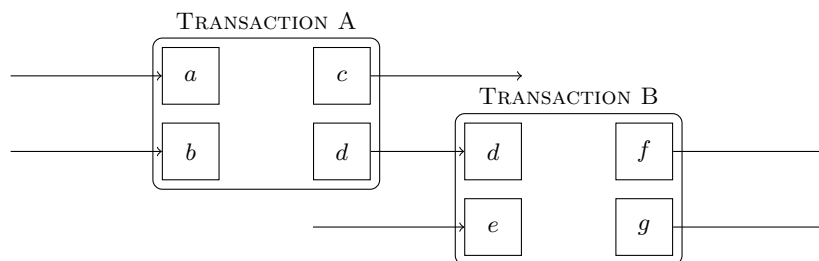


Fig. 2: Fragment of a transaction graph in a transparent UTXO-based ledger.

Figure 2 illustrates the mechanics. UTXOs *a*, *b*, *d*, and *e* are consumed and can thus be discarded to save disk space. UTXOs *c*, *f*, and *g* have not been spent yet and can therefore be used as inputs in future transactions. Every coin can be consumed at most once. The ledger is transparent in two senses. First, the UTXOs themselves are public and so the observer can see their public keys<sup>3</sup> even if he cannot produce a signature for them. Second, the UTXOs are traceable across transactions. While this transparency is beneficial for auditability, it also undermines privacy.

Mutator sets serve as a drop-in replacement for the transparent UTXO set. The consumed coins in a transaction are replaced by removal records and the generated coins by addition records. Both the removal records and the addition records are binding commitments to coins, and so it is possible to prove statements about the committed coins. Specifically, the transaction contains zero-knowledge proofs establishing that a) the spending policies of all inputs are satisfied (for instance, because the owner signed off on the transaction); and b) the coin logic is satisfied (for instance, all outputs have positive amounts and their sum equals that of all inputs). The mutator set guarantees that every coin can be removed at most once.

As coins are hidden behind commitments, and as removal records are unlinkable with addition records, it is impossible to track them. All the while, the ledger’s soundness is guaranteed with zero-knowledge proofs.

<sup>3</sup> The phrase “public key” is used loosely to refer to a binding commitment to a spending policy that must be satisfied by the transaction initiator.






Fig. 3: Fragment of an opaque ledger based on mutator sets.

## 5 Construction

### 5.1 Merkle Mountain Range

A Merkle mountain range (MMR) [38,16] is an authenticated data structure that compactly represents a growing set of items along with generating membership proofs for elements of this set. We provide here only an informal overview. The novelty of this section consists of fitting the folklore notion into our definition of an accumulation scheme.

Informally, a Merkle mountain range is a list of Merkle trees in descending order by size. Whenever two trees of equal height exist, they are folded into one. To add an item to the MMR, add it as a singleton to the small end of the list, and apply the folding operation as often as necessary. Every item ever added is a leaf in *some* tree, and so every item has an authentication path relative to *some* Merkle root.

To fit this object into our definition of accumulation schemes, we populate the types required by it:

- The data structure accumulates items of data type  $T$ .
- A commitment consists of a list of Merkle roots and a leaf count:  $K = [D] \times \mathbb{N}$ .
- A membership proof is an index and an authentication path, the latter of which is a list of hash digests:  $M = \mathbb{N} \times [D]$ .
- An addition record is given by the hash digest of the added element along with its index:  $A = D \times \mathbb{N}$ .
- A modification record is given by the the new leaf, its index, and its authentication path:  $E = T \times \mathbb{N} \times [D]$ .

Likewise, we populate the functions defined by the accumulation scheme interface:

- $\text{init} : [\{1\}] \rightarrow K$  generates an empty list of Merkle roots along with a leaf count set to zero. This commitment represents the empty list.
- $\text{prove} : K \times T \rightarrow M$  produces a membership proof for an item  $t \in T$  that will be valid under the new set commitment if the item is added to the data structure.

- **add** :  $K \times T \rightarrow K \times A$  increments the leaf count, appends the item to the peaks list, reduces the number of peaks by folding equal-height roots as necessary, and outputs a new set commitment along with an addition record.
- **verify** :  $K \times T \times M \rightarrow \{\text{True}, \text{False}\}$  verifies the membership proof which supposedly establishes that the given item is a member of the data structure at the given index.
- **modify** :  $K \times M \times T \times T \rightarrow K \times E$  takes a commitment, a membership proof, an old element, and a new element; and outputs a set commitment list obtained by replacing the indicated element with the new one, along with a modification record. Note that the membership proof for the new item is the same as for the old one.
- **update** :  $K \times T \times M \times (A \vee E) \rightarrow M$  takes a set commitment, the item, its membership proof, and an addition or modification record, and produces a new membership proof. It does this by extending the old authentication path by zero or more digests if an item was added, or modifying some suffix of it if an item in the same tree was modified.

Removals can be simulated by modifying items to set a deletion bit.

**Theorem 1 (MMR Completeness).** *MMR has perfect completeness.*

*Proof.* The statement follows from construction. □

**Theorem 2 (MMR Soundness).** *MMR is honest-commitment sound with error bounded by the hash collision resistance insecurity of  $H$ :  $\text{Err}_{\text{hcsnd}}^{\text{MMR}} \leq \text{InSec}_{\text{coll}}^H$ .*

*Proof.* In order for  $\text{MMR.verify}(\kappa, t, \mu)$  to output **True** for an element  $t \notin S(O)$ ,  $\text{Merkle.verify}(\kappa.p_{[j]}, g, t, \mu.auth\_path)$  has to output **True** for some tree index  $j$  that corresponds to  $\kappa.n$ , and for some (local-tree) leaf index  $g$  corresponding to (MMR-wide) item index  $j$ . Since  $t \notin S(O)$ , no pair  $(g, j)$  satisfies that  $t$  is the  $g$ th leaf of the  $j$ th Merkle tree. So the valid Merkle verification must come from a collision of hashes somewhere.

A successful adversary **A** in the honest-commitment soundness game gives rise to a successful collision finder **B**. The description of this collision-finder is identical to the description of **Game 2** except that it collects all preimage-digest pairs of the hash function as they are computed. After the execution is done, **B** simply sorts for digest. This sorted list will contain one pair with distinct preimages and identical digests, so **B** outputs the preimages. The overhead of **B** is  $\log Q$  where  $Q$  is the number of times the hash function is called. This fact establishes that **B** still runs in polynomial time and therefore has success probability bounded by  $\text{InSec}_{\text{coll}}^H$ . □

## 5.2 Append-Only Commitment List

An append-only commitment list is a data structure that contains a growing list of *commitments* to data items, with the obvious remark that these commitments are themselves data items as well. The state of an append-only commitment list is fully determined by one object:

- a list  $V : [D]$  whose elements are hiding commitments to the accumulated items.

The commitment scheme is defined as follows. To commit to a message  $t \in T$ , obtain randomness  $r \xleftarrow{\$} R$  drawn uniformly at random from a sufficiently large set of symbols  $R$ , and compute the commitment  $c \in D$  as  $c \leftarrow H(t||r)$ . To verify that a commitment  $c$  binds to a message  $t$ , it suffices to supply  $r$  as it enables the test  $c \stackrel{?}{=} H(t||r)$ .

Adding an item  $t \in T$  to the set consists of simply appending a new commitment. Having added an item, it is possible to prove membership of the item  $t$  to the set in a way that can be verified efficiently, provided that the item's *index of addition*  $m \in \mathbb{N}$  is known. Therefore, this index must be part of the item's membership proof, along with the randomness.

The scheme AOCL has the interface of an accumulation scheme, with the exception that removals are disallowed. The functions, `add` and `verify` include logic specific to the commitments. The remaining functions are omitted because obvious.

Algorithm 1: AOCL.add	Algorithm 2: AOCL.verify
<pre> <b>1</b> <b>define</b> AOCL.add(<math>\kappa = V, t; r</math>)   <b>as:</b> <b>2</b>   <math>m \leftarrow \#V</math> <b>3</b>   <math>c \leftarrow H(t  r)</math> <b>4</b>   <math>V \leftarrow V  c</math> <b>5</b>   <b>return</b>       <math>\kappa = V, \mu = (m, r), \alpha = c</math> </pre>	<pre> <b>1</b> <b>define</b> AOCL.verify(<math>\kappa =</math>   <math>V, t, \mu = (m, r)</math>) <b>as:</b> <b>2</b>   <math>c \leftarrow H(t  r)</math> <b>3</b>   <b>return</b> <math>V_{[m]} \stackrel{?}{=} c</math> </pre>

**Security.** The completeness of the append-only timestamped commitment list follows from constuction. Its honest-commitment soundness is less straightforward. We consider Game 2 with the restriction that  $O$  does not contain any remove operations.

**Theorem 3.** *As an accumulation scheme, AOCL has honest-commitment soundness error bounded by*

$$\text{Err}_{\text{snd}}^{\text{AOCL}} \leq \text{InSec}_{\text{coll}}^{\text{H}} . \quad (8)$$

*Proof.* Assume that  $(\text{add}, t) \notin O$  but `verify` returns `True`. This event implies a collision since for some  $t^* \neq t$  and some  $r^*$ ,  $c = H(t||r) = H(t^*||r^*)$ . The pair  $t^*, r^*$  is obtained by recording all preimage-digest pairs as the hash function is computed. A successful adversary can be transformed into a successful collision-finder, whose success probability is bounded by definition by  $\text{InSec}_{\text{coll}}^{\text{H}}$ .  $\square$

### 5.3 Sliding Window Bloom Filter

A *sliding window Bloom filter* represents a growing set of *timestamped* items  $(l, t) \in \mathbb{N} \times T = T^*$ . The state of a filter is fully defined by an infinite array of bits  $F \in \{0, 1\}^*$ . The bits in the array are initially set to zero but are set to one as elements are added to the set. Since only finite sets are ever represented, only a finite section of the bit array is ever represented.

A sliding window Bloom filter is parameterized by four integers:

- the *window width*  $w \in \mathbb{N}$ , which determines the size of the interval where bits can be set as items are added;
- the *pointer count*  $k \in \mathbb{N}$ , which determines the maximum number of bits flipped by any item;
- the *batch size*  $b \in \mathbb{N}$ , which determines the number of items in between window slides; and
- the *step size*  $s \in \mathbb{N}$ . Assume the step size divides the window width.

Additionally, we associate a hash function  $H_w : \mathbb{N} \times T \times \mathbb{R} \times \mathbb{N} \rightarrow \{0, \dots, w-1\} \subset \mathbb{N}$  with the window size  $w$ .

A filter defines two set operations:

- **add** :  $\{0, 1\}^* \times T^* \times \mathbb{R} \rightarrow (\{0, 1\}^* \times \mathbb{N}) \times \mathbb{N}^k$ , which adds a timestamped and randomized item to the set represented by the filter.
- **probe** :  $\{0, 1\}^* \times T^* \times \mathbb{R} \rightarrow \{\text{True}, \text{False}\}$ , which tests if a given item is a member of the set with true positives, false positives, and true negatives, but no false negatives.

What sets a sliding window Bloom filter apart from a regular Bloom filter is that the array of bits has infinite length but the window of settable bits shifts in accordance with the timestamp  $l$  that items come with. Specifically, the **add** operation consists of sampling  $k$  indices in  $\{0, \dots, w-1\}$  pseudorandomly, translating the indices by  $\lfloor \frac{l}{b} \rfloor \cdot s$ , and setting the indicated bits.

<b>Algorithm 3: SWBF.add</b>	<b>Algorithm 4: SWBF.probe</b>
<pre> 1 <b>define</b> SWBF.add(<math>F, (l, t), r</math>)   <b>as:</b> 2   <math>J \leftarrow \emptyset</math> 3   <b>for</b> <math>i \in \{0, \dots, k-1\}</math> : 4     <math>j \leftarrow H_w(l \parallel t \parallel r \parallel i)</math> 5     <math>J \leftarrow J \cup j</math> 6     <math>F_{\lfloor j + \lfloor \frac{l}{b} \rfloor \cdot s \rfloor} \leftarrow 1</math> 7   <b>return</b> <math>F, J</math> </pre>	<pre> 1 <b>define</b> SWBF.probe(<math>F, (l, t), r</math>)   <b>as:</b> 2   <b>for</b> <math>i \in \{0, \dots, k-1\}</math> : 3     <math>j \leftarrow H_w(l \parallel t \parallel r \parallel i)</math> 4     <b>if</b> <math>F_{\lfloor j + \lfloor \frac{l}{b} \rfloor \cdot s \rfloor} = 0</math> : 5       <b>return</b> false 6   <b>return</b> true </pre>

To probe whether an element is a member of the set, check the indicated array elements. If any one element is 0, then the element definitely was not

added. Crucially, this operation requires the timestamp  $l$  because otherwise the offset is not known.

The false positive probability corresponds to the probability that  $k$  given indices are nonzero. This probability is

$$\left(1 - \left(1 - \frac{1}{w}\right)^{\frac{kbw}{s}}\right)^k \approx \left(1 - e^{-\frac{kb}{s}}\right)^k. \quad (9)$$

By setting  $s$  to a significant fraction of  $kb$ , and  $w$  much larger than  $kb$ , it is possible to engineer a cryptographically small false positive probability. The optimal value of  $k$  as a function of  $b$  and  $s$  is  $k = \frac{s \cdot \ln 2}{b}$ . Plugging that ratio into the expression (9) gives a probability of roughly  $2^{-k}$ .

#### 5.4 Non-Compact Mutator Sets

In this section we build a mutator set from an append-only commitment list and a sliding window Bloom filter. This mutator set lacks a crucial property: compactness. The next section remedies this deficiency.

The general strategy is to use canonical commitments to items as addition records and keep track of them in the append-only commitment list. Simultaneously, the removal records contain a list of pseudorandom indices of set bits in the sliding window Bloom filter, which tracks which items are being removed. A membership proof to the mutator set consists of a membership proof to the append-only commitment list, and a *non-membership* proof to the sliding window Bloom filter. This construction turns the cryptographically small false positive probability into a negligible completeness error.

The item acquires a timestamp when it is added to the set. The timestamp is part of the membership proof  $\mu$ . The timestamp is necessary to slide the Bloom filter window, which in turn is necessary to avoid saturating the filter. Therefore, the further apart items are added, the more noticeably will their removal records reflect this distance and relative order. Phrased differently, unlinkability is optimal for items added immediately after one another; and degrades with their distance.

To ensure non-negativity, the property that prevents the set commitment from being updated with a removal record for a non-member, it is necessary to add a non-interactive zero-knowledge proof  $\pi \in \Pi$  to the removal record. Informally, this proof establishes that the index set  $J$  of set bits was derived from some item, timestamp, randomizer, and matching canonical commitment that lives in the AOCL. Moreover, it is important that this proof system has knowledge soundness with respect to the secrets: specifically,  $\text{ZKPoK}\{(l, t, r, c) : c = \mathbf{H}(t||r) \wedge c \in V \wedge J = \{\mathbf{H}_w(l||t||r||i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq i < k\}\}$ . Let  $\varsigma$  denote the soundness error for this proof system; but note that completeness and zero-knowledge are taken for granted. For the purpose of comparing equality between removal records, the proof  $\pi$  is ignored; in particular, this means that malleating the proof is not considered an attack on the binding property of **drop**.

The types are defined as follows.

- $MS.T = T$
- $MS.K = [C] \times \{0, 1\}^*$
- $MS.M = \mathbb{N} \times \mathbb{R}$
- $MS.A = C \times \mathbb{N}$
- $MS.R = \mathbb{N}^k \times \Pi$

The following algorithm definitions use subscript 1 in anticipation of another variant with improved features.

<hr/> <p><b>Algorithm 5: <math>MS_1</math>.init</b></p> <pre> 1 <b>define</b> <math>MS_1</math>.init(<math>i</math>) <b>as:</b> 2     <math>\lambda \leftarrow \#i</math> 3     <b>return</b> <math>\kappa = ([], [0]^\infty)</math> </pre> <hr/>	<hr/> <p><b>Algorithm 6: <math>MS_1</math>.commit</b></p> <pre> 1 <b>define</b> <math>MS_1</math>.commit(<math>t, r</math>) <b>as:</b> 2     <math>\_, \alpha \leftarrow AOCL.add([], t; r)</math> 3     <b>return</b> <math>\alpha</math> </pre> <hr/>
<hr/> <p><b>Algorithm 7: <math>MS_1</math>.prove</b></p> <pre> 1 <b>define</b>   <math>MS_1</math>.commit(<math>\kappa = (V, F), t, r</math>)   <b>as:</b> 2     <math>\_, \mu, \_ \leftarrow AOCL.add(V, t; r)</math> 3     <b>return</b> <math>\mu</math> </pre> <hr/>	<hr/> <p><b>Algorithm 8: <math>MS_1</math>.verify</b></p> <pre> 1 <b>define</b> <math>MS_1</math>.verify(<math>\kappa =</math>   (<math>V, F</math>), <math>t, \mu = (l, r)</math>) <b>as:</b> 2     <b>return</b>     AOCL.verify(<math>V, t, (l, r)</math>) =     True <math>\wedge</math>     SWBF.probe(<math>F, (l, t), r</math>) =     False </pre> <hr/>
<hr/> <p><b>Algorithm 9: <math>MS_1</math>.add</b></p> <pre> 1 <b>define</b> <math>MS_1</math>.add(<math>\kappa = (V, F), \alpha</math>)   <b>as:</b> 2     <math>V \leftarrow V \parallel \alpha</math> 3     <b>return</b> <math>\kappa = (V, F)</math> </pre> <hr/>	<hr/> <p><b>Algorithm 10: <math>MS_1</math>.drop</b></p> <pre> 1 <b>define</b> <math>MS_1</math>.drop(<math>\kappa =</math>   (<math>V, F</math>), <math>t, \mu = (l, r)</math>) <b>as:</b> 2     <math>\_, J \leftarrow</math>     SWBF.add(<math>F, (l, t), r</math>) 3     <math>\pi \leftarrow \Pi</math>.prove(<math>\{(l, t, r, c) :</math>     <math>c = H(t \parallel r) \wedge c \in V \wedge J =</math>     <math>\{H_w(l \parallel t \parallel r \parallel i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq</math>     <math>i &lt; k\}</math>) 4     <b>return</b> <math>\kappa = (V, F)</math> </pre> <hr/>

While there are records  $\alpha$  and  $\rho$ , neither record is necessary to update the membership proof  $\mu$ . Therefore the update function is just the identity.

**Theorem 4 (completeness of  $MS_1$ ).** *As a mutator set,  $MS_1$  is complete with completeness error  $\left(1 - \left(1 - \frac{1}{w}\right)^{\frac{kbw}{s}}\right)^k \approx \left(1 - e^{-\frac{kb}{s}}\right)^k$ .*

*Proof.* Suppose  $\text{Game}_{\text{cempt}}(\lambda, O_1, O_2, t)$  output False in some execution for a valid sequence of operations  $O_1 \parallel (\text{add}, t) \parallel O_2$ . This proposition implies that  $\text{verify}(\kappa, t, \mu)$  on line Line 21 of Algorithm 1 fails, which in turn implies that either  $\text{AOCL.verify}(V, t, (m, r)) = \text{False}$  or  $\text{SWBF.probe}(F, (l, t)) = \text{True}$  on line Line 2 of Algorithm 8. The AOCL has perfect completeness and so the left clause

---

**Algorithm 11: MS<sub>1</sub>.remove**


---

```

1 define MS1.remove( $\kappa = (V, F), \rho$ ) as:
2   if  $H.verify(\rho, \pi, \{(l, t, r, c) : c = H(t||r) \wedge c \in V \wedge J =$ 
   |  $\{H_w(l||t||r||i) + \lfloor \frac{l}{b} \rfloor \cdot s | 0 \leq i < k\}) = \text{False} :$ 
3   |   return  $\perp$ 
4   else
5   |    $F^*, - \leftarrow \text{SWBF.add}(F, (l, t), r)$ 
6   |   if  $F^* \neq F :$ 
7   |   |   return  $\kappa = (V, F^*)$ 
8   |   else
9   |   |   return  $\perp$ 

```

---

of this conjunction is guaranteed to return True. The right clause returning true corresponds to the false positive probability of the SWBF, which is given by Eqn. 9.  $\square$

**Theorem 5 (security of MS<sub>1</sub>).** *MS<sub>1</sub> is secure; specifically:*

- *commit, in combination with the standard opening algorithm, has bounded binding insecurity  $\text{InSec}_{\text{bnd}}^{\text{commit}} \leq \text{InSec}_{\text{coll}}^{\text{H}}$ ;*
- *drop, in combination with the standard opening algorithm, has bounded binding insecurity  $\text{InSec}_{\text{bnd}}^{\text{drop}} \leq \frac{Q(Q-1)k!}{2w^k} + \text{Adv}_{\text{PRF}}^{\text{H}_w}$ , where  $Q$  is the number of queries to  $\text{H}_w$ ;*
- *the function  $\text{set\_commit} : \{A, R\}^* \times \{\emptyset\} \rightarrow K$ ,*

$$(O, \emptyset) \mapsto \begin{cases} \text{init}(\lambda) & \Leftarrow O = \emptyset \\ \text{add}(\text{set\_commit}(O^*, \emptyset), \alpha) & \Leftarrow O = \alpha || O^* \wedge \alpha : A \\ \text{remove}(\text{set\_commit}(O^*, \emptyset), \rho) & \Leftarrow O = \rho || O^* \wedge \rho : R \end{cases}$$

*in combination with the standard opening algorithm, has bounded binding*

$$\text{insecurity } \text{InSec}_{\text{bnd}}^{\text{set\_commit}} \leq \left(1 - \left(1 - \frac{1}{w}\right)^{\frac{kbw}{s}}\right)^k + \text{Adv}_{\text{PRF}}^{\text{H}_w};$$

- *MS<sub>1</sub> has bounded non-negativity insecurity  $\text{InSec}_{\text{nonneg}}^{\text{MS}_1} \leq \varsigma + \text{InSec}_{\text{coll}}^{\text{H}}$ .*

*Proof. Commit.* The bounded binding insecurity of `commit` follows from the fact that its definition coincides with that of a canonical commitment scheme. The adversary who successfully equivocates on a given commitment can be used to build a collision-finder, so  $\text{InSec}_{\text{bnd}}^{\text{commit}}(\lambda) \leq \text{InSec}_{\text{coll}}^{\text{H}}$ .

*Drop.* If the hash function  $\text{H}_w : \Sigma^* \rightarrow \{0, \dots, w-1\}$  is uniformly random, then so is the implicit mapping from  $(l, t, r, i)$  to  $J \in \{0, \dots, w-1\}^k$  as a *list* of indices, *i.e.*, where the order matters. In the worst case there are  $k$  unique indices and then there are  $k!$  ways to order them. So every index *set* defines up to  $k!$  different but equivalent *lists*. Therefore, the probability of observing a collision of *sets* is bounded by  $\frac{Q(Q-1)k!}{2w^k}$ , where  $Q$  is the number of queries to  $\text{H}_w$ . A successful adversary in the binding game for `drop` can be transformed into an algorithm that finds collisions of index sets, and so its success probability

must be bound by this quantity. An adversary with success probability in excess of this amount can be used to distinguish the hash function (interpreted as a PRF with implicit key  $r$ ) from random. The advantage of this distinguisher is bounded by  $\text{Adv}_{\text{PRF}}^{\text{H}_w}$  by assumption. By the union bound, the binding adversary's success probability is bounded by  $\frac{Q(Q-1)k!}{2w^k} + \text{Adv}_{\text{PRF}}^{\text{H}_w}$ .

*Set commit.* The AOCL is perfectly binding to the set of addition records because it is a list. A pair of different sets of removal records  $O_1 \neq O_2$  that induce the same bits in the SWBF can be reduced to a pair of non-intersecting sets with the same properties in addition to  $O_1 \cap O_2 = \emptyset$ . Assuming  $\text{H}_w$  is uniformly random it is possible to calculate the probability that all bits indicated by some  $\rho \in O_1$  are already set by the removal records in  $O_2$ . In fact, this calculation repeats the derivation of the completeness error, which is  $\left(1 - \left(1 - \frac{1}{w}\right)^{\frac{kbw}{s}}\right)^k \approx \left(1 - e^{-\frac{kb}{s}}\right)^k$ . Furthermore, the adversary who succeeds in excess of this probability can be used to build a distinguisher in the PRF game to distinguish  $\text{H}_w$  from random, whose advantage is bounded by  $\text{Adv}_{\text{PRF}}^{\text{H}_w}$ . By the union bound, the binding adversary's success probability is bounded by

$$\left(1 - \left(1 - \frac{1}{w}\right)^{\frac{kbw}{s}}\right)^k + \text{Adv}_{\text{PRF}}^{\text{H}_w} .$$

*Non-negativity.* Suppose the adversary wins Game 4, meaning that  $\text{remove}(\kappa, \rho^*) \neq \perp$  even though all items still in the mutator set generate removal records distinct from  $\rho^*$ . By using the extractor implied by the knowledge-soundness of  $\pi$  one extracts the witness  $(l, t, r, c)$  such that  $\text{H}(t||r) = c \in V$ . There are two possibilities:

1.  $(t, (l, r)) \notin \mathcal{T}$ . In this case we have an attack on the binding property of `commit` because  $\mathcal{T}$  must contain some other element that also generates  $c \in V$ . This attack on binding in turn implies an attack on the collision resistance of  $\text{H}$ .
2.  $(t, (l, r)) \in \mathcal{T}$ . This is a contradiction because on the one hand  $\rho^* = \text{drop}(\kappa, t, \mu = (l, r))$  because this is implied by  $\text{remove}(\kappa, \rho^*) \neq \perp$  in the middle of Line 12 in Algorithm 4; on the other hand  $\rho^* \neq \text{drop}(t, \mu = (l, r))$  is asserted in the same line.

In summary, if the adversary  $\mathbf{A}$  is successful, and if the extractor  $\mathbf{E}$  is successful, then some collision-finder  $\mathbf{C}$  for  $\text{H}$  is successful. One derives a bound on the success probability of  $\mathbf{A}$  from this:

$$1 = \Pr[\mathbf{A}\checkmark \wedge \mathbf{E}\checkmark \Rightarrow \mathbf{C}\checkmark] \tag{10}$$

$$= \Pr[\neg(\mathbf{A}\checkmark \wedge \mathbf{E}\checkmark) \vee \mathbf{C}\checkmark] \tag{11}$$

$$= \Pr[\mathbf{A}\mathbf{x} \vee \mathbf{E}\mathbf{x} \vee \mathbf{C}\checkmark] \tag{12}$$

$$\leq \Pr[\mathbf{A}\mathbf{x}] + \Pr[\mathbf{E}\mathbf{x}] + \Pr[\mathbf{C}\checkmark] \tag{13}$$

$$\leq 1 - \Pr[\mathbf{A}\checkmark] + \varsigma + \text{InSec}_{\text{coll}}^{\text{H}} \tag{14}$$



$$\Pr[\mathbf{A}\checkmark] \leq \varsigma + \text{InSec}_{\text{coll}}^{\text{H}} . \quad (15)$$

□

## 5.5 Compact Mutator Sets

A pair of MMRs transform the non-compact mutator set  $\text{MS}_1$  into a compact one ( $\text{MS}_2$ ). The general idea behind the transformation is this:

- Use one MMR to compactify the AOCL  $V$ , and adapt the algorithms that access it accordingly.
  - Use another MMR to compactify the *inactive part* of the SWBF  $F$ , and adapt the algorithms that access it accordingly. For the purpose of computing Merkle trees, the chunk size is  $s$ . In other words, each leaf contains a contiguous chunk of  $s$  bits from  $F$ . The active part is represented explicitly as an array of  $w$  bits. The dividing line between inactive and active lies at index  $\lfloor \frac{N}{b} \rfloor \cdot s$ , where  $N$  is the total number of items added to the mutator set.
  - Replace the commitment  $\kappa = (V, F) : [\text{D}] \times \{0, 1\}^*$  with  $(\kappa_V, \kappa_F, F_A) : \text{MMR}.K \times \text{MMR}.K \times \{0, 1\}^w = K$ .
  - Expand the membership proof  $\mu = (l, r) : \mathbb{N} \times \mathbb{R}$  with:
    - an authentication path  $\text{auth\_path}_V : [\text{D}]$  which is the authentication path for element  $m$  in the MMR of  $V$ ; and
    - a dictionary sending indices to chunk-and-authentication-path pairs:  $\text{target\_chunks} : \mathbb{N} \rightarrow (\{0, 1\}^s \times [\text{D}])$ . This dictionary is used to access bits in the inactive part of the filter, specifically *those* bits that *will be* set if the item is removed:  $J = \{\lfloor \frac{l}{b} \rfloor \cdot s + \text{H}_w(l \| t \| r \| i) \mid 0 \leq i < k\}$ . The keys in this dictionary correspond to *chunk indices*  $c_i$  derived from the *bit indices*  $b_i$  via  $c_i = \lfloor \frac{b_i}{s} \rfloor$ . There are no entries in this dictionary for bits in the active portion of the filter.
- The new membership proof is therefore  $(l, r, \text{auth\_path}_V, \text{target\_chunks}) : \mathbb{N} \times \mathbb{R} \times [\text{D}] \times (\mathbb{N} \rightarrow (\{0, 1\}^s \times [\text{D}])) = M$ .
- Modify the zero-knowledge proof system  $\Pi$  to prove the correct execution of  $\text{MMR.verify}(\kappa_V, c, \mu_V)$  rather than the membership  $c \in V$ .
  - Expand the removal record  $\rho = (J, \pi) : \mathbb{N}^k \times \Pi$  with a dictionary of leaf-and-authentication-path pairs where the leaf represents the chunk of the filter  $F$  needed to set the inactive bits indicated by  $J$ :  $\rho = (J, \pi, \text{target\_chunks}) : \mathbb{N}^k \times \Pi \times (\mathbb{N} \rightarrow (\{0, 1\}^s \times [\text{D}])) = R$ . There are no leaves or authentication paths for bits in the active part of the filter.

The algorithms are defined explicitly below, except for  $\text{MS}_2.\text{commit}$  which is identical to  $\text{MS}_1.\text{commit}$ . Unfortunately, using the MMR to compress commitments requires unrolling the calls to the proper functions of AOCL and SWBF.

**Security.** The use of MMRs does not affect completeness. In terms of security, if the MMRs are sound then a successful adversary against the compact mutator set  $\text{MS}_2$  can be reduced to a successful adversary against the non-compact mutator set  $\text{MS}_1$ .

---

**Algorithm 12: MS<sub>2</sub>.init**

---

```
1 define MS2.init( $\lambda$ ) as:
2    $\lambda \leftarrow \#\lambda$ 
3    $\kappa_V \leftarrow \text{MMR.init}(\lambda)$ 
4    $\kappa_F \leftarrow \text{MMR.init}(\lambda)$ 
5    $F_A \leftarrow [0]^w$ 
6   return  $\kappa = (\kappa_V, \kappa_F, F_A)$ 
```

---

---

**Algorithm 13: MS<sub>2</sub>.add**

---

```
1 define
   MS2.add( $\kappa = (\kappa_V, \kappa_F, F_A), \alpha$ )
   as:
2    $\kappa_{V,-} \leftarrow \text{MMR.add}(\kappa_V, \alpha)$ 
3   if  $\kappa_V \equiv 0 \pmod{b}$  :
4     // window slides
     slid_chunk  $\leftarrow F_{A[0:s]}$ 
5      $F_A \leftarrow F_{A[s:]} \parallel [0]^{w-s}$ 
6      $\kappa_{F,-} \leftarrow$ 
       MMR.add(chunk)
7   return  $\kappa = (\kappa_V, \kappa_F, F_A)$ 
```

---

---

**Algorithm 14: MS<sub>2</sub>.prove**

---

```
1 define MS2.prove( $\kappa = (\kappa_V, \kappa_F, F_A), t, r$ ) as:
2   // compute commitment to item
    $l \leftarrow \kappa_V.n$ 
3    $c \leftarrow H(t \parallel r)$ 
   // add commitment to list
4    $\kappa_V, \mu_V, \alpha \leftarrow \text{MMR.add}(\kappa_V, c)$ 
   // test if window slides; if so update filter MMR
5   target_chunks  $\leftarrow \text{dict}()$ 
6   if  $\kappa_V.n \equiv 1 \pmod{b}$  :
7     chunk  $\leftarrow F_{A[0:s]}$ 
8      $F_A \leftarrow F_{A[s:]} \parallel [0]^s$ 
9      $\kappa_F, \mu_F, - \leftarrow \text{MMR.add}(\kappa_F, \text{chunk})$ 
     // prepare filter MMR authentication paths
10    if  $j < \lfloor \frac{\kappa_V.n}{b} \rfloor \cdot s$  for some  $j \in J = \{H_w(t \parallel i) + l \cdot s \mid 0 \leq i < k\}$  :
11      | target_chunks( $j$ )  $\leftarrow (\text{chunk}, \mu_F)$ 
12  return  $\mu = (l, r, \mu_V, \text{target\_chunks})$ 
```

---

---

**Algorithm 15: MS<sub>2</sub>.verify**

---

```
1 define MS2.verify( $\kappa = (\kappa_V, \kappa_F, F_A), t, \mu = (l, r, \mu_V, \text{target\_chunks})$ ) as:
2    $c \leftarrow H(t \parallel r)$ 
3   if MMR.verify( $\kappa_V, \text{auth\_path}_V, c$ ) = False : return false
4   for  $j \in J = \{H_w(t \parallel i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq i < k\}$  :
5     if  $j > \lfloor \frac{\kappa_V.n}{b} \rfloor \cdot s$  :
6       | if  $F_{A[j - \lfloor \frac{l}{b} \rfloor \cdot s]} = 1$  : return False
7     else
8       | chunk, path  $\leftarrow \text{target\_chunks}(j)$ 
9       | if chunk $_{[j \bmod s]}$  = 1 : return False
10      | if MMR.verify( $\kappa_F, \text{path}, \text{chunk}$ ) = False : return False
11  return True
```

---

---

**Algorithm 16: MS<sub>2</sub>.drop**

---

```
1 define MS2.drop( $\kappa = (\kappa_V, \kappa_F, F_A), t, \mu = (l, r, \mu_V, target\_chunks)$ ) as:
2    $J \leftarrow \{H_w(l||t||r||i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq i < k\}$ 
3    $\pi \leftarrow \Pi.prove(\{(l, t, r, c, \mu_V) : c = H(t||r) \wedge MMR.verify(\kappa_V, c, \mu_V) \wedge J =$ 
    $\{H_w(l||t||r||i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq i < k\}\})$ 
4   return  $\rho = (J, \pi, target\_chunks)$ 
```

---

---

**Algorithm 17: MS<sub>2</sub>.remove**

---

```
1 define MS2.remove( $\kappa = (\kappa_V, \kappa_F, F_A), \rho = (\rho.J, \rho.\pi, \rho.target\_chunks)$ ) as:
2    $set\_some\_bit \leftarrow \text{False}$ 
3   if  $\Pi.verify(\rho.\pi, \{(l, t, r, c, \mu_V) : c = H(t||r) \wedge MMR.verify(\kappa_V, c, \mu_V) \wedge J =$ 
    $\{H_w(l||t||r||i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq i < k\}\}) = \text{False} :$ 
4     return  $\perp$ 
5   for  $j \in J :$ 
6     if  $j > \lfloor \frac{\kappa_V \cdot n}{b} \rfloor \cdot s :$ 
7       if  $F_A[j - \lfloor \frac{l}{b} \rfloor \cdot s] = 0 :$ 
8          $set\_some\_bit \leftarrow \text{True}$ 
9          $F_A[j - \lfloor \frac{l}{b} \rfloor \cdot s] \leftarrow 1$ 
10      else
11         $path, chunk \leftarrow target\_chunks(\lfloor \frac{j}{b} \rfloor)$ 
12        if  $MMR.verify(\kappa_F, chunk, path) = \text{False} :$ 
13          return  $\perp$ 
14         $new\_chunk \leftarrow chunk$ 
15        if  $new\_chunk[j \bmod s] = 0 :$ 
16           $set\_some\_bit \leftarrow \text{True}$ 
17           $new\_chunk[j \bmod s] \leftarrow 1$ 
18           $\kappa_F \leftarrow MMR.modify(\kappa_F, path, chunk, new\_chunk)$ 
19   if  $set\_some\_bit :$ 
20     return  $\kappa = (\kappa_V, \kappa_F, F_A)$ 
21   else
22     return  $\perp$ 
```

---

---

**Algorithm 18: MS<sub>2</sub>.update**

---

```
1 define MS2.update( $\kappa = (\kappa_V, \kappa_F, F_A), t, \mu = (l, r, \mu_V, \text{target\_chunks}), v$ ) as:
2   if  $v : A$  :
3      $\rightarrow, \mu_V \leftarrow \text{MMR.update}(\kappa_V, t, (l, \mu_V), \alpha = (v, \kappa_V.n))$ 
4     if  $\kappa_V.n + 1 \equiv 0 \pmod{b}$  :
5       // chunk slides
6        $J \leftarrow \{H_w(l||t||r||i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq i < k\}$ 
7       if  $J \cap \{\lfloor \frac{\kappa_V.n}{b} \rfloor \cdot s, \dots, \lfloor \frac{\kappa_V.n}{b} \rfloor \cdot s + s - 1\} \neq \emptyset$  :
8         // need to keep track of this chunk
9          $\text{slid\_chunk} \leftarrow F_{A[0:w]}$ 
10         $\text{slid\_mp} \leftarrow \text{MMR.prove}(\kappa_F, \text{slid\_chunk})$ 
11         $\text{target\_chunks}(\kappa_F) \leftarrow (\text{slid\_mp}, \text{slid\_chunk})$ 
12    else
13      // v is a removal record
14       $\rightarrow, \rho.\text{target\_chunks} \leftarrow v$ 
15      for  $(\rho.\text{chunk\_index}, \rho.\text{chunk\_mp}, \rho.\text{chunk}) \in \rho.\text{target\_chunks}$  :
16        // prepare modification in SWBF MMR
17         $\varepsilon \leftarrow (\rho.\text{chunk}, \rho.\text{chunk\_index}, \rho.\text{chunk\_path})$ 
18        for  $(\text{own\_chunk\_index}, \text{own\_chunk\_mp}, \text{own\_chunk}) \in \text{target\_chunks}$  :
19          if  $\rho.\text{chunk\_index} = \text{own\_chunk\_index}$  :
20             $\text{own\_chunk} \leftarrow \rho.\text{chunk}$ 
21          else
22             $\text{own\_chunk\_mp} \leftarrow$ 
23               $\text{MMR.update}(\kappa_F, \text{own\_chunk}, \text{own\_chunk\_mp}, \varepsilon)$ 
24            // update own dictionary
25             $\text{target\_chunks}(\text{own\_chunk\_index}) \leftarrow$ 
26               $(\text{own\_chunk\_mp}, \text{own\_chunk})$ 
27    return  $\mu = (l, r, \mu_V, \text{target\_chunks})$ 
```

---

**Theorem 6 (security of MS<sub>2</sub>).** MS<sub>2</sub> is secure; specifically:

- **commit**, in combination with the standard opening algorithm, has bounded binding insecurity  $\text{InSec}_{\text{bnd}}^{\text{commit}} \leq \text{InSec}_{\text{coll}}^{\text{H}}$ ;
- **drop**, in combination with the standard opening algorithm, has bounded binding insecurity  $\text{InSec}_{\text{bnd}}^{\text{drop}} \leq \frac{Q(Q-1)k!}{2w^k} + \text{Adv}_{\text{PRF}}^{\text{H}_w}$ , where  $Q$  is the number of queries to  $\text{H}_w$ ;
- the function  $\text{set\_commit} : \{A, R\}^* \times \{\emptyset\} \rightarrow K$ ,

$$(O, \emptyset) \mapsto \begin{cases} \text{init}(\lambda) & \Leftarrow O = \emptyset \\ \text{add}(\text{set\_commit}(O^*, \emptyset), \alpha) & \Leftarrow O = \alpha \| O^* \wedge \alpha : A \\ \text{remove}(\text{set\_commit}(O^*, \emptyset), \rho) & \Leftarrow O = \rho \| O^* \wedge \rho : R \end{cases}$$

in combination with the standard opening algorithm, has bounded binding

insecurity  $\text{InSec}_{\text{bnd}}^{\text{set\_commit}} \leq \left(1 - \left(1 - \frac{1}{w}\right)^{\frac{kbw}{s}}\right)^k + \text{Adv}_{\text{PRF}}^{\text{H}_w} + 2 \cdot \text{InSec}_{\text{coll}}^{\text{H}}$ ;

- MS<sub>2</sub> has bounded non-negativity insecurity  $\text{InSec}_{\text{nonneg}}^{\text{MS}_2} \leq 2 \cdot \varsigma + 3 \cdot \text{InSec}_{\text{coll}}^{\text{H}}$ .

*Proof. Commit.* MS<sub>2</sub>.commit is identical to MS<sub>1</sub>.commit and so is its security.

*Drop.* For the purpose of testing equality both the proof  $\pi$  and the dictionary  $\text{target\_chunks}$  are ignored. As a result, the security argument is identical to that of MS<sub>1</sub>.drop.

*Set commit.* Suppose the adversary finds two sequences of operations  $O_1 \neq O_2$  resulting in the same set commitment  $\kappa = (\kappa_V, \kappa_F, F_A)$ . Then one (or more) of three possibilities must hold:

1.  $O_1$  and  $O_2$  define two different AOCLs even though their MMR set commitment  $\kappa_V$  is identical. This implies a break of the MMR’s honest-commitment soundness.
2.  $O_1$  and  $O_2$  define two different SWBFs even though the MMR set commitment to the inactive parts are identical. This implies a break of the MMR’s honest-commitment soundness.
3.  $O_1$  and  $O_2$  define identical AOCLs and SWBFs. In this case  $O_1$  and  $O_2$  give rise to an attack against MS<sub>1</sub>.set\_commit.

So relative to MS<sub>1</sub>.set\_commit, MS<sub>2</sub>.set\_commit incurs a term  $2 \cdot \text{Err}_{\text{hcsnd}}^{\text{MMR}} \leq 2 \cdot \text{InSec}_{\text{coll}}^{\text{H}}$  security degradation.

*Non-negativity.* Suppose that an adversary  $A_2$  wins the non-negativity game for MS<sub>2</sub>. Build an adversary  $A_1$  against the non-negativity game for MS<sub>1</sub> as follows. Take the output from  $A_2$  and strip all the Merkle authentication paths and inactive SWBF chunks. Modify the proof  $\pi_2$  on the removal record by *a*) extracting the witness  $(l, t, r, c, \mu_V)$  using the extractor  $E$ ; and *b*) generating  $\pi_1$  with  $\text{II.prove}(\{(l, t, r, c) : c = \text{H}(t||r) \wedge c \in V \wedge J = \{\text{H}_w(l||t||r||i) + \lfloor \frac{l}{b} \rfloor \cdot s \mid 0 \leq i < k\})$ . There are three conditions on which  $A_1$  fails even though  $A_2$  succeeds:

- The extractor  $E$  fails to produce a valid witness; the probability of this event is  $\varsigma$ .
- The witness produced by  $E$  is valid  $c \notin V$  even though  $\text{MMR.verify}(\kappa_V, c, \mu_V)$ . This implies a break of the honest-commitment soundness of the MMR.

- The chunks with bits that are flipped in **Line 17** of `MS2.remove` do not correspond to the actual SWBF because of a malicious membership proof in the SWBF MMR. This implies a break of the honest-commitment soundness of the MMR.

If none of these conditions are met, then  $A_1$  must be successful. This leads to a bound on the success probability of  $A_2$ :

$$\Pr[A_2 \checkmark] \leq \varsigma + 2 \cdot \text{Err}_{\text{hcsnd}}^{\text{MMR}} + \Pr[A_1 \checkmark] \quad (16)$$

$$\leq 2 \cdot \varsigma + 3 \cdot \text{InSec}_{\text{coll}}^{\text{H}} \quad (17)$$

□

**Theorem 7 (succinctness).** *MS<sub>2</sub> is succinct.*

*Proof.* By construction. All algorithms touch at most a number of nodes in the MMRs that scales logarithmically with the number of addition and removal records. □

Synchronization is definitely succinct when only additions are concerned because in this case the MMRs are only added to. When there can also be removal updates, the SWBF membership proofs must be updated to accommodate for modifications as well. This fact would undermine succinct synchronization if the modifications were arbitrary, but in the case of mutator sets the modifications induced by removal records are not arbitrary. The range of chunks that could be modified is contiguous and can even be small relative to a large number of chunks. Moreover, and specifically to a blockchain context, old UTXOs are less likely to be spent than new UTXOs, and so the older the chunk the less likely it is to change.

## 5.6 Unlinkable Mutator Sets

The security theorem statements enable a search for parameters in pursuit of a target security level. What they do not do is quantify the amount of privacy generated by a given set of parameters. That is the purpose of this section.

Within one batch, the addition records have identically distributed removal records. It may be a viable strategy to increase the step size  $s$  and decrease the active window size  $w$  until they meet. At this point, the windows of settable bits of distinct batches are disjoint and there is essentially one Bloom filter per batch.

However, when  $s < w$  it is possible that the set of set bits does not uniquely determine which batch a removal record must have come from. This situation is qualitatively more appealing because from the adversary's point of view the batches themselves are only probabilistically delineable. In other words, the adversary is not only incapable of assigning a given removal record (provided it has this property) to its batch index with certainty; but he is also incapable of listing a given record's would-be siblings assuming its batch index is given.

In general, the privacy increases with the active window size  $w$ . However, practice requires  $w$  be small enough to be stored on the intended machinery. A natural question that arises is, what is the effect of  $w$  on privacy? One way to answer this question is to plot a histogram of removal entropy (see [Section 3.5](#)).

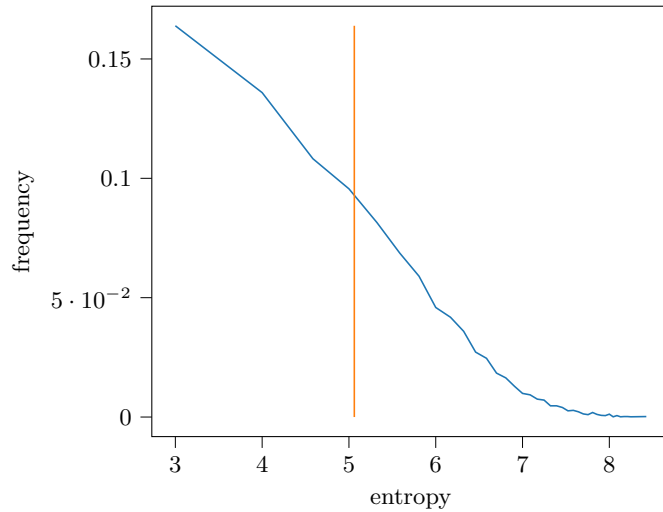


Fig. 4: Frequency of observed entropy.

**Figure 4** was generated for parameter set ( $w = 2^{20}, b = 8, s = 2^{12}, k = 45$ ) which incidentally achieves completeness error  $2^{-160}$ . Note that the batch size is  $b = 8$ , and as removal records are identically distributed with their batch-siblings, the entropy can never be lower than  $\log_2 8 = 3$ . However, as the entropy is frequently larger, the expectation (indicated by the orange line) is substantially higher. For these parameters, the expected entropy of removal records is slightly larger than 5, roughly equivalent to hiding perfectly behind 31 decoys.

## 5.7 Implementation Aspects

*Reversibility.* In a blockchain context, it is possible that what seems to be the consensus now is later overturned as a longer chain is disseminated. In this case nodes must revert the changes induced by the orphaned blocks and apply those induced by the new canonical chain tip. If the mutator set is part of the consensus rules, as described in [Section 4](#), some set bits need to be unset, but only if no other historical removal record set the bit in question. It is easy to support these reversions when the Bloom filter is invertible.

*Sparsity.* The calculation of the false positive probability in [Section 5.3](#) suggests that the optimal value of  $k$  in terms of  $s$  and  $b$  is  $k = \frac{s \cdot \ln 2}{b}$ . But  $k$  corresponds to the number of indices of set bits sampled pseudorandomly and, im-

portantly, whose integral sampling must be proved. Even if an arithmetization-oriented hash function is used for this purpose, it may still be desirable to reduce the number of required invocations of this primitive at the cost of other trade-offs. In particular, when  $k$  is smaller than this optimal value, the Bloom filter may be *sparse* — containing far more zeros than ones even in the inactive part. A straightforward way to represent this *sparse Bloom filter* efficiently is to store the indices of the set bits, and not the bit array itself. Note that this representation affects only performance and does not affect any calculations regarding security or completeness. For the parameters suggested in Section 5.6, the active window consists of at most  $\frac{w}{s} \cdot k = 11520$  integers of at most 20 bits, and on average half that. In our implementation this data structure is around 20 kiB in size.

*Batching.* Updates tend to come in batches rather than individually, for instance because a new block confirms a bunch of transactions rather than one at time. While it is possible to separate a batch of updates into its constituent elements and apply them individually, the update operations stand to benefit from batching them together and reusing intermediate calculations. There are compelling use cases for all permutations of the regular expression /update [((batch of)?(membership proof)), ((batch of)?(removal record)), (set commitment)] for(batch of)? [(addition), (removal)] (reversion)?/.

*Sender and Receiver Randomness.* As it is described in Sections 5.4 and 5.5, the randomizer  $r$  serves two purposes: *a*) it gives the canonical commitment, and thus the addition record, semantic security; and *b*) it randomizes the set of set bits when a removal record is generated. A non-interactive proof establishes that  $r$  links the removal record to the addition record, but as this proof is zero-knowledge, external observers do not learn this link. This construction suffices for exposition, but has two drawbacks. First, the party who chooses  $r$  can determine (with reasonable likelihood) the pattern of set bits, giving rise to maliciously increased false positive probability. Second, in a blockchain context, while the sender of a transaction cannot spend the UTXO (after all, the UTXO requires its own separate unlocking logic) he does know the pattern of set bits and can therefore observe when the receiver spends it. Both deficiencies are remedied with the following fix: the receiver selects a secret random preimage  $r_p \xleftarrow{\$} D$  and disseminates its hash digest  $r_d = H(r_p)$ , for instance as part of his receiving address. The sender samples his own secret randomness  $r_s \xleftarrow{\$} D$ . The canonical commitment is computed as  $\alpha = H(t||r_s||r_d)$ , but the bit indices are derived from  $H_w(l||t||r_s||r_p||i)$  and the zero-knowledge proof additionally establishes that  $r_p$  and  $r_d$  are correctly related. The sender communicates  $r_s$  via an off-chain protocol or uploads a public key encryption of this value to the blockchain such that only the intended receiver can read this secret.

## 6 Conclusion

Mutator sets capture a notion conceptually in between accumulation schemes and mixnets. While this exposition in this paper focused on the application to



blockchain protocols, mutator sets likely also apply to various other schemes and protocols that also rely on an accumulator or on a mixnet. For instance:

- Accumulation schemes can generate dynamic ring and group signature schemes. While there is perfect anonymity among signers that belong to the group, the group itself well defined even from the point of view of external observers. Using a mutator set instead, the group itself can be made fuzzy as external observers can no longer observe who joins and who leaves the group. Mutator sets can thus generate *anonymous group and ring signatures*.
- Accumulation schemes can generate dynamic cryptographically authenticated access control policies, or whitelists. As before, the set of users with access is well defined even from the external observer’s point of view. Using mutator sets, the whitelist can hide who is granted access and for whom it is revoked. In other words, mutator sets can generate *anonymous dynamic whitelists*.
- Mixnets are one of the two main constructions for private electronic voting; the other one being additively homomorphic encryption. In traditional mixnet-based voting, authorities are required to shuffle and rerandomize the set of encrypted votes. They are not trusted for integrity because they are required to prove it, but they are trusted for availability. Moreover, the secret decryption key must exist somewhere, even if it is spread across multiple shards. Using a mutator set instead of a mixnet addresses both of these deficiencies: voters can mix the votes themselves, and open their own committed votes. This sketch gives rise to *cryptographic voting without encryption*.

Our construction for mutator sets relies only on hash functions for security and in particular does not use any number-theoretic hard problems. While we do take the zero-knowledge proof for granted, we do note that there adequate proof systems that likewise rely only on hash functions for security. As a result, our construction is plausibly secure against quantum adversaries in addition to classical ones. An interesting open question remains what performance improvements or security tradeoffs come about as a result of migrating to a construction that does rely on number-theoretic hard problems, if such a construction can be found.

## References

1. Ayebie, E.B., Souidi, E.M.: New code-based cryptographic accumulator and fully dynamic group signature. *Des. Codes Cryptogr.* **90**(12), 2861–2891 (2022), <https://doi.org/10.1007/s10623-022-01007-5>
2. Baldimtsi, F., Camenisch, J., Dubovitskaya, M., Lysyanskaya, A., Reyzin, L., Samelin, K., Yakoubov, S.: Accumulators with applications to anonymity-preserving revocation. In: *IEEE EuroS&P*, 2017. pp. 301–315. IEEE (2017), <https://doi.org/10.1109/EuroSP.2017.13>
3. Baldimtsi, F., Karantaidou, I., Raghuraman, S.: Oblivious accumulators. *IACR Cryptol. ePrint Arch.* p. 1001 (2023), <https://eprint.iacr.org/2023/1001>

4. Baric, N., Pfitzmann, B.: Collision-free accumulators and fail-stop signature schemes without trees. In: Fumy, W. (ed.) EUROCRYPT '97. Lecture Notes in Computer Science, vol. 1233, pp. 480–494. Springer (1997). [https://doi.org/10.1007/3-540-69053-0\\_33](https://doi.org/10.1007/3-540-69053-0_33), [https://doi.org/10.1007/3-540-69053-0\\_33](https://doi.org/10.1007/3-540-69053-0_33)
5. Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., Virza, M.: Zerocash: Decentralized anonymous payments from bitcoin. In: IEEE SP 2014. pp. 459–474. IEEE Computer Society (2014). <https://doi.org/10.1109/SP.2014.36>, <https://doi.org/10.1109/SP.2014.36>
6. Benaloh, J.C., de Mare, M.: One-way accumulators: A decentralized alternative to digital signatures (extended abstract). In: Helleseht, T. (ed.) EUROCRYPT '93. Lecture Notes in Computer Science, vol. 765, pp. 274–285. Springer (1993). [https://doi.org/10.1007/3-540-48285-7\\_24](https://doi.org/10.1007/3-540-48285-7_24), [https://doi.org/10.1007/3-540-48285-7\\_24](https://doi.org/10.1007/3-540-48285-7_24)
7. Bloom, B.H.: Space/time trade-offs in hash coding with allowable errors. Commun. ACM **13**(7), 422–426 (1970). <https://doi.org/10.1145/362686.362692>, <https://doi.org/10.1145/362686.362692>
8. Bünz, B., Kiffer, L., Luu, L., Zamani, M.: Flyclient: Super-light clients for cryptocurrencies. In: IEEE SP 2020. pp. 928–946. IEEE (2020), <https://doi.org/10.1109/SP40000.2020.00049>
9. Camenisch, J., Kohlweiss, M., Soriente, C.: An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In: Jarecki, S., Tsudik, G. (eds.) PKC 2009. Lecture Notes in Computer Science, vol. 5443, pp. 481–500. Springer (2009). [https://doi.org/10.1007/978-3-642-00468-1\\_27](https://doi.org/10.1007/978-3-642-00468-1_27), [https://doi.org/10.1007/978-3-642-00468-1\\_27](https://doi.org/10.1007/978-3-642-00468-1_27)
10. Camenisch, J., Lysyanskaya, A.: Dynamic accumulators and application to efficient revocation of anonymous credentials. In: Yung, M. (ed.) Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings. Lecture Notes in Computer Science, vol. 2442, pp. 61–76. Springer (2002). [https://doi.org/10.1007/3-540-45708-9\\_5](https://doi.org/10.1007/3-540-45708-9_5), [https://doi.org/10.1007/3-540-45708-9\\_5](https://doi.org/10.1007/3-540-45708-9_5)
11. Chaum, D.: Untraceable electronic mail, return addresses, and digital pseudonyms. Commun. ACM **24**(2), 84–88 (1981), <https://doi.org/10.1145/358549.358563>
12. Danezis, G., Dingledine, R., Mathewson, N.: Mixminion: Design of a type III anonymous remailer protocol. In: IEEE S&P 2003. pp. 2–15. IEEE Computer Society (2003), <https://doi.org/10.1109/SECPRI.2003.1199323>
13. Derler, D., Hanser, C., Slamanig, D.: Revisiting cryptographic accumulators, additional properties and relations to other primitives. In: Nyberg, K. (ed.) CT-RSA 2015. Lecture Notes in Computer Science, vol. 9048, pp. 127–144. Springer (2015), [https://doi.org/10.1007/978-3-319-16715-2\\_7](https://doi.org/10.1007/978-3-319-16715-2_7)
14. Díaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: Dingledine, R., Syverson, P.F. (eds.) Privacy Enhancing Technologies PET 2002. Lecture Notes in Computer Science, vol. 2482, pp. 54–68. Springer (2002). [https://doi.org/10.1007/3-540-36467-6\\_5](https://doi.org/10.1007/3-540-36467-6_5), [https://doi.org/10.1007/3-540-36467-6\\_5](https://doi.org/10.1007/3-540-36467-6_5)
15. Dingledine, R., Mathewson, N., Syverson, P.F.: Tor: The second-generation onion router. In: Blaze, M. (ed.) USENIX 2004. pp. 303–320. USENIX (2004), <http://www.usenix.org/publications/library/proceedings/sec04/tech/dingledine.html>
16. Dryja, T.: Utreexo: A dynamic hash-based accumulator optimized for the Bitcoin UTXO set. Cryptology ePrint Archive (2019)

17. Fauzi, P., Meiklejohn, S., Mercer, R., Orlandi, C.: Quisquis: A new design for anonymous cryptocurrencies. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part I. Lecture Notes in Computer Science, vol. 11921, pp. 649–678. Springer (2019). [https://doi.org/10.1007/978-3-030-34578-5\\_23](https://doi.org/10.1007/978-3-030-34578-5_23), [https://doi.org/10.1007/978-3-030-34578-5\\_23](https://doi.org/10.1007/978-3-030-34578-5_23)
18. Friedenbach, M.: Compact spv proofs via block header commitments, <https://www.mail-archive.com/bitcoin-development@lists.sourceforge.net/msg04318.html>
19. Fuchsbauer, G., Orrù, M.: Non-interactive mumblewumble transactions, revisited. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part I. Lecture Notes in Computer Science, vol. 13791, pp. 713–744. Springer (2022), [https://doi.org/10.1007/978-3-031-22963-3\\_24](https://doi.org/10.1007/978-3-031-22963-3_24)
20. Fuchsbauer, G., Orrù, M., Seurin, Y.: Aggregate cash systems: A cryptographic investigation of mumblewumble. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. Lecture Notes in Computer Science, vol. 11476, pp. 657–689. Springer (2019), [https://doi.org/10.1007/978-3-030-17653-2\\_22](https://doi.org/10.1007/978-3-030-17653-2_22)
21. Jedusor, T.E.: Mumblewumble (2016), <https://docs.beam.mw/Mumblewumble.pdf>
22. Kattis, A., Bonneau, J.: Proof of necessary work: Succinct state verification with fairness guarantees. IACR Cryptol. ePrint Arch. p. 190 (2020), <https://eprint.iacr.org/2020/190>
23. Khedr, W.I., Khater, H.M., Mohamed, E.R.: Cryptographic accumulator-based scheme for critical data integrity verification in cloud storage. IEEE Access **7** (2019). <https://doi.org/10.1109/ACCESS.2019.2917628>
24. Kiayias, A., Lamprou, N., Stouka, A.: Proofs of proofs of work with sublinear complexity. In: Clark, J., Meiklejohn, S., Ryan, P.Y.A., Wallach, D.S., Brenner, M., Rohloff, K. (eds.) FC 2016. Lecture Notes in Computer Science, vol. 9604, pp. 61–78. Springer (2016), [https://doi.org/10.1007/978-3-662-53357-4\\_5](https://doi.org/10.1007/978-3-662-53357-4_5)
25. Kiayias, A., Miller, A., Zindros, D.: Non-interactive proofs of proof-of-work. In: Bonneau, J., Heninger, N. (eds.) FC 2020. Lecture Notes in Computer Science, vol. 12059, pp. 505–522. Springer (2020), [https://doi.org/10.1007/978-3-030-51280-4\\_27](https://doi.org/10.1007/978-3-030-51280-4_27)
26. Libert, B., Ling, S., Nguyen, K., Wang, H.: Zero-knowledge arguments for lattice-based accumulators: Logarithmic-size ring signatures and group signatures without trapdoors. In: Fischlin, M., Coron, J. (eds.) EUROCRYPT 2016, Part II. Lecture Notes in Computer Science, vol. 9666, pp. 1–31. Springer (2016), [https://doi.org/10.1007/978-3-662-49896-5\\_1](https://doi.org/10.1007/978-3-662-49896-5_1)
27. Libert, B., Ramanna, S.C., Yung, M.: Functional commitment schemes: From polynomial commitments to pairing-based accumulators from simple assumptions. In: Chatzigiannakis, I., Mitzenmacher, M., Rabani, Y., Sangiorgi, D. (eds.) ICALP 2016. LIPIcs, vol. 55, pp. 30:1–30:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016), <https://doi.org/10.4230/LIPIcs.ICALP.2016.30>
28. Merkle, R.C.: A digital signature based on a conventional encryption function. In: Pomerance, C. (ed.) CRYPTO '87. Lecture Notes in Computer Science, vol. 293, pp. 369–378. Springer (1987). [https://doi.org/10.1007/3-540-48184-2\\_32](https://doi.org/10.1007/3-540-48184-2_32), [https://doi.org/10.1007/3-540-48184-2\\_32](https://doi.org/10.1007/3-540-48184-2_32)
29. Miers, I., Garman, C., Green, M., Rubin, A.D.: Zerocoin: Anonymous distributed e-cash from bitcoin. In: IEEE SP 2013. pp. 397–411. IEEE Computer Society (2013). <https://doi.org/10.1109/SP.2013.34>, <https://doi.org/10.1109/SP.2013.34>
30. Miller, A.: The high value hashway, <https://bitcointalk.org/index.php?topic=98986.0>

31. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system. Metzdowd Cryptography Mailing List (2008)
32. Nguyen, L.: Accumulators from bilinear pairings and applications. In: Menezes, A. (ed.) CT-RSA 2005. Lecture Notes in Computer Science, vol. 3376, pp. 275–292. Springer (2005), [https://doi.org/10.1007/978-3-540-30574-3\\_19](https://doi.org/10.1007/978-3-540-30574-3_19)
33. Nyberg, K.: Fast accumulated hashing. In: Gollmann, D. (ed.) FSE 1996. Lecture Notes in Computer Science, vol. 1039, pp. 83–87. Springer (1996), [https://doi.org/10.1007/3-540-60865-6\\_45](https://doi.org/10.1007/3-540-60865-6_45)
34. Papamanthou, C., Tamassia, R., Triandopoulos, N.: Authenticated hash tables based on cryptographic accumulators. *Algorithmica* **74**(2), 664–712 (2016), <https://doi.org/10.1007/s00453-014-9968-3>
35. Ramabaja, L., Avdullahu, A.: The bloom tree (2020)
36. Saberhagen, N.V.: Cryptonote v 2.0 (2013), <https://www.getmonero.org/resources/research-lab/pubs/cryptonote-whitepaper.pdf>
37. Sako, K., Kilian, J.: Receipt-free mix-type voting scheme - A practical solution to the implementation of a voting booth. In: Guillou, L.C., Quisquater, J. (eds.) EUROCRYPT '95. Lecture Notes in Computer Science, vol. 921, pp. 393–403. Springer (1995), [https://doi.org/10.1007/3-540-49264-X\\_32](https://doi.org/10.1007/3-540-49264-X_32)
38. Todd, P.: Merkle mountain ranges, <https://github.com/opentimestamps/opentimestamps-server/blob/master/doc/merkle-mountain-range.md>