

# Scaling Verifiable Computation Using Efficient Set Accumulators

Alex Ozdemir

Riad S. Wahby

Dan Boneh

{aozdemir, rsw, dabo}@cs.stanford.edu

Stanford University

## Abstract

Verifiable outsourcing systems offload a large computation to a remote server, but require that the remote server provide a succinct proof, called a SNARK, that proves that the server carried out the computation correctly. Real-world applications of this approach can be found in several blockchain systems that employ verifiable outsourcing to process a large number of transactions off-chain. This reduces the on-chain work to simply verifying a succinct proof that transaction processing was done correctly. In practice, verifiable outsourcing of state updates is done by updating the leaves of a Merkle tree, recomputing the resulting Merkle root, and proving using a SNARK that the state update was done correctly.

In this work, we use a combination of existing and novel techniques to implement an RSA accumulator inside of a SNARK, and use it as a replacement for a Merkle tree. We specifically optimize the accumulator for compatibility with SNARKs. Our experiments show that the resulting system can dramatically reduce costs compared to existing approaches that use Merkle trees for committing to the current state. These results apply broadly to any system that needs to offload batches of state updates to an untrusted server.

## 1 Introduction

Verifiable outsourcing [4, 13, 15, 16, 21, 32, 45, 47, 49, 52, 56, 60, 76, 77, 92, 101–103, 106–110, 116, 118, 119] is a technique that enables a weak client to outsource a computation to a powerful server. The server returns the result of the computation along with a proof that the computation was done correctly. The proof must be succinct, which means that it must be short and cheap to verify. Verifiable outsourcing is relevant in a number of scenarios, including weak IoT devices, wearables, and low-power devices.

More recently, verifiable outsourcing has been deployed in blockchain environments, because on-chain work is expensive—literally. Here, a batch of  $k$  transactions, say  $k = 1000$ , is outsourced to an untrusted server, called an *aggregator*, for processing. The aggregator (1) verifies that the transactions are valid (e.g., properly signed), (2) computes the updated global state resulting from these transactions, and

(3) generates a succinct proof that the aggregator correctly executed steps (1) and (2). The updated state and the succinct proof are then sent to the blockchain. In this approach, the (expensive) on-chain work is reduced to only verifying the proof—which is fast, taking time independent of the number of transactions  $k$ —and then recording the updated state. Example systems that operate this way include Rollup [7], Coda [86], Matter [83], and Zexe [29].

The process described above is called **verifiable outsourcing of state update** [32]. In more detail, the state is a set of elements  $S = \{x_1, \dots, x_M\}$  from some universe  $\mathcal{X}$ . The blockchain (or a low-power device) stores only a succinct digest of  $S$ , e.g., the root of a Merkle tree whose leaves comprise the elements of  $S$ . The untrusted but powerful aggregator stores the full set  $S$ , in the clear. (Note that we treat  $S$  as public data—privacy is orthogonal to our goal, which is scalability). When processing a batch of transactions as described above, the aggregator updates  $S$  to produce a new set  $S'$ , then computes a new Merkle digest for  $S'$  that it sends to the blockchain to be verified and recorded. The aggregator’s proof establishes that its starting state  $S$  is consistent with the current digest, that correctly applying transactions yields the ending state  $S'$ , and that the new digest is consistent with  $S'$ .

The succinct proof needed here is called a *SNARK* [19], which we define in more detail in the next section. Constructing efficient SNARKs and optimizing their implementation is a very active area of research [13, 15, 16, 49, 63, 68, 92], with several new systems just in the last year [11, 37, 43, 44, 61, 62, 82, 117]. A common thread in all of these systems is that the proving costs are enormous. In particular, proving imposes multiple-orders-of-magnitude slowdown compared to native execution [92, 101, 111]; this can be defrayed via parallel execution, e.g., in clusters [45, 116] or on GPUs [103, 107].

Perhaps more importantly, for widely deployed SNARKs, proving correctness of large computations requires an amount of RAM *proportional to the computation’s execution time* [16, 92]. The result is that, even when proving is distributed across hundreds of workers, the largest reachable computation sizes are relatively small: only about 2 billion steps [116]. This imposes a strict upper bound on the number of transactions  $k$  that can be processed in a single batch.

This state of affairs has motivated a large body of work on computational primitives that yield efficient proofs. Examples include arithmetic [77, 92, 103], control flow [92, 103, 111], persistent state [4, 32, 49, 56, 100], and random-access memory [12, 13, 16, 32, 77, 111]. Our work continues in this vein, with a focus on reducing proving costs for computations involving persistent state or random-access memory.

**Our work.** A Merkle tree [87] is an example of an *accumulator* [17], a cryptographic primitive that lets one commit to a set  $S$ , and later prove that an element  $x$  is a member of  $S$ . Although Merkle trees are used pervasively in today’s general-purpose verifiable state update applications, in this work we show that a Merkle tree is not the best choice for large batches of state updates when  $S$  is moderately to very large, say  $|S| \geq 2^{10}$ . In particular, we show that replacing Merkle trees with RSA-based accumulators [24, 40, 78] significantly improves proving time and/or reachable computation size. Our contributions are:

- We define a new operation for RSA accumulators, which we call MultiSwap, that provides a precise sequential semantics for batched verifiable state updates (§3).
- We synthesize existing and novel techniques for efficiently implementing MultiSwap (and, more generally, RSA accumulators) in the context of SNARKs (§4). These techniques include a hash function that outputs provable prime numbers, and a new division-intractable hash function. Our approach makes use of very recent advances in manipulating RSA accumulators [24].
- We apply our techniques in two contexts (§5). The first, called Rollup [7], is a technique for batching cryptocurrency transactions off-chain in order to save on-chain work. The second is a general-purpose RAM abstraction with long-lived state (i.e., over many proofs), which builds upon and improves prior work [12, 13, 16, 32, 111].
- We implement and evaluate (§6, §7). In particular, we compare our RSA accumulator implementation to Merkle trees in two benchmarks: one that measures only set operations, and one that implements a Rollup-style distributed payment application. We also compare our RAM abstraction with existing work via a cost model analysis.

In the set operations benchmark, we find that RSA accumulators surpass  $2^{20}$ -element Merkle trees for batches of  $\approx 300$  or more operations, and allow for  $6.1\times$  more operations to be performed in the largest proof sizes we consider. In the Rollup application, RSA accumulators surpass  $2^{20}$ -element Merkle trees for  $\approx 150$  transactions, and allow  $2.8\times$  more transactions in the largest proofs. For RAM, we find that for a RAM of size  $2^{20}$ , RSA accumulators surpass Merkle trees for  $\approx 300$ – $600$  accesses, depending on write load.

## 2 Background and definitions

**Multisets.** A *multiset* is an unordered collection that may contain multiple copies of any element.  $S_1 \uplus S_2$  denotes the union of multisets  $S_1$  and  $S_2$ , i.e., the multiset  $S_3$  where each element  $x \in S_3$  has multiplicity equal to the sum of the multiplicities of  $x$  in  $S_1$  and  $S_2$ .  $S_1 \boxminus S_2$  denotes the strict difference of multisets  $S_1$  and  $S_2$ , i.e., the multiset  $S_3$  where each element  $x \in S_3$  has multiplicity equal to the difference of multiplicities of  $x$  in  $S_1$  and  $S_2$ . Note that  $S_1 \boxminus S_2$  is only defined if  $S_2 \subseteq S_1$ .

**RSA groups.** An *RSA group* is the group  $\mathbb{Z}_N^\times$ , i.e., the multiplicative group of invertible integers modulo  $N$ , where  $N$  is the product of two secret primes. We define the *RSA quotient group* for  $N$  as the group  $\mathbb{Z}_N^\times / \{\pm 1\}$ . In this group, the elements  $x$  and  $N - x$  are the same, meaning that all elements can be represented by integers in the interval  $[1, \lfloor N/2 \rfloor]$ . It is believed that this group has no element of known order, other than the identity.

**Proofs and arguments.** Informally, a *proof* is a protocol between a prover  $\mathcal{P}$  and a PPT verifier  $\mathcal{V}$  by which  $\mathcal{P}$  convinces  $\mathcal{V}$  that  $\exists v : \mathfrak{R}(t, v) = 1$ , for a relation  $\mathfrak{R}$ ,  $t$  an input from  $\mathcal{V}$ , and  $v$  a (possibly empty) *witness* from  $\mathcal{P}$ . A proof satisfies the following properties:

- *Completeness:* If  $\exists v : \mathfrak{R}(t, v) = 1$ , then an honest  $\mathcal{P}$  convinces  $\mathcal{V}$  except with probability at most  $\epsilon_c \ll 1/2$ .
- *Soundness:* If  $\nexists v : \mathfrak{R}(t, v) = 1$ , no cheating prover  $\mathcal{P}^*$  convinces  $\mathcal{V}$  except with probability at most  $\epsilon_s \ll 1/2$ .

If soundness holds only against PPT  $\mathcal{P}^*$ , this protocol is instead called an *argument*. When the witness  $v$  exists, one may also require the proof system to provide *knowledge soundness*. Informally this means that whenever  $\mathcal{P}$  convinces  $\mathcal{V}$  that  $\exists v : \mathfrak{R}(t, v) = 1$ ,  $v$  exists and  $\mathcal{P}$  “knows” a witness  $v$  (slightly more formally, there exists a PPT algorithm, an *extractor*, that can produce a witness via oracle access to  $\mathcal{P}$ ).

**Proof of exponentiation.** Let  $\mathbb{G}$  be a finite group of unknown order. Wesolowski [115] describes a protocol that allows  $\mathcal{P}$  to convince  $\mathcal{V}$  that  $y = x^n$  in  $\mathbb{G}$ , namely a protocol for the relation  $\mathfrak{R}$  given by  $\mathfrak{R}((n, x, y), \cdot) = 1 \iff y = x^n \in \mathbb{G}$ . The protocol is: on input  $(n, x, y)$ ,  $\mathcal{V}$  sends to  $\mathcal{P}$  a random  $\ell$  chosen from the first  $2^k$  primes.  $\mathcal{P}$  sends back  $Q = x^{\lfloor n/\ell \rfloor} \in \mathbb{G}$ , and  $\mathcal{V}$  accepts only if  $Q^\ell \cdot x^{n \bmod \ell} = y \in \mathbb{G}$  holds.

This protocol is complete by inspection. Wesolowski shows that it is sound if the group  $\mathbb{G}$  satisfies the *adaptive root assumption*, roughly, it is infeasible for an adversary to find a random root of an element of  $\mathbb{G}$  chosen by the adversary. The RSA quotient group  $\mathbb{Z}_N^\times / \{\pm 1\}$  is conjectured to satisfy this assumption when  $\mathcal{P}$  cannot factor  $N$  [23].

**Division-intractable hashing.** Recall that a hash function  $H : \mathcal{X} \rightarrow \mathcal{D}$  is collision resistant if it is infeasible for a PPT adversary to find distinct  $x_0, x_1$  such that  $H(x_0) = H(x_1)$ . Informally,  $H$  is *division intractable* if the range of  $H$  is  $\mathbb{Z}$ , and it is infeasible for a PPT adversary to find  $\hat{x}$  and a set  $\{x_i\}$  in

$X$  such that  $\hat{x} \notin \{x_i\}$  and  $H(\hat{x})$  divides  $\prod_i H(x_i)$ . A collision-resistant hash function that outputs prime numbers is division intractable. We construct a different division intractable hash function in Section 4.2.

**Pocklington primality certificates.** Let  $p$  be a prime, and  $r < p$  and  $a$  be positive integers. Define  $p' = p \cdot r + 1$ . Pocklington’s criterion [34] states that if  $a^{p' \cdot r} \equiv 1 \pmod{p'}$  and  $\gcd(a^r - 1, p') = 1$ , then  $p'$  is prime. In this case, we say that  $(p, r, a)$  is a *Pocklington witness* for  $p'$ .

Pocklington’s criterion is useful for constructing primality certificates. For a prime  $p_n$ , this certificate comprises

$$(p_0, \{(r_i, a_i)\}_{0 < i \leq n})$$

where  $p_i = p_{i-1} \cdot r_i + 1$ . To check this certificate, first verify the primality of the small prime  $p_0$  (e.g., using a deterministic primality test), then verify the Pocklington witness  $(p_{i-1}, r_i, a_i)$  for  $p_i$ ,  $0 < i \leq n$ . If each  $r_i$  is nearly as large as  $p_i$ , the bit lengths double at each step, meaning that the total verification cost is dominated by the cost of the final step.

## 2.1 Accumulators

A *cryptographic accumulator* [17] commits to a collection of values (e.g., a vector, set, or multiset) as a succinct digest. This digest is *binding*, meaning informally that it is computationally infeasible to equivocate about the collection represented by the digest. In addition, accumulators admit succinct proofs of membership and, in some cases, non-membership.

**Merkle trees.** The best-known vector accumulator is the Merkle tree [87]. To review, this is a binary tree that stores a vector in the labels of its leaves; the label associated with an internal node of this tree is the result of applying a collision-resistant hash  $H$  to the concatenation of the children’s labels; and the digest representing the collection is the label of the root node.

A membership proof for the leaf at index  $i$  is a path through the tree, i.e., the labels of the siblings of all nodes between the purported leaf and the root. Verifying the proof requires computing the node labels along the path and comparing the final value to the digest (the bits of  $i$  indicate whether each node is the right or left child of its parent). Updating a leaf’s label is closely related: given a membership proof for the old value, the new digest is computed by swapping the old leaf for the new one, then computing the hashes along the path. Merkle trees do not support succinct non-membership proofs.

The cost of verifying  $k$  membership proofs for a vector comprising  $2^m$  values is  $k \cdot m$  evaluations of  $H$ . The cost of  $k$  leaf updates is  $2 \cdot k \cdot m$  evaluations. Membership proofs and updates cannot be batched for savings.

**RSA accumulators.** The RSA multiset accumulator [40, 78] represents a multiset  $S$  with the digest

$$\llbracket S \rrbracket = g^{\prod_{s \in S} H(s)} \in \mathbb{G},$$

where  $g$  is a fixed member of an RSA quotient group  $\mathbb{G}$  and  $H$  is a division-intractable hash function (§2). Inserting a new element  $s$  into  $S$  thus requires computing  $\llbracket S \rrbracket^{H(s)}$ .

To prove membership of  $s \in S$ , the prover furnishes the value  $\pi = \llbracket S \rrbracket^{1/H(s)}$ , i.e., a  $H(s)$ ’th root of  $\llbracket S \rrbracket$ . This proof is verified by checking that  $\pi^{H(s)} = \llbracket S \rrbracket$ .

Non-membership proofs are also possible [78], leveraging the fact that  $s' \notin S$  if and only if  $\gcd(H(s'), \prod_{s \in S} H(s)) = 1$ . This means that the Bézout coefficients  $a, b$ , i.e., integers satisfying

$$a \cdot H(s') + b \cdot \prod_{s \in S} H(s) = 1$$

are a non-membership witness, since the above implies that

$$\llbracket S \rrbracket^b \cdot (g^{a'})^{H(s')} = g$$

Because  $a$  is large and  $b$  is small, the proof  $(g^a, b)$  is succinct.

Insertions, membership proofs, and non-membership proofs can all be batched [24] via Wesolowski proofs (§2). For example, since  $\llbracket S \uplus \{s_i\} \rrbracket = \llbracket S \rrbracket^{\prod_i s_i}$ , computing an updated digest directly requires an exponentiation by  $\prod_i s_i$ . In contrast, checking the corresponding proof only requires computing and then exponentiating by  $\prod_i s_i \pmod{\ell}$ , for  $\ell$  a prime of less than 200 bits. This means that the exponentiation (but not the multiplication) to verify a batch proof has constant size.

## 2.2 Verifiable computation and SNARKs

Several lines of built systems [13, 15, 16, 21, 32, 47, 49, 60, 77, 92, 101–103, 107, 108, 119] enable the following high-level model.<sup>1</sup> A verifier  $\mathcal{V}$  asks a prover  $\mathcal{P}$  to convince it that  $y = \Psi(x)$ , where  $\Psi$  is a program taking input  $x$  and returning output  $y$ . To do so,  $\mathcal{P}$  produces a short certificate that the claimed output is correct. Completeness holds with  $\epsilon_c = 0$ ; soundness holds as long as  $\mathcal{P}$  is computationally bounded, with  $\epsilon_s$  negligible in a security parameter (§2).

Roughly speaking, these systems comprise two parts. In the *front-end*,  $\mathcal{V}$  compiles  $\Psi$  into a system of equations  $C(X, Y, Z)$ , where  $X, Y$ , and  $Z$  are (vectors of) formal variables.  $\mathcal{V}$  constructs  $C$  such that  $z$  satisfying  $C(X = x, Y = y, Z = z)$  exists (that is, the formal variable  $X$  is bound to the value  $x$ , and so on) if and only if  $y = \Psi(x)$ . The *back-end* comprises cryptographic and complexity-theoretic machinery by which  $\mathcal{P}$  convinces  $\mathcal{V}$  that a witness  $z$  exists for  $X = x$  and  $Y = y$ .

This paper focuses on compilation in the front-end. We target back-ends derived from GGPR [63] via Pinocchio [92] (including [15, 16, 68]), which we briefly describe below. Our work is also compatible with other back-ends, e.g., Zatar [101], Ligerio [2], Bulletproofs [36], Sonic [82], and Aurora [14].<sup>2</sup>

<sup>1</sup>The description in this section owes a textual and notational debt to the description in Buffet [111], which works in the same model.

<sup>2</sup>We do not target STARK [11] (which uses a different  $C$  representation) or systems built on GKR [65] and CMT [47], e.g., vRAM [119], Hyrax [112], and Libra [117] (which restrict  $C$  in ways this work does not comprehend).

GGPR, Pinocchio and their derivatives instantiate *zero-knowledge Succinct Non-interactive ARguments of Knowledge with preprocessing* (zkSNARKs), which are argument protocols satisfying completeness, knowledge soundness, and zero knowledge (§2),<sup>3</sup> where knowledge soundness and zero knowledge apply to the assignment to  $Z$ . In addition, these protocols satisfy *succinctness*: informally, proof length and verification time are both sublinear in  $|C|$  (here, proofs are of constant size, while  $\mathcal{V}$ 's work is  $O(|X| + |Y|)$ ). These protocols include a preprocessing phase, in which  $\mathcal{V}$  (or someone that  $\mathcal{V}$  trusts) processes  $C$  to produce a *structured reference string* (SRS), which is used by  $\mathcal{P}$  to prove and  $\mathcal{V}$  to verify. The cost of the preprocessing phase and the length of the SRS are  $O(|C|)$ . The cost of the proving phase is  $O(|C| \log |C|)$  in time and  $O(|C|)$  in space (i.e., prover RAM).

The system of equations  $C(X, Y, Z)$  is a *rank-1 constraint system* (R1CS) over a large finite field  $\mathbb{F}_p$ . An R1CS is defined by three matrices,  $A, B, C \in \mathbb{F}_p^{|C| \times (1 + |X| + |Y| + |Z|)}$ . Its satisfiability is defined as follows: for  $W$  the column vector of formal variables  $[1, X, Y, Z]^T$ ,  $C(X, Y, Z)$  is the system of  $|C|$  equations  $(A \cdot W) \circ (B \cdot W) = C \cdot W$ , where  $\circ$  denotes the Hadamard (element-wise) product. In other words, an R1CS  $C$  is a conjunction of  $|C|$  constraints in  $|X| + |Y| + |Z|$  variables, where each constraint has the form “linear combination times linear combination equals linear combination.”

These facts outline a computational setting whose costs differ significantly from those of CPUs. On a CPU, bit operations are cheap and word-level arithmetic is slightly more costly. In an R1CS, addition is free, word-level multiplication has unit cost, and bitwise manipulation and many inequality operations are expensive; details are given below.

**Compiling programs to constraints.** A large body of prior work [13, 16, 32, 77, 92, 101–103, 110, 111] deals with efficiently compiling from programming languages to constraints.

An important technique for non-arithmetic operations is the use of *advice*, variables in  $Z$  whose values are provided by the prover. For example, consider the program fragment  $x \neq 0$ , which cannot be concisely expressed in terms of rank-1 constraints. Since constraints are defined over  $\mathbb{F}_p$ , this assertion might be rewritten as  $X^{p-1} = 1$ , which is true just when  $X \neq 0$  by Fermat’s little theorem. But this is costly: it requires  $O(\log p)$  multiplications. A less expensive way to express this constraint is  $Z \cdot X = 1$ ; the satisfying assignment to  $Z$  is  $X^{-1} \in \mathbb{F}_p$ . Since every element of  $\mathbb{F}_p$  other than 0 has a multiplicative inverse, this is satisfiable just when  $X \neq 0$ .

Comparisons, modular reductions, and bitwise operations make heavy use of advice from  $\mathcal{P}$ . For example, the program fragment  $y = x1 \ \& \ x2$ , where  $x1$  and  $x2$  have bit width  $b$  and

$\&$  is bitwise AND, is represented by the following constraints:

$$\begin{aligned} Z_{1,0} + 2 \cdot Z_{1,1} + \dots + 2^{b-1} \cdot Z_{1,b-1} &= X_1 \\ Z_{2,0} + 2 \cdot Z_{2,1} + \dots + 2^{b-1} \cdot Z_{2,b-1} &= X_2 \\ Z_{3,0} + 2 \cdot Z_{3,1} + \dots + 2^{b-1} \cdot Z_{3,b-1} &= Y \\ Z_{1,0} \cdot (1 - Z_{1,0}) &= 0 \\ &\dots \\ Z_{1,b-1} \cdot (1 - Z_{1,b-1}) &= 0 \\ Z_{2,0} \cdot (1 - Z_{2,0}) &= 0 \\ &\dots \\ Z_{2,b-1} \cdot (1 - Z_{2,b-1}) &= 0 \\ Z_{1,0} \cdot Z_{2,0} &= Z_{3,0} \\ &\dots \\ Z_{1,b-1} \cdot Z_{2,b-1} &= Z_{3,b-1} \end{aligned}$$

Here, the variables  $Z_{1,0} \dots Z_{1,b-1}$  contain a purported bitwise expansion of  $X_1$ , and likewise  $Z_{2,0} \dots Z_{2,b-1}$  and  $Z_{3,0} \dots Z_{3,b-1}$  for  $X_2$  and  $Y$ , respectively. The first three constraints ensure that the assignment to  $Z$  meets this requirement provided that each  $Z_{i,j}$  is assigned either 0 or 1; the remaining constraints ensure the latter. This operation is known as *bit splitting*; its cost for a  $b$ -bit value is  $b + 1$ , so the above program fragment costs  $3 \cdot b + 3$  constraints in total. Comparisons and modular reductions also require bit splitting.

Compiling conditionals to constraints requires expanding all branches into their corresponding constraints and selecting the correct result. Loops are similar; loop bounds must be statically known. For example, the program fragment

```
if (x1 != 0) { y = x2 + 1 } else { y = x2 * 3 }
```

compiles to the constraints

$$\begin{aligned} Z_1 \cdot X_1 &= Z_2 & (1) \\ Z_3 \cdot (Z_2 - 1) &= 0 & (2) \\ (1 - Z_3) \cdot X_1 &= 0 & (3) \\ (1 - Z_3) \cdot (Y - X_2 - 1) &= 0 & (4) \\ Z_3 \cdot (Y - 3 \cdot X_2) &= 0 & (5) \end{aligned}$$

This works as follows: if  $X_1 = 0$ ,  $Z_2 = 0$  by (1), so  $Z_3 = 0$  by (2) and  $Y = X_2 + 1$  by (4). Otherwise,  $Z_3 = 1$  by (3), so  $Z_2 = 1$  by (2),  $Z_1 = X_1^{-1}$  by (1), and  $Y = 3 \cdot X_2$  by (5).

**Multiprecision arithmetic.** xJsnark [77] describes techniques for compiling multiprecision arithmetic to efficient constraint systems. In brief, large integers are represented as a sequence of *limbs* in  $\mathbb{F}_p$ . The limb width,  $b_l$ , is defined such that a  $b$ -bit number  $a$  is represented as  $\eta = \lceil b/b_l \rceil$  limbs  $\{\hat{a}_i\}$ , where  $a = \sum_{i=0}^{\eta-1} \hat{a}_i \cdot 2^{b_l \cdot i}$ . For correctness, the compiler must track the maximum value of each number and ensure that  $C$  contains constraints that encode a sufficient number of limbs.

<sup>3</sup>We do not explicitly target zero-knowledge applications in this work, but our techniques may also be useful in that setting.

Multiprecision operations rely heavily on advice from  $\mathcal{P}$ . At a high level,  $\mathcal{P}$  supplies the result of a multiplication or addition, and the compiler emits constraints to check that result. Subtractions and divisions are checked by verifying the inverse addition or multiplication, respectively. xJsnark describes a range of optimizations that reduce the required number of constraints. We leave details to [77], because they are not necessary to understand our further optimizations (§4.3).

### Random-access memory

Programs that make use of RAM—in particular, programs whose memory accesses depend on the input, and thus cannot be statically analyzed—present a challenge for compiling to constraints. Prior work demonstrates three solutions. We now describe each, and compare costs and functionality below.

**Linear scan.** The most direct approach to emulating RAM in constraints is to perform a linear scan [77, 92]. Concretely,  $Y = \text{LOAD}(Z)$  compiles to a loop that scans through an array, comparing the loop index to  $Z$  and, if they match, setting  $Y$  to the corresponding value. (STORE is analogous.)

**The Pantry approach.** In Pantry [32], the authors borrow a technique from the memory-checking literature [20] based on Merkle trees [87] (see also §2.1). In particular, Pantry stores the contents of RAM in the leaves of a Merkle tree whose root serves as ground truth for the state of memory.

For a LOAD,  $\mathcal{P}$  furnishes advice comprising a purported value from memory, plus a Merkle path authenticating that value. The corresponding constraints encode verification of the Merkle path, i.e., a sequence of hash function invocations and an equality check against the Merkle root. For a STORE,  $\mathcal{P}$  furnishes, and the constraints verify, the same values as for a LOAD. In addition, the constraints encode a second sequence of hash function invocations that compute a new Merkle root corresponding to the updated memory state.

**The BCGT approach.** Ben-Sasson et al. [12] introduce, and other work [13, 16, 77, 111] refines, an approach building on the observation [3] that one can check a sequence of RAM operations using an *address-ordered transcript*, i.e., the sequence of RAM operations sorted by address accessed, breaking ties by execution order. In such a transcript, each LOAD is preceded either by the corresponding STORE or by another LOAD from the same address; correctness of RAM dictates that this LOAD should return the same value as the preceding operation. (A LOAD from an address to which no value was previously stored returns a default value, say, 0.)

Leveraging this observation, correctness of memory operations is compiled to constraints as follows. First, every access to memory appends a tuple  $(\text{IDX}_i, \text{OP}_i, \text{ADDR}_i, \text{DATA}_i)$  to an *execution-ordered transcript*; here,  $\text{IDX}_i = i$  is the index of the memory operation and  $\text{OP}_i$  is either LOAD or STORE. Then  $\mathcal{P}$  furnishes a purported address-ordered transcript  $\mathcal{T}$ , and

the constraints check its correctness by ensuring that (1) transcript  $\mathcal{T}$  is a permutation of the execution-ordered transcript, (2) each sequential pair of entries in transcript  $\mathcal{T}$  is indeed correctly ordered, and (3) each sequential pair of entries in transcript  $\mathcal{T}$  is *coherent*, i.e., each LOAD returns the value of the previous STORE (or the default if no such STORE exists). Check (1) is implemented with a *routing network* [18, 113].

**Costs and functionality.** Roughly speaking, for tiny memories linear scan is cheapest; otherwise, BCGT-style RAM is.<sup>4</sup> In more detail, assume a memory of size  $2^m$ , accessed  $k$  times. For a linear scan, each RAM operation costs  $O(2^m)$  constraints. (i.e.,  $2^m$  copies of constraints encoding conditional assignment). For Pantry, each LOAD entails  $m$  copies of constraints encoding a collision-resistant hash function and each STORE entails  $2m$  such copies, where such hash functions entail a few hundred to a few thousand constraints (§6; [15, 32, 77]). For BCGT, each RAM operation costs  $O(\log k)$  constraints for the routing network,  $O(m)$  constraints for address comparison, and  $O(1)$  constraints for coherence checking, all with good constants [111, Fig. 5].

Although Pantry-style RAM is costly, it offers functionality that the other two do not: the ability to pass the full state of a large RAM from one computation to another. Pantry accomplishes this by including in  $X$  the Merkle root corresponding to the initial RAM state; this has constant size (usually one element of  $\mathbb{F}_p$ ). In contrast, BCGT and linear scan would both require  $2^m$  values in  $X$  for a  $2^m$ -sized RAM; as discussed above, this would incur  $2^m$  cost for  $\mathcal{V}$  in verification. (Prior work [15, 16] uses this approach to *partially* initialize RAM.)

## 3 Swap sequences via batched operations

In this section, we define a new primitive, which we call MultiSwap, that exposes a sequential update semantics for RSA accumulators (§2.1). MultiSwap takes an accumulator and a list of pairs of elements, removing the first element from each pair and inserting the second. The key property of this primitive is that it is defined in terms of *batched* insertions and removals. In Section 4, we show how these batched operations are efficiently implemented as a system of constraints (§2.2).

In more detail, let  $S$  and  $S'$  be multisets, and let  $(x_1, y_1), \dots, (x_n, y_n)$  be a sequence of operations, called *swaps*, that replaces each  $x_i$  by  $y_i$  *in order*:  $(x_1, y_1)$  applied to  $S$  produces some new set  $S_1 = S \boxminus \{x_1\} \uplus \{y_1\}$ ; then  $(x_2, y_2)$  applied to  $S_1$  produces  $S_2 = S_1 \boxminus \{x_2\} \uplus \{y_2\}$ , etc. Our goal is to verify that when the above sequence is applied to  $S$ , the result is  $S' = S_n$ . Recall from Section 2.1 that RSA accumulators admit efficient *batched* insertions (deletions are analogous; §4). Our question is: how can we use this *un-ordered* primitive to implement one with *ordered* semantics?

<sup>4</sup>An exception is a computation with an enormous number of memory accesses where Pantry would win. But the number of accesses to reach this asymptote is well beyond the reach of practical proof systems.

Consider the following naïve solution: first verify the deletions, then verify the insertions. In other words, verify that there exists some  $S_{\text{mid}}$  such that  $S \boxminus \{x_i\} = S_{\text{mid}}$  and  $S_{\text{mid}} \boxplus \{y_i\} = S'$ . The problem with this approach is that it does not permit certain valid sequences, i.e., those in which a later swap deletes an item inserted by an earlier swap. (To see why, notice that  $S_{\text{mid}}$  only exists if all  $x_i \in S$ .)

Instead, our solution first verifies all the insertions, and then verifies all the deletions, irrespective of the order in which the operations are listed. In other words, it verifies the predicate

$$\exists S_{\text{mid}} : S \boxplus \{y_i\} = S_{\text{mid}} \quad \wedge \quad S_{\text{mid}} \boxminus \{x_i\} = S' \quad (6)$$

(Note that  $S_{\text{mid}} \boxminus \{x_i\} = S'$  is equivalent to  $S' \boxplus \{x_i\} = S_{\text{mid}}$ .) Intuitively, Equation (6) holds just when each element of an unordered multiset of swaps  $\{(x_i, y_i)\}$  can be applied to  $S$  in some order to produce  $S'$ . As we discuss below, this multiset may include *cycles*, subsets of swaps that have no net effect.

We now give a precise semantics for MultiSwap. Let  $\text{MultiSwap}(S, \sigma, S')$  denote the predicate that holds just when Equation (6) is satisfied. Let  $\sigma$  denote an unordered multiset of swaps  $\{(x_i, y_i)\}$ . A swap  $(x_i, y_i)$  is *valid* for  $S^*$  if  $x_i \in S^*$ . We say that  $\sigma$  is *sequentially consistent* with respect to  $S$  if there exists some ordering on  $\sigma$  such that all swaps are valid when applied in that order starting from  $S$ . Furthermore, we say that  $\sigma$  *produces  $S'$  from  $S$*  if  $S'$  is the product of such an application order to  $S$ , and we say that  $\sigma^c$  is a *cycle* if it comprises  $\{(c_0, c_1), (c_1, c_2), \dots, (c_n, c_0)\}$ .

**Lemma 1.** *MultiSwap( $S, \sigma, S'$ ) holds if and only if there exist any number of cycles  $\sigma_i^c$  and cycle-free  $\sigma' \subseteq \sigma$  such that  $\sigma = \sigma' \boxplus \boxplus_i \sigma_i^c$ ,  $\sigma'$  is sequentially consistent with respect to  $S$ , and  $\sigma'$  produces  $S'$  from  $S$ .*

The proof of Lemma 1 is in Appendix A. Section 5 applies MultiSwap to problems that need sequential semantics for batched verifiable state updates.

## 4 Batched operations from constraints

In the previous section we described how the MultiSwap primitive is built from batched insertions and removals. In this section we describe these batched operations, the primitives that they are built on, and how those primitives are implemented as a set of constraints  $\mathcal{C}$  (§2.2).

Recall (§2.1) that RSA accumulators support batched insertions through an interactive protocol whose final check is

$$Q^\ell \cdot \llbracket S \rrbracket^{\prod_i H_\Delta(y_i) \bmod \ell} = \llbracket S \boxplus \{y_i\} \rrbracket \quad (7)$$

where  $\llbracket \cdot \rrbracket$  denotes a digest;  $S$ , the initial multiset;  $\ell$ , a random prime challenge;  $\{y_i\}$ , the inserted elements;  $H_\Delta$ , a division-intractable hash function; and  $Q$ , a witness from  $\mathcal{P}$ . Removing elements  $\{x_i\}$  is similar, except that  $S \boxminus \{x_i\}$  is regarded as the initial multiset and  $S$  the final one.<sup>5</sup>

<sup>5</sup>Proofs of non-membership (§2.1) use similar primitives; we do not discuss them in detail because they are not necessary for MultiSwap.

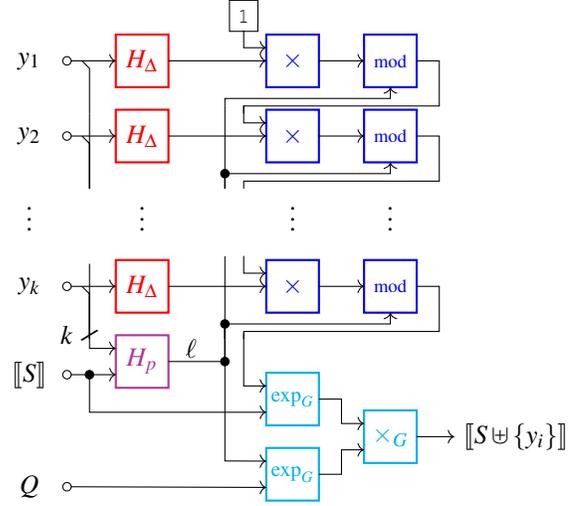


Figure 1: Insertion proof verification procedure (§4), which checks that  $Q$  is a valid Wesolowski proof (§2) for the exponentiation  $\llbracket S \rrbracket^{\prod_i H_\Delta(y_i)}$  on challenge  $\ell$ . To do so, it computes  $\ell = H_p(y_1, \dots, y_k)$  (purple box, bottom left), computes  $\prod_i H_\Delta(y_i) \bmod \ell$  (red and blue boxes, top), and then computes the LHS of the verification equation (cyan boxes, bottom right).  $H_\Delta$  is a division-intractable hash function (§4.2),  $H_p$  is a hash to a prime (§4.1), and  $G$  is an RSA quotient group (§2).

To instantiate this interactive protocol in constraints, we apply the Fiat-Shamir heuristic [55], i.e.,  $\mathcal{C}$  computes the challenge  $\ell$  by hashing all of the inputs to the protocol.<sup>6</sup> Figure 1 illustrates the insertion proof’s verification procedure. MultiSwap requires two proofs (one for insertion and one for removal); for this purpose, we hash all inputs to both proofs to derive a common challenge, as is standard [50].

In the rest of this section we explain how to efficiently implement the blocks of Figure 1 in constraints. In particular, we explain how to implement  $H_p$ , the prime hash function used to compute  $\ell$  (§4.1) and  $H_\Delta$ , the division-intractable hash function used to hash each element (§4.2). We also describe optimizations for multiprecision operations (§4.3). Finally, we discuss  $\mathcal{P}$ ’s cost for generating the witness input  $Z$  to  $\mathcal{C}$  (§2.2), notably, the digests  $S \boxplus \{y_i\}$  and  $S \boxminus \{x_i\}$  and the corresponding witnesses  $Q$  for insertion and removal (§4.4).

### 4.1 Hashing to primes

The hash function  $H_p$  (Fig. 1) generates the challenge  $\ell$  used in the Wesolowski proofs of batch insertion and removal. These proofs are sound when  $\mathcal{P}$  has negligible probability of guessing all factors of  $\ell$  before evaluating  $H_p$  [115], which is true for  $\ell$  chosen at random from the first  $2^\lambda$  primes (§2).

<sup>6</sup>This requires that we model the concrete hash function that outputs  $\ell$  as a random oracle [8]; similar assumptions are common in practice.

In our context, a more efficient approach is for  $H_p$  to output slightly larger primes, which are guaranteed by construction to have  $\lambda$  bits of entropy.<sup>7</sup> Soundness is identical.

In standard settings (i.e., outside of constraints), a typical approach (§8) for hashing to a random prime is rejection sampling. Here, the input is fed to a collision-resistant hash whose output seeds a pseudorandom generator (PRG), then the PRG’s outputs are tested in sequence until a prime is found. Verifying correct execution requires, at the very least, testing primality of the purported output. This is typically done with a probabilistic primality test like Miller-Rabin [94]. Such tests, however, generally require many iterations for soundness, where each iteration involves an exponentiation modulo the prime being tested. This would be far too costly if implemented directly in constraints.

Instead, we take advantage of advice from  $\mathcal{P}$  (§2.2). At a high level,  $\mathcal{P}$  helps to recursively construct a Pocklington certificate (§2) for  $H_p$ ’s output, where each intermediate prime  $p_i$  is the result of hashing  $H_p$ ’s input. (This is related to prior approaches; see §8.) This strategy is economical when implemented in constraints, because it uses very little pseudorandomness and requires only *one* exponentiation modulo the resulting prime, plus a few smaller exponentiations.

We now describe the recursive step used to construct  $p_i$  from  $p_{i-1}$ . Further below, we describe the base case and give implementation details. Recall (§2) that a Pocklington witness for  $p_i$  comprises  $(p_{i-1}, r_i, a_i)$  such that  $p_i = p_{i-1} \cdot r_i + 1$ . (If  $p_i$  is prime, some  $a_i$  must exist.) Notice that, given  $p_{i-1}$ , one can find  $p_i$  by testing candidate  $r_i$  values until  $p_{i-1} \cdot r_i + 1$  is prime. To implement this in constraints, we let  $r_i = 2^{b_{n_i}} \cdot h_i + n_i$ , where  $n_i$  is a  $b_{n_i}$ -bit number provided by  $\mathcal{P}$  as advice and  $h_i$  is a  $b_{h_i}$ -bit pseudorandom number (we discuss its generation below).  $\mathcal{P}$  furnishes a corresponding  $a_i$  and  $\mathcal{C}$  includes constraints that compute  $p_i$  and  $r_i$ , and check the witness.

The base case is  $p_0 = 2^{b_{n_0}} \cdot h_0 + n_0$ , for  $h_0$  a pseudorandom number and  $n_0$  supplied by  $\mathcal{P}$ . We fix  $b_{n_0} + b_{h_0} = 32$ , i.e.,  $p_0 < 2^{32}$ , and the constraints test primality of  $p_0$  using a deterministic 3-round Miller-Rabin test that works for all values up to  $2^{32}$  [71]. This test requires 3 exponentiations modulo  $p_0$  with exponents less than 32 bits; these are inexpensive.

We choose bit widths  $b_{n_i}$  such that a valid  $n_i$  exists with overwhelming probability, then choose  $b_{h_i}$  subject to the constraint that  $b_{h_i} + b_{n_i} < \log p_{i-1}$ , which ensures that  $r_i < p_{i-1}$  as required (§2). The entropy of each  $p_i$  is  $\sum_{j=0}^i b_{h_j}$ ; three rounds suffice for 128 bits of entropy using the parameters listed in Figure 2.  $\mathcal{C}$  generates  $h_i$  by hashing the input to  $H_p$  with a hash function  $H$  modeled as a random oracle.

Each iteration yields a prime approximately twice as wide as the prior iteration’s; meanwhile, the cost of each iteration is dominated by an exponentiation. This means that our approach has cost less than that of two exponentiations modulo the final prime. In contrast, using Miller-Rabin to check a

<sup>7</sup>In this section, we use entropy to mean (the negative logarithm of)  $\mathcal{P}$ ’s probability of guessing the correct value, i.e., the guessing entropy.

Iteration, $i$	0	1	2	3
maximum $p_i$ bitwidth	32	63	125	177
$b_{h_i}$	23	20	50	39
$b_{n_i}$	9	11	12	13

Figure 2: Bitwidths for recursive primality proofs in our system. While the  $b_{h_i}$  sum to 132, there are only 128 bits of entropy because each  $h_i$  has its high bit fixed to 1 (§4.1).

135-bit prime (which have roughly 128 bits of entropy) would require 40 exponentiations modulo that prime to give  $\approx 2^{-80}$  probability of outputting a composite (because Miller-Rabin is a probabilistic primality test). Our approach thus saves more than an order of magnitude and provably outputs a prime.

One final optimization is to force the most significant bit of each  $h_i$  to 1; this establishes a lower bound on each  $p_i$  and on  $\ell$  (which is the final  $p_i$ ). As we discuss in Section 4.3, having this lower bound reduces the cost of modular reductions. The tradeoff is a small loss in entropy, namely, 1 bit per iteration. Even so, three rounds suffice to produce a 177-bit prime<sup>8</sup> with 128 bits of entropy.

## 4.2 Division-intractable hashing

Coron and Naccache show [48] that a hash function  $H$  that outputs sufficiently large integers is division intractable when modeled as a random oracle. Informally, this is because in any randomly-selected set of large (say, 2000 bit) numbers, each element has a *distinct*, moderately sized (say, 200 bit) prime factor with high probability.

Security of this hash function rests on the fact that the density of integers in the interval  $[0, \alpha]$  with factors all less than  $\mu$  approaches  $\beta^{-\beta+o(1)}$  as  $\alpha \rightarrow \infty$ , where  $\beta = \frac{\log \alpha}{\log \mu}$ . We conjecture that this density also holds for a large interval around  $\alpha$ , namely,  $[\alpha, \alpha + \alpha^{1/8}]$ . (This is closely related to a conjecture on which the elliptic curve factorization method relies; there, the interval is  $[\alpha - \sqrt{\alpha}, \alpha + \sqrt{\alpha}]$  [69].)

Our hash function is defined as follows: let  $\Delta$  be a public 2048-bit integer chosen at random, and let  $H$  be a hash function with codomain  $[0, 2^{256})$  with 128-bit collision resistance. Then  $H_\Delta(x) = H(x) + \Delta$ . Security of this construction follows from the analysis of [48] in the random oracle model, assuming the conjecture stated above. Concretely, we conjecture that an adversary making  $q$  queries to  $H_\Delta$  has probability roughly  $q \cdot 2^{-128}$  of breaking division intractability.

$H_\Delta$ ’s advantage over prior work is that its implementation in constraints is much smaller. The system parameter  $\Delta$  is baked into the constraints, and the only dynamic values to compute are the base hash  $H(x)$  and the sum  $H(x) + \Delta$ ; using known techniques [77], this sum is inexpensive. Moreover, since all hashes must be reduced modulo the challenge  $\ell$  (Eq. (7))

<sup>8</sup>Even though the prime  $\ell$  comprises only 177 bits,  $\mathcal{C}$  represents it with 192 (Fig. 3), which is the next multiple of the limb width  $b_l$  (32 bits; §2.2).

and  $H_\Delta(x) \bmod \ell = (H(x) + (\Delta \bmod \ell)) \bmod \ell$ , the (costly) reduction  $\rho \bmod \ell$  can be checked once in the constraints and the result can be re-used for each  $H_\Delta(x)$ . We note that while this approach gives smaller  $C$  than hashing to primes (because  $H_\Delta$  and modular reductions are cheaper), it increases  $\mathcal{P}$ 's work (because  $H_\Delta$ 's bit length is longer; §4.4).

### 4.3 Multiprecision arithmetic optimizations

We describe two optimizations for multiprecision arithmetic in constraints, building on ideas described in Section 2.2.

**Computing greatest common divisor.** We observe that addition and multiplication checks can be leveraged to verify a statement  $\gcd(x, y) = d$  by checking three equations over  $\mathbb{Z}$ :

$$\begin{aligned} \exists a, b & \quad a \cdot x + b \cdot y = d & (8) \\ \exists x' & \quad x' \cdot d = x \\ \exists y' & \quad y' \cdot d = y \end{aligned}$$

In constraints, the existential variables above correspond to advice provided by  $\mathcal{P}$ . Verifying coprimality ( $\gcd(x, y) = 1$ ) reduces to condition (8), i.e., materializing the multiplicative inverse of  $x$  modulo  $y$ . We use this simplification in Section 4.1 to verify a Pocklington witness (§2).

**Optimizing division and modular reduction.** Prior work implements division and modular reduction for a dividend  $x$  and divisor  $d$  by having the prover provide, as advice, the quotient  $q$  and remainder  $r < d$  such that  $x = q \cdot d + r$ ; this equality is then checked with multiprecision arithmetic (§2.2). For correctness,  $C$  must enforce upper bounds on the bit widths of  $q$  and  $r$  via bit splitting (§2.2), which requires as many constraints as the sum of the bit widths of  $q$  and  $r$ .

Since  $r$  can range from 0 to  $d - 1$ , its width is just that of  $d$ . The width of  $q$ , however, is slightly more subtle. Since  $q$ 's value is  $\lfloor x/d \rfloor$ , a conservative choice is to assume  $q$  is as wide as  $x$ . But this choice is imprecise:  $q$  is only as wide as  $\lceil \log_2(\lfloor x_{\max}/d_{\min} \rfloor) \rceil$ , where  $x_{\max}$  denotes  $x$ 's maximum possible value, and  $d_{\min}$  denotes  $d$ 's minimum possible value. (Intuitively, this is because  $q$  is small when  $d$  is large.)

As in prior work [77], our system uses a dataflow analysis to track the maximum value of each number, in order to determine the required representation size. To bound  $q$ 's width more tightly using the above expression, we augment this dataflow analysis to also track *minimum* values.

### 4.4 Optimizing the cost of advice generation

The prior sections have treated  $\mathcal{P}$  as an advice oracle. We now discuss  $\mathcal{P}$ 's cost in computing this advice. Prior work [111, 116] shows that  $\mathcal{P}$ 's (single-threaded) cost per constraint is  $\approx 1$  millisecond (this includes, e.g., an elliptic curve point multiplication per constraint [16, 63, 68, 92]). Computing most advice values—including for multiprecision

operations and prime hashing—is negligible by comparison. The exceptions are the witnesses for Wesolowski proofs (§2) used by batch insertion and removal operations (§2.1). (Recall that one of each operation is required for a MultiSwap; §3.)

The witness for a batch insertion  $\llbracket S \uplus \{y_i\} \rrbracket = \llbracket S \rrbracket^{\prod_i H_\Delta(y_i)}$  is the value  $\llbracket S \rrbracket^{\lfloor (\prod_i H_\Delta(y_i))/\ell \rfloor}$ . This exponent has length  $\approx 2048 \cdot k$  bits for  $k$  elements inserted. In microbenchmarks, GMP [64] computes a 2048-bit exponentiation modulo a 2048-bit  $N$  in  $\approx 2.5$  milliseconds (i.e., about  $2.5 \times \mathcal{P}$ 's per-constraint proving cost), so computing this value costs roughly the same as  $2.5 \cdot k$  constraints, which is inconsequential (§5, Fig. 3).

Batch removal is much more expensive: in this case  $\mathcal{P}$  needs to compute a witness for  $\llbracket S \ominus \{x_i\} \rrbracket = \llbracket S \rrbracket^{1/\prod_i H_\Delta(x_i)}$ , i.e.,

$$g^{\lfloor (\prod_{s \in S \ominus \{x_i\}} H_\Delta(s))/\ell \rfloor} \quad (9)$$

No known method is faster than directly computing this expression, because the order of  $\mathbb{G}$  is unknown (recall that this computation is in  $\mathbb{G} = \mathbb{Z}_N^\times / \{\pm 1\}$  where  $N$  has unknown factorization; §2). Meanwhile, this exponent has bit length  $\approx 2048 \cdot M$ , for  $M$  the *total size* of the multiset  $S$  (i.e., it costs roughly the same to compute as  $2.5 \cdot M$  constraints).

Even for large accumulators, this cost may be reasonable: as we show in Section 7, MultiSwap can easily save tens of millions of constraints compared to Merkle trees. On the other hand, proof generation can be parallelized [116], whereas at first glance the exponentiation in (9) appears to be strictly serial [22, 97]. We observe, however, that since  $g$  is fixed, a pre-computation phase can be used to sidestep this issue [33]. Specifically, for some upper bound  $2^m$  on the maximum size of the accumulator, the above exponent is at most  $2^{2048 \cdot 2^m}$ , so pre-computing the values  $g_i = g^{2^{i \cdot 2^m}}$ ,  $0 \leq i < 2048$  (via successive squaring) turns the above exponentiation into a 2048-way multi-exponentiation [88] (which can be computed in parallel): for each  $g_i$ , the exponent is a  $2^m$ -bit chunk of the value  $\lfloor (\prod_{s \in S \ominus \{x_i\}} H_\Delta(s))/\ell \rfloor$ . Further parallelism is possible simply by computing more  $g_i$  with closer spacing.

This precomputation also enables a time-space tradeoff, via windowed multi-exponentiation [88, 105]. In brief, when computing a multi-exponentiation over many bases, first split the bases into groups of size  $t$  and compute for each group a table of size  $2^t$ . This turns  $t$  multiplications into a table lookup and one multiplication, for a factor of  $t$  speedup.  $t = 20$  is reasonable, and reduces the cost of computing the exponentiation in (9) to roughly the equivalent of  $M/8$  constraints.

The above pre-computation is a serial process that requires  $\approx 2048 \cdot 2^m$  squarings in  $\mathbb{G}$ . Assuming that 2048 squarings takes  $\approx 2.5$  milliseconds (i.e., the same amount of time as a general 2048-bit exponentiation, which is pessimistic), this precomputation takes  $\approx 2^m \cdot 2.5$  milliseconds. For  $m = 20$ , this is  $\approx 45$  minutes; for  $m = 25$ , it is  $\approx 1$  day. Note, however, that this pre-computation is entirely untrusted, so it can be done once by anyone and reused indefinitely for the same  $g$ .

Finally, the exponentiation in (9) requires materializing  $\prod_{s \in S \subseteq \{s_i\}} H_{\Delta}(s)$  and then dividing by  $\ell$ . For  $M = 2^{20}$ , the product is  $2^{31}$  bits. Much of its computation can be parallelized, but the final step is a single multiplication of two,  $2^{30}$ -bit values. In microbenchmarks, GMP computes this product in  $\approx 10$  seconds; floor division of a  $2^{31}$ -bit value by a 192-bit value takes  $\approx 0.3$  seconds. Thus, neither is a bottleneck.

## 5 Applications of MultiSwap

In this section we discuss two applications of MultiSwap and compare constraint costs for these applications when implemented using Merkle swaps and MultiSwaps.

**MultiSwap Costs.** The first two rows of Figure 3 model the costs of Merkle swaps and swaps computed via MultiSwap.

A Merkle swap requires hashing the old and new values and Merkle path verifications for each (§2.1), so the number of hash invocations is logarithmic in the number of leaves.

For a MultiSwap, each swap requires a  $H_{\Delta}$  invocation (§4.2), which comprises an invocation of the underlying hash  $H$  and multiprecision arithmetic to compute the result and multiply it mod  $\ell$  (§4, Fig. 1). In addition, each swap is an input to  $H_p$ , which requires another hash invocation. All of these costs are independent of the number of elements in the accumulator. MultiSwap also costs a large constant overhead, however; this is to generate  $\ell$  (§4.1) and check two Wesolowski proofs via modular exponentiations (§2, §4).

### 5.1 Verifiable outsourcing for smart contracts

Blockchain systems [26] like Ethereum [53] enable *smart contracts*: computations defined by a blockchain’s users and executed as part of the block validation procedure. One application of smart contracts is implementing a form of verifiable state update (§1): for global state  $\Gamma$  (stored on the blockchain) and a transaction  $\gamma$  (submitted by a user), the computation (1) checks that  $\gamma$  is valid according to some predicate, and if so (2) updates the global state to a new value  $\Gamma'$ .

Consider, for example, a distributed payment system where  $\Gamma$  comprises a list of users and their public keys and balances. Transactions let users send payments to one another. When Alice wishes to send a payment, she constructs a transaction  $\gamma$  that includes (1) the target user; (2) the amount to send; and (3) a digital signature over the prior two items; she submits this to the smart contract, which verifies it and updates  $\Gamma$ .

A major practical limitation of this approach is that computation, storage, and network traffic are extremely expensive for smart contracts.<sup>9</sup> One solution to this issue, *Rollup* [7], is an instance of verifiable computation (§2.2): the smart contract delegates the work of checking transactions to an

<sup>9</sup>Anecdotally, recent Ethereum prices [54] result in storage costs of more than \$1 per kilobyte. Similarly, per-transaction costs are frequently in the \$0.25 to \$1 range even when executing minimal computation.

untrusted *aggregator*, and then checks a proof that this work was done correctly. To effect this, users submit transactions  $\gamma_i$  to the aggregator rather than directly to the smart contract. The aggregator assembles these transactions into a batch  $\{\gamma_i\}$ , then generates a proof  $\pi$  certifying the correct execution of a computation  $\Psi$  that verifies the batch and updates the global state from  $\Gamma$  to  $\Gamma'$ . Finally, the aggregator submits  $\pi$  and  $\Gamma'$  to the smart contract, which verifies the proof and stores the updated state. Checking this proof is substantially cheaper for the smart contract than verifying each transaction individually, and the exorbitant cost of smart contract execution justifies the aggregator’s cost in generating the proof [110].

In more detail, the constraints  $C$  corresponding to  $\Psi$  (§2.2) take the current state  $\Gamma$  as the input  $X$  and the updated state  $\Gamma'$  as the output  $Y$ .  $\mathcal{P}$  (i.e., the aggregator) supplies the batch  $\{\gamma_i\}$  as part of the witness (i.e., the advice vector  $Z$ ), meaning that the smart contract can verify the proof without reading  $\{\gamma_i\}$ . This saves both computation and network traffic.

Notably, though, even reading  $\Gamma$  and  $\Gamma'$  is too expensive for the smart contract, as is storing  $\Gamma$  on the blockchain. (Recall that verifying a proof requires work proportional to the size of the inputs and outputs; §2.2.) The original Rollup design [7] addresses this by storing  $\Gamma$  in a Merkle tree (§2.1). The inputs and outputs of  $C$  are just Merkle roots, and only this root is stored on the blockchain. Each leaf of this tree contains a tuple  $(pk, bal, \#tx)$  comprising a user’s public key, their balance, and a transaction count (which prevents replaying past transactions). The constraints that verify a transaction in  $C$  thus require two Merkle tree updates (which are also inclusion proofs), one each for the payer and payee.

We observe that a single MultiSwap (§3) can replace all of the Merkle tree updates for a batch of transactions. In particular, MultiSwap’s semantics guarantee sequential consistency of the transactions with respect to  $\Gamma$  and  $\Gamma'$ . And whereas the per-swap cost of Merkle swaps increase logarithmically with the number of accounts stored in  $\Gamma$ , the per-swap cost of MultiSwap is essentially independent of the number of users. This means that for large batches of transactions and/or large numbers of users, a MultiSwap-based Rollup requires far fewer constraints than a Merkle-based one.

**Costs.** The middle two rows of Figure 3 show costs for Rollup using Merkle and MultiSwap. Both cases pay to verify the payer’s signature and ensure that the payer’s balance is sufficient. The difference is in the swap costs, which are discussed above (§5); Rollup requires two swaps per transaction, one each to update the payer’s and payee’s accounts.

### 5.2 Efficient persistent RAM

Recall from Section 2.2 that Pantry-style RAM, while expensive, offers unique functionality: the ability to pass the full state of RAM from one proof to another. This enables computations over persistent state [32], recursively verifiable state machine execution [15, 86], and other useful applications.

		Number of constraints	
System		Per-Operation Costs	Per-Proof Costs
Merkle swap		$2(c_{H_e} + m \cdot c_H)$	
MultiSwap (§3, §4)		$2(c_{H_e} + c_{H_{in}} + c_{split} + c_{+\ell}(f) + c_{\times\ell})$	$4c_{e_G}( \ell ) + 2c_{\times_G} + c_{H_p} + c_{mod_\ell}(b_{H_\Delta})$
Payments (Merkle swap)		Merkle swap $\times 2 + c_{sig} + c_{tx}$	
Payments (MultiSwap)		MultiSwap $\times 2 + c_{sig} + c_{tx}$	MultiSwap
RAM (Merkle-based [32])		$(1 + w)(c_{H_e} + m \cdot c_H)$	
RAM (MultiSwap)		MultiSwap $+ c_{mem-check}$	MultiSwap

$\lambda$	security parameter (128)	$f$	field width ( $\log_2  \mathbb{F} $ ) (255)
$b_{H_\Delta}$	bits in division-intractable hash output (2048)	$b_G$	group element bits ( $\log_2  G $ ) (2048)
$c_{H_e}$	cost of multiset item hash to $\mathbb{F}$ (varies)	$c_H$	cost of $\mathbb{F}^2 \rightarrow \mathbb{F}$ hash (varies)
$c_{H_p}$	cost of prime generation (217703)	$ \ell $	prime challenge bits (192)
$c_{split}$	cost of strict bitsplit in $\mathbb{F}$ (388)	$c_{\times_G}$	operation cost in $G$ (7563)
$c_{sig}$	cost of signature check (12000)	$w$	write fraction (RAM) (varies)
$c_{tx}$	cost of tx validity check (255)	$c_{\times_\ell}$	cost of multiplication, mod $\ell$ (479)
$m$	$\log_2$ of accumulator capacity (varies)	$c_{H_{in}}$	per-operation cost of full-input hash (varies)
$c_{mem-check}$	cost of memory checks, $21 + \log_2 k + 2 \cdot m$ for $k$ operations [111, Fig. 5; 77, Appx. B.A] ( $< 125$ )		
$c_{+\ell}(b)$	cost of addition with two inputs of maximum width $b$ , mod $\ell$ ( $16 + b$ )		
$c_{mod_\ell}(b)$	cost of reduction mod $\ell$ , with a $b$ -bit input ( $16 + b$ )		
$c_{e_G}(b)$	cost of exponentiation with a $b$ -bit exponent, in $G$ ( $7044b$ )		

Figure 3: Constraint count models for Merkle swaps (§2.1), MultiSwap (§3, §4), Payments (§5.1), and Persistent RAM (§5.2). The approximate value of each parameter in our implementation (§6, §7) is given in parentheses. See Section 5 for discussion.

Unfortunately, the high cost (in constraints) of hash functions (§6) limits the number of Pantry-style RAM operations that can be used in a computation—especially for large RAMs [32, 77, 111]. In this section, we show how to use the batched RSA accumulator construction of Section 4 to address this issue. Our design yields a persistent RAM abstraction whose per-access constraint cost is lower than Pantry’s even at modest RAM sizes, and is nearly insensitive to RAM size.

To begin, notice that Pantry’s RAM abstraction essentially stores memory values in a fixed-size Merkle tree, executing a membership proof for each LOAD and a swap for each STORE. Moreover, since our goal is efficiency, our design will ideally check all memory operations using a small number of batched accumulator operations (§4).

This seems to suggest the following (incorrect) approach. First, replace the Merkle tree with an RSA accumulator, representing memory locations as  $\langle \text{addr}, \text{data} \rangle$  tuples. Then, verify all LOAD and STORE operations in a batch using MultiSwap (§3) as follows. For each LOAD from address  $\delta$ ,  $\mathcal{P}$  supplies as advice the value  $v$  purportedly stored at  $\delta$ , and the constraints encode a swap that replaces the tuple  $\langle \delta, v \rangle$  with itself. For each STORE of the value  $v'$  to address  $\delta$ ,  $\mathcal{P}$  supplies as advice the value  $v$  purportedly being overwritten, and the constraints encode the swap  $(\langle \delta, v \rangle, \langle \delta, v' \rangle)$ .

The reason this approach is incorrect is that it does not enforce the consistency of LOAD operations with program execution. In particular, recall (§3) that  $\text{MultiSwap}(S, \sigma, S')$  only

guarantees that  $S'$  is produced by a sequentially-consistent *cycle-free* subsequence  $\sigma' \subseteq \sigma$ . Since LOAD operations are self-cycles, they are not included in  $\sigma'$ . This use of MultiSwap thus *only* guarantees that  $\sigma$  correctly encodes STORE operations—LOADs can return *any* value.

We might attempt to fix this issue by checking LOAD operations using membership proofs. But this is inefficient: checking such a proof requires the constraints to materialize an accumulator that contains the value being loaded; meanwhile, the LOAD might correspond to a prior STORE, in which case the accumulator against which the proof must be checked would first have to be computed. In other words, this strategy makes batching accumulator operations impossible.

Our key insight is that a *hybrid* of the Pantry and BCGT approaches solves this issue. At a high level, our design enforces the correctness of LOAD and STORE operations using an address-ordered transcript (§2.2) while ensuring that this transcript is consistent with the initial and final state of RAM using batched accumulator operations. As above, each memory location is stored in the accumulator as an  $\langle \text{addr}, \text{data} \rangle$  tuple. As in BCGT-style RAM, the constraints build an execution-ordered transcript,  $\mathcal{T}$ , and the constraints ensure that  $\mathcal{T}$  is correctly ordered, coherent, and a permutation of the execution-ordered transcript.

For the initial state of RAM, the constraints enforce consistency by ensuring that the first time an address  $\delta$  is accessed in  $\mathcal{T}$ , the tuple  $\langle \delta, v \rangle$  is removed from the accumulator. If the

first access is a LOAD,  $v$  is the corresponding DATA value from  $\mathcal{T}$ . Otherwise,  $\mathcal{P}$  supplies as advice a claimed  $v$  value such that  $\langle \delta, v \rangle$  is in the accumulator. (For now, we assume that memory location  $\delta$  has *some* corresponding tuple in the accumulator; we discuss uninitialized memory below.) Observe that this ensures consistency, because a removal is only possible if  $\langle \delta, v \rangle$  is indeed in the accumulator.

For the final state of RAM, the constraints enforce consistency by ensuring that the last time an address  $\delta$  is accessed in  $\mathcal{T}$ , the tuple  $\langle \delta, v' \rangle$  is inserted into the accumulator. The value  $v'$  is the corresponding DATA value from  $\mathcal{T}$ . Together with the above, this ensures that all of the accesses to address  $\delta$  collectively result in the swap  $(\langle \delta, v \rangle, \langle \delta, v' \rangle)$ .

Constraints for the above checks work as follows. First, for entry  $i$  in  $\mathcal{T}$ , the constraints compute  $h_{i,\text{del}} = H_{\Delta}(\langle \text{ADDR}_i, v \rangle)$  and  $h_{i,\text{ins}} = H_{\Delta}(\langle \text{ADDR}_i, v' \rangle)$  (§4.2). Then, for each sequential pair of entries  $i, i+1$  in  $\mathcal{T}$ , if  $\text{ADDR}_i \neq \text{ADDR}_{i+1}$ , then entry  $i$  must be the last access to  $\text{ADDR}_i$  and entry  $i+1$  must be the first access to  $\text{ADDR}_{i+1}$ . Finally, the constraints compute  $\prod_{i \in \mathcal{F}} h_{i,\text{del}} \bmod \ell$  and  $\prod_{i \in \mathcal{L}} h_{i,\text{ins}} \bmod \ell$  (§4), the values inserted into and removed from the accumulator, respectively, for  $\mathcal{F}$  the first-accessor set and  $\mathcal{L}$  the last-accessor set.

**Handling uninitialized memory.** A remaining issue is how to handle the case where memory is uninitialized. Recall that in the BCGT approach, a LOAD not preceded by a STORE to the same address is serviced with a default value, say, 0. That does not work here, because this approach relies crucially on swapping old values for new ones, to ensure consistency with both the initial and final accumulators.

A straightforward solution is to ensure that every memory location is initialized, by executing a setup phase that constructs an accumulator containing the tuple  $\langle \delta, 0 \rangle$  for every address  $\delta$ . The cost of constructing this accumulator is high when the address space is large, since it amounts to one exponentiation per entry in RAM. Note, however, that this computation can be parallelized using the pre-computed values described in Section 4.4, and admits the same time-space tradeoff described in that section.<sup>10</sup>

**Costs.** The constraint costs of memory accesses are shown in the bottom two rows of Figure 3. The Merkle-based RAM requires two proofs of membership for each STORE, but only one for each LOAD [32], so it is slightly cheaper than a Merkle swap—but logarithmic in RAM size.

The RSA accumulator-based RAM uses one MultiSwap for all LOADs and STOREs, with attendant per-operation

<sup>10</sup>An alternative solution is to implement, in essence, a shadow memory [89] indicating which addresses are valid. This is effected by storing a canary value  $\text{valid}[\delta]$  in the accumulator for each address  $\delta$  for which some tuple  $\langle \delta, \cdot \rangle$  exists. If  $\Psi$  attempts to LOAD or STORE from a memory location  $\delta$  for which no value exists,  $\mathcal{P}$  supplies a proof of non-membership (§2.1) for  $\text{valid}[\delta]$ , plus a default value. This obviates the setup phase, but requires additional constraints to (1) compute  $H_{\Delta}(\text{valid}[\text{ADDR}_i])$  for each entry in  $\mathcal{T}$ , (2) check a batched non-membership proof, (3) check a batched insertion of  $\text{valid}[\cdot]$  values (which can be combined with the swap check), and (4) enforce correctness of the default value. Further exploration is future work.

costs (which are independent of RAM size; §5). It also incurs extra per-operation costs to check  $\mathcal{T}$  as described above; these are logarithmic in the number of accesses but concretely very inexpensive (§2.2, [111, Fig. 5; 77, Appx. B-A]).

## 6 Implementation

We implement a library comprising multiprecision arithmetic, Pocklington prime certification, RSA accumulators, and Merkle trees. This library builds on Bellman [9], a Rust library for constructing (compilers to) constraints.

We also implement or adapt four hash functions: MiMC [1], which costs 731 constraints (91 rounds of the  $x^7$  permutation); Poseidon [67], which costs 542 constraints; Pedersen [70, 93], which costs 2753 constraints (based on the JubJub elliptic curve [28]), and SHA-256 [57], which costs 45567 constraints. We adapt the latter three hashes from Sapling [99].

Based on our library, we implement two end-to-end constraint systems: one encoding just a sequence of swaps of variable length, and one that verifies transactions for a distributed payment system (§5.1).

Finally, we implement custom Bellman constraint synthesizers (`ConstraintSystems`, in the jargon of Bellman) that allow us to quickly measure a constraint system’s size and  $\mathcal{P}$ ’s cost computing a corresponding witness.

We use a 2048-bit RSA quotient group (§2) modulo the RSA-2048 challenge number [74, 98], and choose a random 2048-bit  $\Delta$  to define the division-intractable hash function  $H_{\Delta}$  (§4.2); we give concrete values in Appendix B. We synthesize all constraints over the BLS12-381 [27] curve.

In total, our implementation comprises  $\approx 10,000$  lines of Rust. We have released it under an open-source license [10].

## 7 Evaluation

We evaluate our MultiSwap implementation, comparing it to Merkle trees by answering the following questions:

- (1) How does the cost of a MultiSwap compare to the cost of Merkle swaps for a batch of swaps? In particular, what is the *break-even point* (i.e., the number of operations beyond which MultiSwap is cheaper), and how do costs compare for a fixed (large) constraint budget?
- (2) What is the effect of hash function cost on the tradeoff between RSA accumulators and Merkle trees?

We answer the first question by synthesizing constraint systems for both MultiSwap and Merkle swaps, at varying set and batch sizes (§7.1). We also synthesize constraints for the Rollup application (§7.2), and compare the persistent RAM application using a cost model (§7.3). For the second question, we evaluate the break-even point for MultiSwap versus the cost of the underlying hash function, for four different hash functions (§7.4).

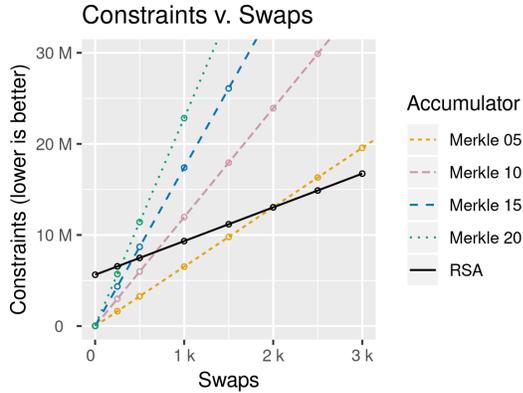


Figure 4: Constraint count v. number of swaps (§7.1). “Merkle  $m$ ” denotes a Merkle tree with  $2^m$  leaves.

Accumulator	Swaps	Accumulator	Transactions
Merkle 05	153,045	Merkle 05	38,283
Merkle 10	83,479	Merkle 10	27,033
Merkle 15	57,392	Merkle 15	20,893
Merkle 20	43,727	Merkle 20	17,026
RSA	268,604	RSA	48,560

(a) Swaps (§7.1).                      (b) Payments (§7.2).

Figure 5: Number of operations verifiable in  $10^9$  constraints (higher is better).

In sum, we find that MultiSwap breaks even for batch sizes of at most a few thousand operations; for large sets, this value is a few hundred. We also find that MultiSwap’s advantage is greater when hashing is more expensive.

**Baseline.** Our baselines are constraint systems (§2.2) that use Merkle trees (§2.1) to store state. For each baseline, we fix capacity to be  $M = 2^m$ , for a range of  $m$  values. In all experiments except persistent RAM, the basic Merkle tree operation is a swap (§5, Fig. 3). Merkle-based RAMs use a mix of membership proofs and swaps (§2.1, §2.2); we discuss further in Section 7.3.

**Setup.** Except in Section 7.4, both Merkle trees and MultiSwap fix the hash function  $H$  (§4.1, §4.2) to be our Poseidon [67] implementation (§6). As we show in Section 7.4, this choice is favorable to the Merkle tree baseline, because Poseidon is inexpensive when implemented in constraints.

**Method.** Our experiments compare costs in terms of number of constraints. For this purpose, we use the custom Bellman synthesizer described in Section 6. Number of constraints is a good cost metric because  $\mathcal{P}$ ’s costs (both in proving time and in the amount of RAM required) are dominated by this value in the back-ends that we target (§2.2).<sup>11</sup>  $\mathcal{V}$ ’s costs are in-

<sup>11</sup>In fact, this metric is a conservative predictor of performance, because  $\mathcal{P}$ ’s computation costs are strictly super-linear with the number of constraints.

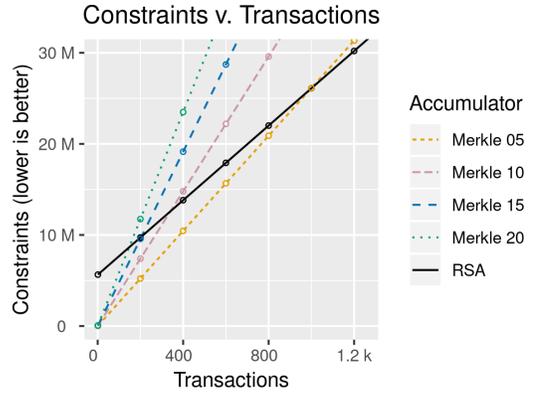


Figure 6: Constraint count v. number of transactions (§7.2). “Merkle  $m$ ” denotes a Merkle tree with  $2^m$  leaves.

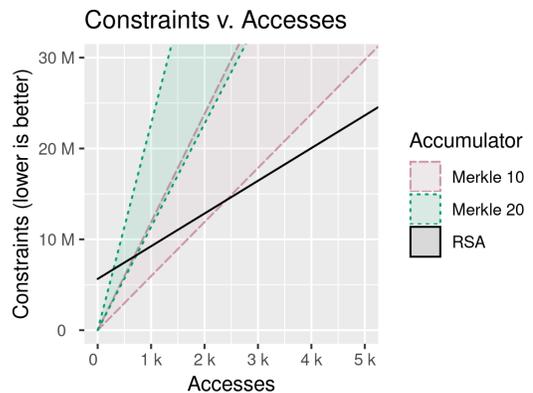


Figure 7: Constraint count v. number of accesses (§7.3). “Merkle  $m$ ” denotes a Merkle tree with  $2^m$  leaves. Ribbons indicate variation according to write load, from 0 to 100%.

dependent of the accumulator type and number of constraints, and are concretely small; we do not evaluate these costs.

One caveat is that, as discussed in Section 4.4, MultiSwap incurs extra proving time for  $\mathcal{P}$  compared to Merkle trees, to compute advice values. For the accumulator sizes we consider (up to  $2^{20}$ ), this additional cost is insignificant ( $\approx 1\%$  of  $\mathcal{P}$ ’s total proving work), and thus does not significantly affect break-even point. We discuss further in Section 9.

## 7.1 MultiSwap versus Merkle swaps

**Benchmark.** This experiment compares the costs of MultiSwap and Merkle trees for a computation comprising only swaps, varying the number of swaps and set size.

**Results.** Figure 4 shows the results. The cost of Merkle trees varies with set size, because the number of hash invocations depends on this value (§2.1; §5, Fig. 3). In contrast, the constraint cost of MultiSwap is independent of the number

of elements contained in the set, and the per-element cost is smaller even for very small sets.

On the other hand, MultiSwap pays a large overhead ( $\approx 5.5$  million constraints) to evaluate  $H_p$  and verify two Wesolowski proofs (§4; §5, Fig. 3). Thus, MultiSwap requires some minimum batch size to break even. For small sets the required batch size is relatively large ( $\approx 2000$ ), whereas for large sets the required batch size is just a few hundred.

Figure 5a shows the number of swaps that fit in  $10^9$  constraints, for different accumulators. (We compare at this size because it is the largest that prior work can handle [116].) Depending on set size, MultiSwap improves reachable batch sizes by  $\approx 1.8\times$  to  $\approx 6.1\times$ .

## 7.2 Application: payment system

**Benchmark.** This experiment compares the costs of MultiSwap and Merkle trees for the Rollup application described in Section 5.1. We measure cost versus the number of *transactions* (a signature verification, a validity check, and two swaps). Signatures use the scheme from ZCash [70].

**Results.** Figure 6 shows the results. In contrast with the previous experiment, here all accumulator types pay a fixed overhead per transaction (this is dominated by signature verification), which reduces MultiSwap’s per-transaction advantage. Even so, the break-even point for the smallest set size is  $\approx 1000$  transactions; it is fewer than 200 for the largest size. (In this application, set size corresponds to number of accounts;  $2^{20}$  is quite reasonable.)

Figure 5b shows the number of transactions that fit in  $10^9$  constraints, for different accumulators. MultiSwap’s advantage ranges from  $\approx 1.3\times$  to  $\approx 2.9\times$ , depending on set size.

## 7.3 Application: persistent RAM

**Benchmark.** This experiment compares the costs of MultiSwap-based and Pantry’s [32] Merkle-based persistent RAM 5.2. Unlike all other experiments, rather than synthesizing we compare based on the cost model of Figure 3 (§5), which is derived from prior work [77, 111]; future work is to port Buffet’s RAM compiler to Bellman and synthesize. We report cost versus RAM size.

**Results.** Figure 7 shows the results. For Merkle-based RAM, bands in the figure represent varying write loads, from 0 (lowest cost) to 100% (highest cost). As in prior experiments, MultiSwap’s cheaper per-operation cost yields a break-even point of at most a few thousand operations. This model includes the cost of memory consistency checks (§2.2, §5.2, Fig. 3); they cost less than 100 constraints per operation and are ultimately insignificant.

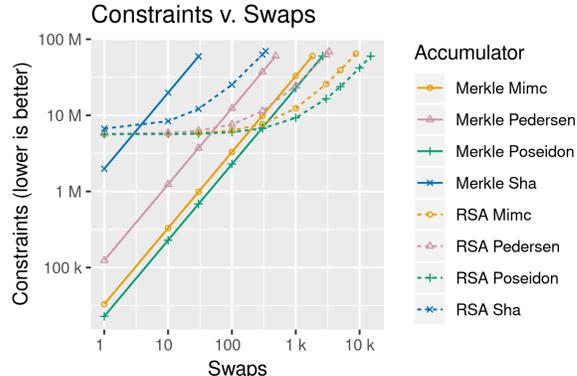


Figure 8: Constraint count v. number of swaps, varying hash function (§7.4). Merkle trees are all of depth 20.

## 7.4 Effect of hash cost

**Benchmark.** This experiment measures the effect of hash cost on MultiSwap’s break-even point for a computation comprising a batch of swaps (§7.1). We fix the Merkle tree’s depth at 20 (others are analogous) and synthesize using MiMC, Poseidon, Pedersen/Jubjub, and SHA-256 (§6).

**Results.** Figure 8 shows the results (note that this is a log-log plot to capture the wide range of sizes). As expected, in all cases Merkle trees are cheaper for small numbers of operations. For the least expensive hash (Poseidon), MultiSwap’s break-even point is the highest; as hash cost increases, so does MultiSwap’s advantage. (We report results in all other experiments with Poseidon, which is the worst case for MultiSwap.)

## 8 Related work

**Verifiable computation.** The literature on verifiable computation is both broad and deep; a somewhat recent survey [114] gives a thorough treatment of the area’s beginnings.

Our work builds most directly on xJsnark’s [77] multiprecision arithmetic and on the RAM primitives first described by Ben-Sasson et al. [12] and further refined by Ben-Sasson et al. [13, 16], in Buffet [111], and in xJsnark. Buffet and xJsnark both extend lines of work concerned with efficiently compiling high-level programs to constraints, including Pepper [102], Ginger [103], Pinocchio [92], and Pantry [32].

Several other works in this area deal with persistent state. Pantry [32] was the first to use Merkle trees for stateful computations, and its persistent RAM primitive inspired ours (§5.2). vSQL [118] builds a verifiable subset of SQL, building on the interactive proofs of Goldwasser et al. [65], Cormode et al. [47], and Thaler [106], and on the polynomial commitments of Papamanthou et al. [91], which build on the work of Kate et al. [75]. In contrast to the persistent RAM and multiset abstractions we develop, vSQL exposes a database abstraction; queries operate on all rows in parallel.

ADSNARK [4] extends the Pinocchio [92] SNARK to support operations on authenticated data provided by a third party. Geppetto [49] also extends Pinocchio, allowing the verifier to commit to inputs for a specific computation and later verify a proof against that commitment, and also enabling data transfer between separate constraint systems bundled into one proof. Fiore et al. [56] take Geppetto’s commitments to inputs a step further, making them computation independent. In contrast to a multiset or persistent RAM abstraction, however, all of these systems require a number of constraints sufficient to read *every* input value—in other words, a multiset of size  $M$  implies at least  $M$  constraints. Further, they do not efficiently support programs whose multiset or RAM accesses depend on inputs and thus cannot be statically analyzed (§2.2).

Spice [100] aims to enable zero-knowledge auditing of concurrent services. Spice differs from this work in two key ways. First, Spice’s core state verification primitive requires a number of constraints linear in the *total size* of the state; this cost is amortized over a batch of state updates. In contrast, MultiSwap operations (§3) have constraint costs that depend only on the number of state *updates*, not on state size. Second, verification costs in Spice scale both with the number of state updates in a batch and with total state size. In our work, verification cost is independent of both the number of accesses to state and the size of state, consistent with our goal of saving work for the verifier (e.g., a smart contract; §5.1).

**Accumulators.** Cryptographic accumulators [17] based on RSA have a long history [5, 40, 78, 81]. The recent work of Boneh et al. [24] builds upon work by Wesolowski [115] to construct *batched* membership and non-membership proofs for these accumulators. Our work builds directly on this line.

Merkle-based accumulators have also seen extensive study [38, 87], and related structures have seen applications, e.g., in the blockchain [95] and PKI contexts [96]. These works all rely crucially on collision-resistant hashing, which is expensive when expressed as constraints (§6, §7).

Two other lines of work build accumulators [39, 42, 51, 90] and vector commitments [41, 79, 80] from bilinear maps. Elliptic curve operations and pairings appear to be very expensive when compiled to constraints [15], but these lines may nevertheless be an interesting direction for further study.

**Prime generation.** A long line of work [30, 31, 66, 72, 73] aims to efficiently generate pseudorandom prime numbers. In some cases, uniformly distributed primes [58] are desirable. All of these proceed in “guess-and-check” fashion, which is inefficient when implemented in constraints (see §4.1). Most closely, Maurer [84, 85] and Shawe-Taylor [104] describe prime generation methods based on Pocklington certificates; Clavier et al. [46] optimize for embedded devices. To our knowledge, no prior work tackles this problem in our context.

## 9 Discussion and conclusion

We have shown that in verifiable state applications with moderate to large state, accessed hundreds to thousands of times, RSA accumulators are less costly than Merkle trees.

There are two caveats: first, RSA accumulators require a trusted setup. In practice, most SNARKs [15, 63, 68, 92] also require a trusted setup, so this is not a significant burden. Moreover, it is possible to mitigate trust requirements by generating an RSA modulus using a multiparty computation [25, 59]. A conjectured alternative that avoids trusted setup is a class group of imaginary quadratic order [24, 35]; exploring efficient constraint implementations is future work.

Second, for very large sets (say,  $> 2^{25}$ )  $\mathcal{P}$ ’s cost (in time) for advice generation is high (§4.4). For small batch sizes, this cost may overwhelm the time saved because of reduced constraint count. Recall, however, that reducing the number of constraints also reduces  $\mathcal{P}$ ’s RAM requirements; meanwhile,  $\mathcal{P}$ ’s advice generation task requires little memory. This means that *even if* an RSA accumulator requires greater total proving *time* than a Merkle tree, the RSA accumulator’s use may still be justified because it reduces the amount of RAM  $\mathcal{P}$  needs to generate a proof. Since RAM is a major bottleneck [111, 116] (§1), such a time-space tradeoff has significant practical benefit. Exploring this tradeoff is future work.

## References

- [1] Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. MiMC: Efficient encryption and cryptographic hashing with minimal multiplicative complexity. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part I*, volume 10031 of *LNCS*, pages 191–219. Springer, Heidelberg, December 2016.
- [2] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 2087–2104. ACM Press, October / November 2017.
- [3] László Babai, Lance Fortnow, Leonid A. Levin, and Mario Szegedy. Checking computations in polylogarithmic time. In *23rd ACM STOC*, pages 21–31. ACM Press, May 1991.
- [4] Michael Backes, Manuel Barbosa, Dario Fiore, and Raphael M. Reischuk. ADSNARK: Nearly practical and privacy-preserving proofs on authenticated data. In *2015 IEEE Symposium on Security and Privacy*, pages 271–286. IEEE Computer Society Press, May 2015.

- [5] Niko Bari and Birgit Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 480–494. Springer, Heidelberg, May 1997.
- [6] Paulo S. L. M. Barreto, Ben Lynn, and Michael Scott. Constructing elliptic curves with prescribed embedding degrees. In Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano, editors, *SCN 02*, volume 2576 of *LNCS*, pages 257–267. Springer, Heidelberg, September 2003.
- [7] barryWhiteHat. roll\_up: Scale ethereum with SNARKs. [https://github.com/barryWhiteHat/roll\\_up](https://github.com/barryWhiteHat/roll_up).
- [8] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [9] Bellman circuit library, community edition. <https://github.com/matter-labs/bellman>.
- [10] Bellman-BigNat. <https://github.com/alex-ozdemir/bellman-bignat>.
- [11] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 701–732. Springer, Heidelberg, August 2019.
- [12] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. Fast reductions from RAMs to delegatable succinct constraint satisfaction problems: extended abstract. In Robert D. Kleinberg, editor, *ITCS 2013*, pages 401–414. ACM, January 2013.
- [13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, August 2013.
- [14] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for R1CS. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 103–128. Springer, Heidelberg, May 2019.
- [15] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 276–294. Springer, Heidelberg, August 2014.
- [16] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 781–796. USENIX Association, August 2014.
- [17] Josh Cohen Benaloh and Michael de Mare. One-way accumulators: A decentralized alternative to digital signatures (extended abstract). In Tor Hellesest, editor, *EUROCRYPT'93*, volume 765 of *LNCS*, pages 274–285. Springer, Heidelberg, May 1994.
- [18] Václav Beneš. *Mathematical theory of connecting networks and telephone traffic*. Mathematics in Science and Engineering. Elsevier Science, 1965.
- [19] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In Shafi Goldwasser, editor, *ITCS 2012*, pages 326–349. ACM, January 2012.
- [20] Manuel Blum, William S. Evans, Peter Gemmell, Sampath Kannan, and Moni Naor. Checking the correctness of memories. In *32nd FOCS*, pages 90–99. IEEE Computer Society Press, October 1991.
- [21] Andrew J. Blumberg, Justin Thaler, Victor Vu, and Michael Walfish. Verifiable computation using multiple provers. Cryptology ePrint Archive, Report 2014/846, 2014. <https://eprint.iacr.org/2014/846>.
- [22] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 757–788. Springer, Heidelberg, August 2018.
- [23] Dan Boneh, Benedikt Bünz, and Ben Fisch. A survey of two verifiable delay functions. Cryptology ePrint Archive, Report 2018/712, 2018. <https://eprint.iacr.org/2018/712>.
- [24] Dan Boneh, Benedikt Bünz, and Ben Fisch. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 561–586. Springer, Heidelberg, August 2019.

- [25] Dan Boneh and Matthew K. Franklin. Efficient generation of shared RSA keys (extended abstract). In Burton S. Kaliski Jr., editor, *CRYPTO'97*, volume 1294 of *LNCS*, pages 425–439. Springer, Heidelberg, August 1997.
- [26] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. SoK: Research perspectives and challenges for bitcoin and cryptocurrencies. In *2015 IEEE Symposium on Security and Privacy*, pages 104–121. IEEE Computer Society Press, May 2015.
- [27] Sean Bowe. BLS12-381: New zk-SNARK elliptic curve construction. <https://electriccoin.co/blog/new-snark-curve/>, March 2017.
- [28] Sean Bowe. Cultivating Sapling: Faster zk-SNARKs. <https://electriccoin.co/blog/cultivating-sapling-faster-zksnarks/>, September 2017.
- [29] Sean Bowe, Alessandro Chiesa, Matthew Green, Ian Miers, Pratyush Mishra, and Howard Wu. Zexe: Enabling decentralized private computation. Cryptology ePrint Archive, Report 2018/962, 2018. <https://eprint.iacr.org/2018/962>.
- [30] Jørgen Brandt and Ivan Damgård. On generation of probable primes by incremental search. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 358–370. Springer, Heidelberg, August 1993.
- [31] Jørgen Brandt, Ivan Damgård, and Peter Landrock. Speeding up prime number generation. In Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto, editors, *ASIACRYPT'91*, volume 739 of *LNCS*, pages 440–449. Springer, Heidelberg, November 1993.
- [32] Benjamin Braun, Ariel J. Feldman, Zuocheng Ren, Srinath Setty, Andrew J. Blumberg, and Michael Walfish. Verifying computations with state. In *SOSP*, November 2013.
- [33] Ernest F. Brickell, Daniel M. Gordon, Kevin S. McCurley, and David Bruce Wilson. Fast exponentiation with precomputation (extended abstract). In Rainer A. Rueppel, editor, *EUROCRYPT'92*, volume 658 of *LNCS*, pages 200–207. Springer, Heidelberg, May 1993.
- [34] John Brillhart, D. H. Lehmer, and J. L. Selfridge. New primality criteria and factorizations of  $2^m \pm 1$ . *Math. Comp.*, 29(130):620–647, April 1975.
- [35] Johannes Buchmann and Safuat Hamdy. A survey on iq cryptography. In *In Proceedings of Public Key Cryptography and Computational Number Theory*, pages 1–15, 2001.
- [36] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
- [37] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. Cryptology ePrint Archive, Report 2019/1229, 2019. <https://eprint.iacr.org/2019/1229>.
- [38] Philippe Camacho, Alejandro Hevia, Marcos A. Kiwi, and Roberto Opazo. Strong accumulators from collision-resistant hashing. In Tzong-Chen Wu, Chin-Laung Lei, Vincent Rijmen, and Der-Tsai Lee, editors, *ISC 2008*, volume 5222 of *LNCS*, pages 471–486. Springer, Heidelberg, September 2008.
- [39] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In Stanislaw Jarecki and Gene Tsudik, editors, *PKC 2009*, volume 5443 of *LNCS*, pages 481–500. Springer, Heidelberg, March 2009.
- [40] Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 61–76. Springer, Heidelberg, August 2002.
- [41] Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 55–72. Springer, Heidelberg, February / March 2013.
- [42] Alexander Chepur, Charalampos Papamanthou, and Yupeng Zhang. Edrax: A cryptocurrency with stateless transaction validation. Cryptology ePrint Archive, Report 2018/968, 2018. <https://eprint.iacr.org/2018/968>.
- [43] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. Cryptology ePrint Archive, Report 2019/1047, 2019. <https://eprint.iacr.org/2019/1047>.
- [44] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. Cryptology ePrint Archive, Report 2019/1076, 2019. <https://eprint.iacr.org/2019/1076>.

- [45] Alessandro Chiesa, Eran Tromer, and Madars Virza. Cluster computing in zero knowledge. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 371–403. Springer, Heidelberg, April 2015.
- [46] Christophe Clavier, Benoit Feix, Loïc Thierry, and Pascal Paillier. Generating provable primes efficiently on embedded devices. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 372–389. Springer, Heidelberg, May 2012.
- [47] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In Shafi Goldwasser, editor, *ITCS 2012*, pages 90–112. ACM, January 2012.
- [48] Jean-Sébastien Coron and David Naccache. Security analysis of the Gennaro-Halevi-Rabin signature scheme. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 91–101. Springer, Heidelberg, May 2000.
- [49] Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. Geppetto: Versatile verifiable computation. In *2015 IEEE Symposium on Security and Privacy*, pages 253–270. IEEE Computer Society Press, May 2015.
- [50] Ronald John Fitzgerald Cramer. *Modular design of secure yet practical cryptographic protocols*. PhD thesis, Universiteit van Amsterdam, January 1997.
- [51] Ivan Damgård and Nikos Triandopoulos. Supporting non-membership proofs with bilinear-map accumulators. Cryptology ePrint Archive, Report 2008/538, 2008. <http://eprint.iacr.org/2008/538>.
- [52] Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, and Bryan Parno. Cinderella: Turning shabby X.509 certificates into elegant anonymous credentials with the magic of verifiable computation. In *2016 IEEE Symposium on Security and Privacy*, pages 235–254. IEEE Computer Society Press, May 2016.
- [53] Ethereum. <https://ethereum.org>.
- [54] ETH Gas Station. <https://ethgasstation.info>.
- [55] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.
- [56] Dario Fiore, Cédric Fournet, Esha Ghosh, Markulf Kohlweiss, Olga Ohrimenko, and Bryan Parno. Hash first, argue later: Adaptive verifiable computations on outsourced data. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1304–1316. ACM Press, October 2016.
- [57] Secure hash standard. NIST FIPS PUB 180-4, August 2015.
- [58] Pierre-Alain Fouque and Mehdi Tibouchi. Close to uniform prime number generation with fewer random bits. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *ICALP 2014, Part I*, volume 8572 of *LNCS*, pages 991–1002. Springer, Heidelberg, July 2014.
- [59] Tore Kasper Frederiksen, Yehuda Lindell, Valery Osheter, and Benny Pinkas. Fast distributed RSA key generation for semi-honest and malicious adversaries. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 331–361. Springer, Heidelberg, August 2018.
- [60] Matthew Fredrikson and Benjamin Livshits. Zø: An optimizing distributing zero-knowledge compiler. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 909–924. USENIX Association, August 2014.
- [61] Ariel Gabizon. AuroraLight: Improved prover efficiency and SRS size in a sonic-like system. Cryptology ePrint Archive, Report 2019/601, 2019. <https://eprint.iacr.org/2019/099>.
- [62] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. <https://eprint.iacr.org/2019/953>.
- [63] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
- [64] GNU multiple precision arithmetic library. <https://gmplib.org>.
- [65] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 113–122. ACM Press, May 2008.

- [66] John Gordon. Strong primes are easy to find. In Thomas Beth, Norbert Cot, and Ingemar Ingemarsson, editors, *EUROCRYPT'84*, volume 209 of *LNCS*, pages 216–223. Springer, Heidelberg, April 1985.
- [67] Lorenzo Grassi, Daniel Kales, Dmitry Khovratovich, Arnab Roy, Christian Rechberger, and Markus Schafneggger. Starkad and poseidon: New hash functions for zero knowledge proof systems. Cryptology ePrint Archive, Report 2019/458, 2019. <https://eprint.iacr.org/2019/458>.
- [68] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
- [69] Jr. Hendrik W. Lenstra. Factoring integers with elliptic curves. *Annals of Mathematics*, 126(3):649–673, 1987.
- [70] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification. <https://github.com/zcash/zips/blob/master/protocol/protocol.pdf>.
- [71] Gerhard Jaeschke. On strong pseudoprimes to several bases. *Mathematics of Computation*, 61(204):915–926, 1993.
- [72] Marc Joye and Pascal Paillier. Fast generation of prime numbers on portable devices: An update. In Louis Goubin and Mitsuru Matsui, editors, *CHES 2006*, volume 4249 of *LNCS*, pages 160–173. Springer, Heidelberg, October 2006.
- [73] Marc Joye, Pascal Paillier, and Serge Vaudenay. Efficient generation of prime numbers. In Çetin Kaya Koç and Christof Paar, editors, *CHES 2000*, volume 1965 of *LNCS*, pages 340–354. Springer, Heidelberg, August 2000.
- [74] Burt Kaliski. RSA factoring challenge. In Henk C. A. van Tilborg, editor, *Encyclopedia of Cryptography*. Springer, 2005.
- [75] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, December 2010.
- [76] Ahmed E. Kosba, Dimitrios Papadopoulos, Charalampos Papamanthou, Mahmoud F. Sayed, Elaine Shi, and Nikos Triandopoulos. TRUESET: Faster verifiable set computations. In Kevin Fu and Jaeyeon Jung, editors, *USENIX Security 2014*, pages 765–780. USENIX Association, August 2014.
- [77] Ahmed E. Kosba, Charalampos Papamanthou, and Elaine Shi. xJsnark: A framework for efficient verifiable computation. In *2018 IEEE Symposium on Security and Privacy*, pages 944–961. IEEE Computer Society Press, May 2018.
- [78] Jiangtao Li, Ninghui Li, and Rui Xue. Universal accumulators with efficient nonmembership proofs. In Jonathan Katz and Moti Yung, editors, *ACNS 07*, volume 4521 of *LNCS*, pages 253–269. Springer, Heidelberg, June 2007.
- [79] Benoît Libert, Somindu C. Ramanna, and Moti Yung. Functional commitment schemes: From polynomial commitments to pairing-based accumulators from simple assumptions. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *ICALP 2016*, volume 55 of *LIPICs*, pages 30:1–30:14. Schloss Dagstuhl, July 2016.
- [80] Benoît Libert and Moti Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 499–517. Springer, Heidelberg, February 2010.
- [81] Helger Lipmaa. Secure accumulators from euclidean rings without trusted setup. In Feng Bao, Pierangela Samarati, and Jianying Zhou, editors, *ACNS 12*, volume 7341 of *LNCS*, pages 224–240. Springer, Heidelberg, June 2012.
- [82] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge snarks from linear-size universal and updateable structured reference strings. Cryptology ePrint Archive, Report 2019/099, 2019. <https://eprint.iacr.org/2019/099>.
- [83] Matter network. <https://demo.matter-labs.io/explorer/>.
- [84] Ueli M. Maurer. Fast generation of secure RSA-moduli with almost maximal diversity. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *EUROCRYPT'89*, volume 434 of *LNCS*, pages 636–647. Springer, Heidelberg, April 1990.
- [85] Ueli M. Maurer. Fast generation of prime numbers and secure public-key cryptographic parameters. *Journal of Cryptology*, 8(3):123–155, September 1995.
- [86] Izaak Meckler and Evan Shapiro. Coda: Decentralized cryptocurrency at scale. <https://cdn.codaprotocol.com/v2/static/coda-whitepaper-05-10-2018-0.pdf>, May 2018.

- [87] Ralph C. Merkle. A digital signature based on a conventional encryption function. In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 369–378. Springer, Heidelberg, August 1988.
- [88] Bodo Möller. Algorithms for multi-exponentiation. In Serge Vaudenay and Amr M. Youssef, editors, *SAC 2001*, volume 2259 of *LNCS*, pages 165–180. Springer, Heidelberg, August 2001.
- [89] Nicholas Nethercote and Julian Seward. How to shadow every byte of memory used by a program. In *VEE*, June 2007.
- [90] Lan Nguyen. Accumulators from bilinear pairings and applications. In Alfred Menezes, editor, *CT-RSA 2005*, volume 3376 of *LNCS*, pages 275–292. Springer, Heidelberg, February 2005.
- [91] Charalampos Papamanthou, Elaine Shi, and Roberto Tamassia. Signatures of correct computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 222–242. Springer, Heidelberg, March 2013.
- [92] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.
- [93] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 129–140. Springer, Heidelberg, August 1992.
- [94] Michael O. Rabin. Probabilistic algorithm for testing primality. *J. Number Theory*, 12(1):128–138, February 1980.
- [95] Leonid Reyzin, Dmitry Meshkov, Alexander Chepurnoy, and Sasha Ivanov. Improving authenticated dynamic dictionaries, with applications to cryptocurrencies. In Aggelos Kiayias, editor, *FC 2017*, volume 10322 of *LNCS*, pages 376–392. Springer, Heidelberg, April 2017.
- [96] Leonid Reyzin and Sophia Yakoubov. Efficient asynchronous accumulators for distributed PKI. In Vassilis Zikas and Roberto De Prisco, editors, *SCN 16*, volume 9841 of *LNCS*, pages 292–309. Springer, Heidelberg, August / September 2016.
- [97] Ronald L. Rivest, Adi Shamir, and David A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, MIT LCS, March 1996.
- [98] The RSA challenge numbers. <https://web.archive.org/web/20130921041734/http://www.emc.com/emc-plus/rsa-labs/historical/the-rsa-challenge-numbers.htm>.
- [99] Sapling cryptography library, community edition. <https://github.com/matter-labs/sapling-crypto>.
- [100] Srinath Setty, Sebastian Angel, Trinabh Gupta, and Jonathan Lee. Proving the correct execution of concurrent services in zero-knowledge. In *OSDI*, October 2018.
- [101] Srinath Setty, Benjamin Braun, Victor Vu, Andrew J. Blumberg, Bryan Parno, and Michael Walfish. Resolving the conflict between generality and plausibility in verified computation. In *EuroSys*, April 2013.
- [102] Srinath T. V. Setty, Richard McPherson, Andrew J. Blumberg, and Michael Walfish. Making argument systems for outsourced computation practical (sometimes). In *NDSS 2012*. The Internet Society, February 2012.
- [103] Srinath T. V. Setty, Victor Vu, Nikhil Panpalia, Benjamin Braun, Andrew J. Blumberg, and Michael Walfish. Taking proof-based verified computation a few steps closer to practicality. In Tadayoshi Kohno, editor, *USENIX Security 2012*, pages 253–268. USENIX Association, August 2012.
- [104] John Shawe-Taylor. Generating strong primes. *Electronics Letters*, 22(16):875–877, 1986.
- [105] Ernst G. Straus. Addition chains of vectors (problem 5125). *Amer. Math. Monthly*, 70:806–808, 1964.
- [106] Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 71–89. Springer, Heidelberg, August 2013.
- [107] Justin Thaler, Mike Roberts, Michael Mitzenmacher, and Hanspeter Pfister. Verifiable computation with massively parallel interactive proofs. In *HotCloud*, June 2012.
- [108] Victor Vu, Srinath T. V. Setty, Andrew J. Blumberg, and Michael Walfish. A hybrid architecture for interactive verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 223–237. IEEE Computer Society Press, May 2013.
- [109] Riad S. Wahby, Max Howald, Siddharth J. Garg, abhi shelat, and Michael Walfish. Verifiable ASICs. In *2016 IEEE Symposium on Security and Privacy*, pages 759–778. IEEE Computer Society Press, May 2016.

- [110] Riad S. Wahby, Ye Ji, Andrew J. Blumberg, abhi shelat, Justin Thaler, Michael Walfish, and Thomas Wies. Full accounting for verifiable outsourcing. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 2071–2086. ACM Press, October / November 2017.
- [111] Riad S. Wahby, Srinath T. V. Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. Efficient RAM and control flow in verifiable outsourced computation. In *NDSS 2015*. The Internet Society, February 2015.
- [112] Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zk-SNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy*, pages 926–943. IEEE Computer Society Press, May 2018.
- [113] Abraham Waksman. A permutation network. *Journal of the ACM*, 15(1):159–163, January 1968.
- [114] Michael Walfish and Andrew J. Blumberg. Verifying computations without reexecuting them: from theoretical possibility to near practicality. *Communications of the Association for Computing Machinery*, February 2015.
- [115] Benjamin Wesolowski. Efficient verifiable delay functions. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 379–407. Springer, Heidelberg, May 2019.
- [116] Howard Wu, Wenting Zheng, Alessandro Chiesa, Raluca Ada Popa, and Ion Stoica. DIZK: A distributed zero knowledge proof system. In William Enck and Adrienne Porter Felt, editors, *USENIX Security 2018*, pages 675–692. USENIX Association, August 2018.
- [117] Tiancheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. Libra: Succinct zero-knowledge proofs with optimal prover computation. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 733–764. Springer, Heidelberg, August 2019.
- [118] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. vSQL: Verifying arbitrary SQL queries over dynamic outsourced databases. In *2017 IEEE Symposium on Security and Privacy*, pages 863–880. IEEE Computer Society Press, May 2017.
- [119] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. vRAM: Faster verifiable RAM with program-independent preprocessing. In *2018 IEEE Symposium*
- on Security and Privacy*, pages 908–925. IEEE Computer Society Press, May 2018.

## A Proof of MultiSwap Consistency

Let  $\sigma$  denote a multiset of swaps. Let  $\text{in}_\sigma$  denote  $\{y : (x, y) \in \sigma\}$  and let  $\text{rm}_\sigma$  denote  $\{x : (x, y) \in \sigma\}$ .

**Claim 1.** *Let  $\sigma$  be a multiset of swaps and  $\sigma^c$  be a cycle.  $\text{MultiSwap}(S, \sigma \uplus \sigma^c, S')$  holds if and only if  $\text{MultiSwap}(S, \sigma, S')$  does.*

*Proof:* We prove both directions simultaneously, by illustrating a bidirectional chain of mutually implicating equalities. We start with the definition of  $\text{MultiSwap}(S, \sigma \uplus \sigma^c, S')$ :

$$\begin{aligned} S' &= S \uplus \text{in}_{\sigma \uplus \sigma^c} \boxminus \text{rm}_{\sigma \uplus \sigma^c} \\ S' &= S \uplus \text{in}_{\sigma^c} \uplus \text{in}_\sigma \boxminus \text{rm}_{\sigma^c} \boxminus \text{rm}_\sigma \quad \text{properties of } \uplus, \boxminus \end{aligned}$$

Since  $\sigma^c$  is a cycle, we have that  $\text{in}_{\sigma^c} = \text{rm}_{\sigma^c}$ , so  $\text{rm}_{\sigma^c} \subseteq S \uplus \text{in}_{\sigma^c}$ , and the removal of  $\text{rm}_{\sigma^c}$  can be moved earlier

$$\begin{aligned} S' &= S \uplus \text{in}_{\sigma^c} \boxminus \text{rm}_{\sigma^c} \uplus \text{in}_\sigma \boxminus \text{rm}_\sigma \\ S' &= S \uplus \text{in}_\sigma \boxminus \text{rm}_\sigma \end{aligned}$$

This last line is exactly our goal: the statement that  $\text{MultiSwap}(S, \sigma, S')$  holds.  $\square$

**Claim 2.** *If  $\sigma$  contains no cycles and  $\text{MultiSwap}(S, \sigma, S')$  holds, then  $\sigma$  is sequentially consistent with respect to  $S$ , producing  $S'$ .*

*Proof:* Let  $n$  be the number of swaps in  $\sigma$ . For a set  $S$  and multiset of swaps  $\tau$ , define the directed multigraph  $G_{S, \tau}$  as a multigraph where the vertices are the universe of multiset elements, the edges point from each removal to its corresponding insertion, and each vertex is labeled with a *multiplicity* equal to the multiplicity of that vertex's element in  $S$ , minus the out-degree, plus the in-degree. Observe that in  $G = G_{S, \sigma}$ , the multiplicity of each vertex is equal to the multiplicity of that element in  $S'$ . Furthermore, by the predicate  $\text{MultiSwap}(S, \sigma, S')$  and the soundness of the proofs of insertions and removal, all multiplicities in  $G$  are non-negative.

We now construct the sequentially valid ordering of  $\sigma$ . Since  $\sigma$  has no swap cycles,  $G$  has no edge cycles. Thus, the edges of  $G$  can be topologically sorted such that all edges to a vertex occur before any edge from that vertex. We lift this edge order to a swap order, observing that in this swap order, all swaps inserting an element occur before all swaps removing it.

It suffices to show that when  $\sigma$  is applied to  $S$  in this order, each swap is valid. Let  $\sigma_i$  denote the first  $i$  elements of  $\sigma$  in the aforementioned order. Thus,  $G_{S, \sigma_i}$  is equal to  $G$ . Furthermore, the order ensures for all  $i > j$  and for all vertices  $v$ , the multiplicity of  $v$  in  $G_{S, \sigma_i}$  is at most the multiplicity of  $v$  in  $G_{S, \sigma_j}$ . Suppose that the  $i^{\text{th}}$  element of this order,  $(x_i, y_i)$  were

invalid, where  $i \leq n$ . This implies that the multiplicity of  $x_i$  in  $G_{S, \sigma_i}$  is negative. This would imply that the multiplicity of  $x_i$  in  $G_{S, \sigma_n} = G$  were negative, a contradiction. Thus no swap  $(x_i, y_i)$  is invalid in this order.  $\square$

**Proof of Lemma 1.** The reverse direction follows immediately from the definition of  $\text{MultiSwap}$ .

We prove the forward direction by (strong) induction on the size of  $\sigma$ . Say that  $\sigma$  has no cycles. Then the lemma follows from Claim 2. Otherwise, let  $\tau$  be a multiset of swaps and let  $\sigma^c$  be a cycle such that  $\sigma = \tau \uplus \sigma^c$ . By Claim 1,  $\text{MultiSwap}(S, \tau, S')$  holds. Then, by the inductive hypothesis,  $\tau$  can be decomposed into cycle-free  $\tau'$  and cycles  $\tau_i^c$  such that  $\tau = \tau' \uplus \uplus_i \tau_i^c$  and  $\tau'$  is sequentially consistent with respect to  $S$ , producing  $S'$ . By observing that  $\tau' \uplus (\uplus_i \tau_i^c) \uplus \sigma^c$  is a decomposition of  $\sigma$  into a cycle-free swap multiset and cycles, we conclude this direction of the proof.  $\square$

## B Parameter Values

Our RSA accumulators work in  $\mathbb{G} = \mathbb{Z}_N^\times / \{\pm 1\}$ , where  $N$  is the RSA-2048 challenge number [98],  $N=0xc7970ceedcc3b0754490201a7aa613cd73911081c790f5f1a8726f463550bb5b7ff0db8e1ea1189ec72f93d1650011bd721aeeacc2acde32a04107f0648c2813a31f5b0b7765ff8b44b4b6fffc93384b646eb09c7cf5e8592d40ea33c80039f35b4f14a04b51f7bdf781be4d1673164ba8eb991c2c4d730bbbe35f592bdef524af7e8daefd26c66fc02c479af89d64d373f442709439de66ceb955f3ea37d5159f6135809f85334b5cb1813addc80cd05609f10ac6a95ad65872c909525bdad32bc729592642920f24c61dc5b3c3b7923e56b16a4d9d373d8721f24a3fc0f1b3131f55615172866bcc30f95054c824e733a5eb6817f7bc16399d48c6361cc7e5.$

We randomly selected a 2048-bit offset  $\Delta$  for our division-intractable hash  $H_\Delta$  (§4.2); we use the value  $\Delta=0xf3709c40772816d668926cae548ffea31f49034ab1b30fb84b595ca6c126a6646a4341abea2f8b07bf8d366801ac293e5a286abb43accdec39ac8f0bc599519cf1e532f9c70b5406c4b652ca7da4e1cb102b69953841ae20d4bcab055c5338487ba00fe95e821abd381b191dfb77bae3e022ccd818d4064882d28481ffa2db45093a4deab05f6ebfbadcf11afe7369caeaaf1f02572348a17f0510b333b8a2d56e67d892f1e1182b26301d9347ae0a900cff2a0979caddb1a86e04a6cbc9704d6549e5b3aef0d5c3dc4aba648ed421b0ba37c3f8e8edc12ef42b86d8e5fbc0dbd903238ca2e9ed6873ccb68e8103b5d01b4249bfbe8e70cb4f4983f41df8c8f.$

Our evaluation (§7) builds on the BLS12-381 elliptic curve [27], which is the Barreto-Lynn-Scott curve [6] with parameter  $z = -0xd20100000010000$  whose subgroup order is  $p = 0x73eda753299d7d483339d80809a1d80553bda402fffe5bfeffffffffff0000001$ . This is the characteristic of the field  $\mathbb{F}_p$  for which we synthesize constraints.