

Multi-Party Revocation in Sovrin: Performance through Distributed Trust

Lukas Helminger^{1,2}, Daniel Kales¹, Sebastian Ramacher³, and Roman Walch^{1,2}

¹ Graz University of Technology, Graz, Austria
{lukas.helminger,daniel.kales,roman.walch}@iaik.tugraz.at

² Know-Center GmbH, Graz, Austria

³ AIT Austrian Institute of Technology, Vienna, Austria
sebastian.ramacher@ait.ac.at

Abstract. Accumulators provide compact representations of large sets and enjoy compact membership witnesses. Besides constant-size witnesses, public-key accumulators provide efficient updates of both the accumulator itself and the witness; however, they come with two drawbacks: they require a trusted setup and – without knowledge of the secret trapdoors – their performance is not practical for real-world applications with large sets. Recent improvements in the area of secure multi-party computation allow us to replace the trusted setup with a distributed generation of the public parameters.

In this paper, we introduce multi-party public-key accumulators dubbed dynamic linear secret-shared accumulators. We present versions of dynamic public-key accumulators in bilinear groups giving access to more efficient witness generation and update algorithms that utilize the shares of the secret trapdoors sampled by the parties generating the public parameters. Specifically, for the t -SDH-based accumulators, we provide a maliciously-secure variant sped up by a secure multi-party computation (MPC) protocol (IMACC’19) built on top of SPDZ. For this scheme, a performant proof-of-concept implementation is provided, which substantiates the practicability of public-key accumulators in this setting. With our implementation in two MPC frameworks, MP-SPDZ and FRESCO, we obtain more efficient accumulators for both medium-sized (2^{10}) and large (2^{14} and above) accumulated sets.

Finally, we explore applications of dynamic linear secret-shared accumulators to revocations schemes of group signatures and credentials system. In particular, we consider it as part of Sovrin’s system for anonymous credentials where credentials are issued by the a foundation of trusted nodes. Hence, our accumulators naturally fit this setting.

Keywords: multiparty computation, dynamic accumulators, distributed trust

1 Introduction

Digital identity management systems become an increasingly important cornerstone of digital workflows. Self-sovereign identity (SSI) systems such as Sovrin⁴ are of central interest as underlined by a recent push in the European Union for a cross-boarder SSI system.⁵ But all these systems face a similar issue, namely that of efficient revocation. Regardless of whether they are built from signatures, group signatures or anonymous credentials, such systems have to consider mechanisms to revoke a user's identity information. Especially for identity management systems with a focus on privacy, revocation may threaten those privacy guarantees. As such various forms of privacy-preserving revocations have emerged in the literature including approaches based on various forms of blacklists or whitelists including [67,78,22,23,24,6,25,39,68,69,42].⁶

One promising approach regarding efficiency represents the blacklist (or whitelist) via cryptographic accumulators which were introduced by Benaloh and de Mare [18]. They allow one to accumulate a finite set \mathcal{X} into a succinct value called the accumulator. For every element in the accumulated set, one can efficiently compute a witness certifying its membership, and additionally, some accumulators also support efficient non-membership witnesses. However, it should be computationally infeasible to find a membership witness for non-accumulated values and a non-membership witness for accumulated values, respectively. Accumulators facilitate privacy-preserving revocation mechanisms, which is especially relevant for privacy-friendly authentication mechanisms like group signatures and credentials. For a blacklist approach, the issuing authority accumulates all revoked users and users prove in zero-knowledge that they know a non-membership witness for their credential. Alternatively, for a whitelist approach, the issuing authority accumulates all users and users then prove in zero-knowledge that they know a membership witness. As both approaches may involve large lists, efficient updates of the accumulator as well as efficient proofs are important for building an overall efficient system. For example, in Sovrin [54] such an accumulator-based approach with whitelists following the ideas of [42] is used. Their credentials contain a unique revocation ID attribute, i_R , which are accumulated. Each user obtains a membership witness proofing that their i_R is contained in the accumulator. Once a credential is revoked, the corresponding i_R gets removed from the accumulator and all users have to update their proofs accordingly. The revoked user is not longer able to prove knowledge of a verifying witness and thus verification fails.

Accumulators in general are an important primitive and building block in many cryptographic protocols. In particular, Merkle trees [63] have seen many applications as accumulators in both the cryptographic literature but also in practice. For example, they have been used to implement Certificate Transparency (CT) [55,56,38] where all issued certificates are publicly logged, i.e.,

⁴ <https://sovrin.org/>

⁵ <https://www.tno.nl/en/about-tno/news/2019/12/essif-lab/>

⁶ For a discussion of approaches for group signatures, see, for example, [75].

accumulated, and then clients can check that the server’s certificate is included in a log by requesting witnesses from the log servers. Accumulators also find application in redactable signatures [73,35], credentials [25,7], ring, and group signatures [57,36,51], anonymous cash [64], authenticated hash tables [72], among many others.

Interestingly, when looking at applications of accumulators deployed in practice, many systems rely on Merkle trees. Most prominently we can observe this fact in CT. Even though new certificates are continuously added to the log, the system is designed around a Merkle tree that gets recomputed all the time instead of updating a dynamic public-key accumulator. The reason is two-fold: first, for dynamic accumulators to be efficiently computable, knowledge of the secret trapdoor used to generate the public parameters is required. Without this information, witness generation and accumulator updates are simply too slow for large sets (as recently observed in [50]). Secondly, in this setting it is of paramount importance that the log servers do not have access to the secret trapdoor. Otherwise malicious servers would be able to present membership witnesses for each and every certificate even if it was not included in the log.

The latter issue can also be observed in other applications of public-key accumulators. The approaches due to Garman et al. [42] and the one used in Sovrin rely on the Strong-RSA accumulator and the t -SDH accumulators, respectively. Both these accumulators have trapdoors: in the first case the factorization of the RSA modulus and in the second case a secret exponent. Therefore, the security of the system requires those trapdoors to stay secret. Hence, in these types of protocols, the generation of the public parameters is problematic. In fact, it requires to put significant trust in the parties generating those parameters. If they would act maliciously and not delete the secret trapdoors, they would be able to break all these protocols in one way or another. To circumvent this problem, Sander [74] proposed a variant of an RSA-based accumulator from RSA moduli with unknown factorization. Alternatively, secure multi-party computation (MPC) protocols make it possible to compute the public parameters and thereby replace the trusted third party. As long as a large enough subset of parties participating is honest, the secret trapdoor is not available to malicious parties. Over the years, efficient solutions for distributed parameter generation have emerged, e.g., for distributed RSA key generation [40], or distributed ECDSA key generation [58]. One very prominent and wildly publicized example of such an approach is the “ceremony” executed for Zcash, where an MPC protocol involving hundreds of participants was used to generate the public parameters for the proof system [21].

Based on the recent progress in efficient MPC protocols, we ask the following question: *what if the parties kept their shares of the secret trapdoor?* Are the algorithms of the public-key accumulators exploiting knowledge of the secret trapdoor faster if performed within an (maliciously-secure) MPC protocol than their variants relying only on the public parameters?

1.1 Our Techniques

We give a short overview of how our construction works which allows us to positively answer this question for accumulators in the discrete logarithm setting. Let us consider the accumulator based on the t -SDH assumption. The construction is based on the fact that given powers $g^{s^i} \in \mathbb{G}$ for all i up to t where $s \in \mathbb{Z}_p$ is unknown, it is possible to evaluate polynomials $f \in \mathbb{Z}_p[X]$ up to degree t at s in the exponent, i.e., $g^{f(s)}$. To evaluate the polynomial in the exponent, one takes the coefficients of the polynomial, i.e., $f = \sum_{i=0}^t a_i X^i$, and computes $g^{f(s)}$ as

$$\prod_{i=0}^t \left(g^{s^i}\right)^{a_i}.$$

The idea of the accumulator is to then define a polynomial involving all elements of the set as roots. The accumulator is then comprised of that polynomial evaluated at s in the exponent. A witness is simply the corresponding factor canceled out, i.e., $g^{f(s)(s-x)^{-1}}$. Verification of the witness is performed by checking whether the corresponding factor and the witness match $g^{f(s)}$ using a pairing.

If s is known, all the computations are significantly more efficient: $f(s)$ can be directly evaluated in \mathbb{Z}_p and then the generation of the accumulator only requires one exponentiation in \mathbb{G} . Similarly, computation of the witness also only requires one exponentiation in \mathbb{G} , since $(s-x)^{-1}$ can first be computed in \mathbb{Z}_p . For the latter, the asymptotic runtime is thereby reduced from $\mathcal{O}(|\mathcal{X}|)$ (i.e., linear in the number of accumulated elements) to $\mathcal{O}(1)$. But the improvement comes at a cost: if s is known, witnesses for non-members can be produced.

On the other hand, if s is first produced by multiple parties in an additively secret-shared fashion, these parties can cooperate in a secret-sharing based MPC protocol. Thereby, all the computations can still benefit from the knowledge of s . Indeed, the parties would compute their share of $g^{f(s)}$ and $g^{f(s)(s-x)^{-1}}$ respectively and thanks to the partial knowledge of s could still perform all operations – except the final exponentiation – in \mathbb{Z}_p .

1.2 Our Contribution

Our contributions can be summarized as follows:

- Starting from the very recent treatment of accumulators in the UC model [27] by Baldimtsi et al. [8], we introduce the notion of *linear secret-shared accumulators*. As the name suggests, it covers accumulators where the trapdoor is available in a linearly secret-shared fashion with multiple parties running the parameter generation as well as the algorithms that profit from the availability of the trapdoor. Since the MPC literature discusses security in the UC model, we also chose to do so for our accumulators.
- Based on recent improvements on distributed key generation of discrete logarithms, we provide dynamic public-key accumulators without trusted setup. During the parameter generation, the involved parties keep their shares of the secret trapdoor. Consequently, we can provide MPC protocols secure in

the semi-honest and the malicious security model, respectively, implementing the algorithms for accumulator generation, witness generation, and accumulator updates exploiting the shares of the secret trapdoor. To the best of our knowledge, this is the first work that uses secure multiparty computation to build distributed cryptographic accumulators and in particular, linear secret-shared accumulators.

- We provide a protocol for both semi-honest and malicious security models, for t -SDH accumulators, and especially for the dynamic accumulator due to Derler et al [34], which is based on the accumulator by Nguyen [70]. In particular, this protocol enables updates to the accumulator independent of the size of the accumulated set. For increased efficiency, we also transport this accumulator to the Type-3 pairing setting. Due to the structure of the bilinear groups setting, the construction nicely generalizes to any number of parties.
- We provide a proof-of-concept implementation of our protocols in two MPC frameworks, MP-SPDZ [77] and FRESCO [3]. We evaluate the efficiency of our protocols and compare them to the performance of an implementation, having no access to the secret trapdoors as usual for the public-key accumulators. We evaluate our protocol in the LAN and WAN setting in the semi-honest and malicious security model for various choices of parties and accumulator sizes. For the latter, we choose sizes up to 2^{14} . Specifically, for the t -SDH accumulator, we observe the expected $\mathcal{O}(1)$ runtimes for witness creation and accumulator updates, which cannot be achieved without access to the trapdoor. Notably, for the tested numbers of up to 5 parties, the MPC-enabled accumulator creation algorithms are faster for 2^{10} elements in the LAN setting than its non-MPC counterpart (without access to the secret trapdoor); for 2^{14} elements the algorithms are also faster in the WAN setting. We expect even greater improvements as the size of the accumulated set grows further.

On top of that, we discuss how our proposed MPC-based accumulators might impact revocation in distributed credential systems such as Sovrin [54]. In this scenario, the trust in the nodes run by the Sovrin foundation members can further be reduced. We also discuss applications like CT and the privacy-preserving extension based on private information retrieval (PIR) [50]. In particular, the size of the witnesses stored in certificates or sent as part of the TLS handshake is significantly reduced without running into performance issues.

1.3 Related Work

When cryptographic protocols are deployed that require the setup of public parameters by a trusted third party, issues similar to those mentioned for public-key accumulators may arise. As discussed before, especially cryptocurrencies had to come up with ways to circumvent this problem for accumulators but also the common reference string (CRS) of zero-knowledge SNARKs [26]. Here, trust in the CRS is of paramount importance on the verifier side to prevent malicious

provers from cheating. But also provers need to trust the CRS as otherwise zero-knowledge might not hold. We note that there are alternative approaches, namely subversion-resilient zk-SNARKS [16] to reduce the trust required in the CRS generator. However, subversion-resilient soundness and zero-knowledge at the same time has been shown to be impossible by Bellare et al. Abdolmaleki et al. [1] provided a construction of zk-SNARKS, which was later improved by Fuchsbauer [41], achieving subversion zero-knowledge, by adding a verification algorithm for the CRS and then only the verifier needs to trust the correctness of the CRS. Groth et al. [47] recently introduced the notion of an updatable CRS where first generic compilers [2] are available to lift any zk-SNARK to an updatable simulation sound extractable zk-SNARK. There the CRS can be updated and if the initial generation or one of the updates was done honestly, neither soundness nor zero-knowledge can be subverted. In the random oracle model (ROM), those considerations become less of a concern and the trust put into the CRS can be minimized, e.g., as done in the construction of STARKs [17].

Approaches that try to fix the issue directly in the formalization of accumulators and corresponding constructions have also been studied. For example, Lipmaa [60] proposed a modified model tailored to the hidden order group setting. In this model, the parameter setup is split into two algorithms, `Setup` and `Gen`, whereas the adversary can control access the trapdoors output by `Setup`, but cannot influence nor access the randomness used by `Gen`. However, secure constructions in this model so far have been provided for rather unstudied assumptions based on modules over Euclidean rings, and are not applicable to the efficient standard constructions we are interested in. More recently, Boneh et al. [20] revisited the RSA accumulator without trapdoor which allows the accumulator to be instantiated from unknown order groups without trusted setup such as class groups of quadratic imaginary orders [48] and hyperelliptic curves of genus 2 or 3 [37]. While this line of research has already shown promising results, some functionalities, e.g., creating membership witnesses, are still not practical for large sets.

The area of secure multiparty computation has seen a lot of interest both in improving the MPC protocols itself to a wide range of practical applications. In particular, SPDZ [33,30] has seen a lot of interest, improvements and extensions [52,53,28,71]. This interest also lead to multiple MPC frameworks, e.g. MP-SPDZ [77], FRESCO [3] and SCALE-MAMBA [4], enabling easy prototyping for researchers as well as developers. For practical applications of MPC, one can observe first MPC-based systems turned into products such as Unbound’s virtual hardware security model (HSM).⁷ For such a virtual HSM, one essentially wants to provide distributed key generation [40] together with threshold signatures [31] allowing to replace a physical HSM. Similar techniques are also interesting for securing wallets for the use in cryptocurrencies, where especially protocols for ECDSA [44,43,59] are of importance to secure the secret key material. Similarly, such protocols are also of interest for securing the secret key material of internet infrastructure such as DNSSEC [29]. Additionally, addressing

⁷ <https://www.unboundtech.com/usecase/virtual-hsm/>

privacy concerns in machine learning algorithms has become increasingly popular recently, with MPC protocols being one of the building blocks to achieve private classification and private model training [66,79,14,65]. Recent works [76] also started to generalize the algorithms that are used as parts of those protocols allowing group operations on elliptic curve groups with secret exponents or secret group elements.

1.4 Paper Organization

This work is structured as follows: first we recap some basics on the UC model and MPC protocols in Section 2. In Section 3 we recall the definition of cryptographic accumulators and instantiations. Then we introduce dynamic linear secret-shared accumulators and constructions in Section 4, which is followed by the evaluation of our implementation in Section 5. Finally, in Section 6, we discuss the application of the construction in the context of certificate transparency.

2 Preliminaries

In this section, we introduce cryptographic primitives and constructions that we subsequently use as building blocks. Notation-wise, let $[n] := \{1, \dots, n\}$ for $n \in \mathbb{N}$. For an algorithm \mathcal{A} , we write $\mathcal{A}(\dots; r)$ to make the random coins explicit. We say that an algorithm is efficient, if it runs in probabilistic polynomial time (PPT). For the cryptographic assumptions, we refer to Appendix A.

2.1 UC security and ABB

In this paper, we mainly work in the UC model first introduced by Canetti [27]. The success of the UC model stems from its universal composition theorem, which, informally speaking, states that it is safe to use a secure protocol as a sub-protocol in a more complex one. This strong statement enables one to analyze and prove the security of involved protocols in a modular way, allowing us to build upon work that was already proven to be secure in the UC model. In preparation for the security analysis of our MPC accumulators, we recall the definition of the UC model.

Definition 1 ([27]). *Let $EXEC_{\Pi, \mathcal{A}, \mathcal{E}}$ respectively $EXEC_{\mathcal{F}, SLM, \mathcal{E}}$ denote the random variables describing the output of environment \mathcal{E} when interacting with an adversary \mathcal{A} and parties performing protocol Π , respectively when interacting with a simulator SLM and an ideal functionality \mathcal{F} . Protocol Π UC emulates the ideal functionality \mathcal{F} if for any adversary \mathcal{A} there exists a simulator SLM such that, for any environment \mathcal{E} the distribution of $EXEC_{\Pi, \mathcal{A}, \mathcal{E}}$ and $EXEC_{\mathcal{F}, SLM, \mathcal{E}}$ are indistinguishable.*

The importance of the UC model for secure multiparty computation stems from the arithmetic black box (ABB) as introduced by Damgård and Nielsen [32]. The ABB models a secure general-purpose computer in the UC model. It allows

performing arithmetic operations on private input provided by the parties. The result of these operations is then revealed to all parties. Working with the ABB provides us with a tool of abstracting arithmetic operations, including addition and multiplication in fields.

2.2 SPDZ and Derived Protocols

Our protocols build upon SPDZ [33,30], a concrete implementation of the abstract ABB. SPDZ itself is based on an additive secret-sharing over a finite field \mathbb{F}_p with information-theoretic MACs making the protocol statistically UC secure against an active adversary corrupting all but one player. We will denote the ideal functionality of SPDZ (the online protocol) by $\mathcal{F}_{\text{SPDZ}}$. For an easy use of the SPDZ protocol later in our accumulators, we give a high-level description of the functionality together with an intuitive notation. We assume that the computations are performed by n parties and by $\langle s \rangle \in \mathbb{F}_p$ denote a secret-shared value between the parties in a finite field with p elements, where p is prime. The ideal functionality $\mathcal{F}_{\text{SPDZ}}$ provides us with the following operations:

Add($\langle a \rangle, \langle b \rangle$): The parties locally compute $\langle a + b \rangle \leftarrow \langle a \rangle + \langle b \rangle$.

Multiply($\langle a \rangle, \langle b \rangle$): A secret-shared multiplication triple set up during preprocessing is used in a 1 round interactive protocol to compute $\langle ab \rangle \leftarrow \langle a \rangle \cdot \langle b \rangle$.

sRand(\mathbb{F}_p): Samples $\langle r \rangle \xleftarrow{R} \mathbb{F}_p$.

Open($\langle a \rangle$) The value a is send to every party.

The ability to add two secret-shared values locally implies linear functions can be computed locally. More precisely, given two constants $a, b \in \mathbb{F}_p$ and a secret-shared value $\langle x \rangle \in \mathbb{F}_p$ the parties can non-interactively compute $\langle c \rangle \leftarrow a\langle x \rangle + b$. Further, we assume for convenience in describing the main protocol that we have access to the following function derived from $\mathcal{F}_{\text{SPDZ}}$.

Inverse($\langle a \rangle$): Joint secret randomness set up during preprocessing is used in a 1 round interactive protocol to compute $\langle a^{-1} \rangle$

Computation of the inverse can be efficiently implemented using a standard form of masking as first done in [9]. That is, given a joint secret randomness $\langle r \rangle \in \mathbb{F}_p$, compute $\langle z \rangle \leftarrow \langle ra \rangle$ and immediately open the result and invert it in plain. The inverse of $\langle a \rangle$ is then $z^{-1}\langle r \rangle$. However, there is a small failure probability if either a or r is zero. In our case, the field size is large enough that the probability of a random element being zero is negligible.

There is one additional sub-protocol which we will often need and therefore explain here. Recently, Smart et al. [76] introduced protocols – in particular based on SPDZ – for group operations of elliptic curve groups supporting secret exponents and secret group elements. The central high-level idea is to use the original SPDZ in the exponent group and for the authentication of the shares of an elliptic curve point a similar protocol as in SPDZ to compute the MACs. For this work, we only need the protocol Multiply-G-P for exponentiating a public point with a secret exponent. We rewrite it in a slightly more general version to

fit our setting. Let \mathbb{G} be a cyclic group of prime order p and $g \in \mathbb{G}$. Further, let $\langle a \rangle \in \mathbb{F}_p$ be a secret-shared exponent.

$\text{Exp}_{\mathbb{G}}(\langle a \rangle, g)$: The parties locally compute $\langle g^a \rangle \leftarrow g^{\langle a \rangle}$.

Since the security proof (in the UC model) of this sub-protocol in [76] does not use any exclusive property of an elliptic curve group, it automatically translates to any cyclic group of prime order.

At this point, we want to summarize that all protocols discussed so far are secure in the UC model, making them safe to use in our accumulators as sub-protocols. Therefore, we will refer to their ideal functionality as $\mathcal{F}_{\text{SPDZ+}}$. As a result, our protocols become secure in the UC model as long as we do not reveal any intermediate values.

3 Accumulators

3.1 Definitions

We rely on the formalization of accumulators by Derler et al. [34]. Based on this formalization, we then state the bilinear accumulator within this framework. We start with the definition of a static accumulator and then recall the definition of a dynamic accumulator.

Definition 2 (Static Accumulator). *A static accumulator is a tuple of efficient algorithms $(\text{Gen}, \text{Eval}, \text{WitCreate}, \text{Verify})$ which are defined as follows:*

$\text{Gen}(1^\kappa, t)$: *This algorithm takes a security parameter κ and a parameter t . If $t \neq \infty$, then t is an upper bound on the number of elements to be accumulated. It returns a key pair $(\text{sk}_A, \text{pk}_A)$, where $\text{sk}_A = \emptyset$ if no trapdoor exists. We assume that the accumulator public key pk_A implicitly defines the accumulation domain D_A .*

$\text{Eval}((\text{sk}_A, \text{pk}_A), \mathcal{X})$: *This algorithm takes a key pair $(\text{sk}_A, \text{pk}_A)$ and a set \mathcal{X} to be accumulated and returns an accumulator $\Lambda_{\mathcal{X}}$ together with some auxiliary information aux .*

$\text{WitCreate}((\text{sk}_A, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x_i)$: *This algorithm takes a key pair $(\text{sk}_A, \text{pk}_A)$, an accumulator $\Lambda_{\mathcal{X}}$, auxiliary information aux and a value x_i . It returns \perp , if $x_i \notin \mathcal{X}$, and a witness wit_{x_i} for x_i otherwise.*

$\text{Verify}(\text{pk}_A, \Lambda_{\mathcal{X}}, \text{wit}_{x_i}, x_i)$: *This algorithm takes a public key pk_A , an accumulator $\Lambda_{\mathcal{X}}$, a witness wit_{x_i} and a value x_i . It returns 1 if wit_{x_i} is a witness for $x_i \in \mathcal{X}$ and 0 otherwise.*

Definition 3 (Dynamic Accumulator). *A dynamic accumulator is a static accumulator with an additional tuple of efficient algorithms $(\text{Add}, \text{Delete}, \text{WitUpdate})$ which are defined as follows:*

$\text{Add}((\text{sk}_A, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x)$: *This algorithm takes a key pair $(\text{sk}_A, \text{pk}_A)$, an accumulator $\Lambda_{\mathcal{X}}$, auxiliary information aux , as well as an element x to be added. If $x \in \mathcal{X}$, it returns \perp . Otherwise, it returns the updated accumulator $\Lambda_{\mathcal{X}'}$ with $\mathcal{X}' \leftarrow \mathcal{X} \cup \{x\}$ and updated auxiliary information aux' .*

$\text{Delete}((\text{sk}_A, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x)$: This algorithm takes a key pair $(\text{sk}_A, \text{pk}_A)$, an accumulator $\Lambda_{\mathcal{X}}$, auxiliary information aux , as well as an element x to be added. If $x \notin \mathcal{X}$, it returns \perp . Otherwise, it returns the updated accumulator $\Lambda_{\mathcal{X}'}$ with $\mathcal{X}' \leftarrow \mathcal{X} \setminus \{x\}$ and updated auxiliary information aux' .

$\text{WitUpdate}((\text{sk}_A, \text{pk}_A), \text{wit}_{x_i}, \text{aux}, x)$: This algorithm takes a key pair $(\text{sk}_A, \text{pk}_A)$, a witness wit_{x_i} to be updated, auxiliary information aux and an x which was added to/deleted from the accumulator, where aux indicates addition or deletion. It returns an updated witness wit'_{x_i} on success and \perp otherwise.

Note that the formalization of accumulators by Derler et al. gives access to a trapdoor if it exists. If the trapdoor is not available, then sk_A is set to \emptyset . Giving those algorithms access to the trapdoor can often be beneficial performance-wise, but requires additional trust assumptions as we have discussed before.

Finally, we recall the notion of collision freeness:

Definition 4 (Collision Freeness). A cryptographic accumulator is collision-free, if for all PPT adversaries \mathcal{A} there is a negligible function $\varepsilon(\cdot)$ such that:

$$\Pr \left[\begin{array}{l} (\text{sk}_A, \text{pk}_A) \leftarrow \text{Gen}(1^\kappa, t), \\ (\text{wit}_{x_i^*}, x_i^*, \mathcal{X}^*, r^*) \leftarrow \mathcal{A}^{\mathcal{O}}(\text{pk}_A) : \\ \text{Verify}(\text{pk}_A, \Lambda^*, \text{wit}_{x_i^*}, x_i^*) = 1 \wedge x_i^* \notin \mathcal{X}^* \end{array} \right] \leq \varepsilon(\kappa),$$

where $\Lambda^* \leftarrow \text{Eval}((\text{sk}_A, \text{pk}_A), \mathcal{X}^*; r^*)$ and the adversary gets access to the oracles

$$\mathcal{O} = \{\text{Eval}((\text{sk}_A, \text{pk}_A), \cdot), \text{WitCreate}((\text{sk}_A, \text{pk}_A), \cdot, \cdot, \cdot)\}$$

and, if the accumulator is dynamic, additionally to

$$\{\text{Add}((\text{sk}_A, \text{pk}_A), \cdot, \cdot, \cdot), \text{Delete}((\text{sk}_A, \text{pk}_A), \cdot, \cdot, \cdot), \\ \text{WitUpdate}((\text{sk}_A, \text{pk}_A), \cdot, \cdot, \cdot)\}.$$

3.2 Pairing-based Accumulator

We recall the t -SDH-based accumulator from [34], which is based on the accumulator by Nguyen [70]. The idea here is to encode the accumulated elements in a polynomial. This polynomial is then evaluated for a fixed element and the result is randomized to obtain the accumulator. A witness consists of the evaluation of the same polynomial with the term corresponding to the respective element cancelled out. For verification, a pairing evaluation is used to check whether the polynomial encoded in the witness is a factor of the one encoded in the accumulator.

As it is typically more efficient to work in the Type-3 setting, we state this accumulator in the asymmetric pairing. The scheme is presented accumulator in Scheme 1.

Remark 1 (On the hash function). Note that for support of arbitrary accumulation domains, the accumulator requires a suitable hash function mapping to \mathbb{Z}_p^* . For the MPC-based accumulators that we will define later, it is clear that the hash function can be evaluated in public. Hence, for simplicity, we omit the hash function in our discussion.

<p>Gen($1^\kappa, t$): Let $\text{BG} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2) \leftarrow \text{BGen}(\kappa)$. Choose $s \xleftarrow{R} \mathbb{Z}_p^*$ and return $\text{sk}_A \leftarrow s$ and $\text{pk}_A \leftarrow (\text{BG}, (g_1^{s^i})_{i=1}^t, g_2^s)$.</p>
<p>Eval((sk_A, pk_A), \mathcal{X}): Parse \mathcal{X} as subset of \mathbb{Z}_p^*. Choose $r \xleftarrow{R} \mathbb{Z}_p^*$. If $\text{sk}_A \neq \emptyset$, compute $\Lambda_{\mathcal{X}} \leftarrow g_1^{r \prod_{x \in \mathcal{X}} (x+s)}$. Otherwise, expand the polynomial $\prod_{x \in \mathcal{X}} (x+X) = \sum_{i=0}^n a_i X^i$, and compute $\Lambda_{\mathcal{X}} \leftarrow ((\prod_{i=0}^n g_1^{s^i})^{a_i})^r$. Return $\Lambda_{\mathcal{X}}$ and $\text{aux} \leftarrow (\text{add} \leftarrow 0, r, \mathcal{X})$.</p>
<p>WitCreate((sk_A, pk_A), $\Lambda_{\mathcal{X}}, \text{aux}, x$): Parse aux as (r, \mathcal{X}). If $x \notin \mathcal{X}$, return \perp. If $\text{sk}_A \neq \emptyset$, compute and return $\text{wit}_x \leftarrow \Lambda_{\mathcal{X}}^{(x+s)^{-1}}$. Otherwise, run $(\text{wit}_x, \dots) \leftarrow \text{Eval}((\text{sk}_A, \text{pk}_A), \mathcal{X} \setminus \{x\}; r)$, and return wit_x.</p>
<p>Verify($\text{pk}_A, \Lambda_{\mathcal{X}}, \text{wit}_x, x$): If $e(\Lambda_{\mathcal{X}}, g_2) = e(\text{wit}_x, g_2^x \cdot g_2^s)$ holds, return 1, otherwise return 0.</p>
<p>Add((sk_A, pk_A), $\Lambda_{\mathcal{X}}, \text{aux}, x$): Parse aux as (r, \mathcal{X}). If $x \in \mathcal{X}$, return \perp. Set $\mathcal{X}' \leftarrow \mathcal{X} \cup \{x\}$. If $\text{sk}_A \neq \emptyset$, compute and return $\Lambda_{\mathcal{X}'} \leftarrow \Lambda_{\mathcal{X}}^{x+s}$ and $\text{aux}' \leftarrow (r, \mathcal{X}', \text{add} \leftarrow 1, \Lambda_{\mathcal{X}}, \Lambda_{\mathcal{X}'})$. Otherwise, return $\text{Eval}((\text{sk}_A, \text{pk}_A), \mathcal{X}'; r)$ with aux extended with $(\text{add} \leftarrow 1, \Lambda_{\mathcal{X}}, \Lambda_{\mathcal{X}'})$.</p>
<p>Delete((sk_A, pk_A), $\Lambda_{\mathcal{X}}, \text{aux}, x$): Parse aux as (r, \mathcal{X}). If $x \notin \mathcal{X}$, return \perp. Set $\mathcal{X}' \leftarrow \mathcal{X} \setminus \{x\}$. If $\text{sk}_A \neq \emptyset$, compute and return $\Lambda_{\mathcal{X}'} \leftarrow \Lambda_{\mathcal{X}}^{(x+s)^{-1}}$ and $\text{aux}' \leftarrow (r, \mathcal{X}', \text{add} \leftarrow -1, \Lambda_{\mathcal{X}}, \Lambda_{\mathcal{X}'})$. Otherwise, return $\text{Eval}((\text{sk}_A, \text{pk}_A), \mathcal{X}'; r)$ with aux extended with $(\text{add} \leftarrow 0, \Lambda_{\mathcal{X}}, \Lambda_{\mathcal{X}'})$.</p>
<p>WitUpdate((sk_A, pk_A), $\text{wit}_{x_i}, \text{aux}, x$): Parse aux as $(r, \mathcal{X}', \text{add}, \Lambda_{\mathcal{X}}, \Lambda_{\mathcal{X}'})$. If $\text{add} = 0$, return \perp. Return $\Lambda_{\mathcal{X}} \cdot \text{wit}_{x_i}^{x-x_i}$ if $\text{add} = 1$. If instead $\text{add} = -1$, return $(\Lambda_{\mathcal{X}'}^{-1} \cdot \text{wit}_{x_i})^{(x-x_i)^{-1}}$. In the last two cases in addition return $\text{aux} \leftarrow (\text{add} \leftarrow 0)$.</p>

Scheme 1: t -SDH-based accumulator in the Type-3 setting.

Correctness is clear. The proof of collision freeness follows from the t -SDH assumption. For completeness, we still restate the theorem from [34] adopted to the Type-3 setting:

Theorem 1. *If the t -SDH assumption holds relative to BGen , then Scheme 1 is collision-free.*

For the proof, we refer to Appendix B.1.

3.3 UC Secure Accumulators

Only recently, Baldimtsi et al. [8] formalized the security of accumulators in the UC framework. Interestingly, they showed, that any correct and collision-free standard accumulator is automatically UC secure. We, however, want to note, that their definitions of accumulators are slightly different then the framework by Derler et al. (which we are using). Hence, we adapt the ideal functionality \mathcal{F}_{Acc} from [8] in the following way.

First our ideal functionality \mathcal{F}_{Acc} consists of two more sub-functionalities. This is due to a separation of the algorithms responsible for the evaluation, addition, and deletion. Secondly, our \mathcal{F}_{Acc} is simplified to our purpose, whereas

\mathcal{F}_{Acc} from Baldimtsi et al. is in their words “an entire menu of functionalities covering all different types of accumulators ...”. Thirdly, we added identity checks to sub-functionalities (where necessary) to be consistent with the given definitions of accumulators.

The resulting ideal functionality is depicted in Functionality 1. Note that the ideal functionality has up to three parties. First, the party which holds the set \mathcal{X} is the accumulator manager \mathcal{AM} , responsible for the algorithms Gen, Eval, WitCreate, Add and Delete. The second party \mathcal{H} owns a witness and is interested in keeping it updated and for this reason, performs the algorithm WitUpdate. The last party \mathcal{V} can be seen as an external party. \mathcal{V} is only able to use Verify to check the membership of an element in the accumulated set.

In the following theorem we adapt the proof from [8] to our setting:

Theorem 2. *Let $\Pi_{Acc} = (\text{Gen}, \text{Eval}, \text{WitCreate}, \text{Verify}, \text{Add}, \text{Delete}, \text{WitUpdate})$ be a correct and collision-free dynamic accumulator scheme (in the sense of Definition 3), and let Verify be deterministic. Then Π_{Acc} UC emulates \mathcal{F}_{Acc} .*

Proof. We will proceed by contraposition. Assume to the contrary that Π_{Acc} does not UC emulate \mathcal{F}_{Acc} , i.e., there exists an environment \mathcal{E} , for all simulators SIM such that \mathcal{E} can distinguish between the distributions of the random variables $EXEC_{\mathcal{F}, SIM, \mathcal{E}}$ and $EXEC_{\Pi, \mathcal{A}, \mathcal{E}}$ with non-negligible probability. Since the last statement holds for all simulators, we can choose one. We want a simulator SIM that interacts with the ideal functionality \mathcal{F}_{Acc} in a way such that their distribution can not be distinguished by any environment from the real world, except when it violates either correctness or collision-freeness. Such a simulator would be a contradiction to our assumption and thereby prove the theorem.

Consider a simulator SIM that uses the standard corruption model from [27]. Further, SIM interacts with the environment \mathcal{E} by forwarding any input to the real adversary \mathcal{A} and conversely forwarding any output from \mathcal{A} directly to \mathcal{E} . When SIM receives the request (GEN, sid) from \mathcal{F}_{Acc} , it replies with the actual accumulator algorithms.

By construction of SIM , the only differences to the real world that are visible for the environment \mathcal{E} are the following instances where \mathcal{F}_{Acc} returns \perp : (i) WitCreate: 4., (ii) Verify: 1.b, (iii) Add: 6. and (iv) WitUpdate : 3.

The occurrence of one of the above cases would immediately imply a violation of the classical definition. More concretely, if Verify 1.b would return \perp , then the collision-freeness would be violated. In the other instances, correctness would not be given any more.

As a direct consequence of Theorems 1 and 2, the accumulator from Scheme 1 is also secure in the UC model of [8] since it is correct and collision-free:

Corollary 1. *Scheme 1 emulates \mathcal{F}_{Acc} in the UC model.*

GEN: On input (gen, sid) from \mathcal{AM} the functionality does the following:

1. If this is not the first gen command, or if sid does not encode the identity of \mathcal{AM} , ignore this command. Otherwise, continue.
2. $t \leftarrow 0$.
3. Initialize an empty list A (keeps track of all accumulator states).
4. Initialize map S , and set $S[0] \leftarrow \emptyset$ (maps operation counters to current accumulated sets).
5. Send (gen, sid) to SIM .
6. Get $(algorithms, sid, (Gen, Eval, WitCreate, Verify, Add, Delete, WitUpdate))$ from SIM . Their expected input output behavior is described in Definition 3. All of them should be polynomial-time and Verify should be deterministic.
7. Run $(sk, pk) \leftarrow Gen(1^\lambda)$.
8. Store sk, pk ; add $\Lambda_\emptyset \leftarrow \emptyset$ to A .
9. Send $(algorithms, sid, (Gen, Eval, WitCreate, Verify, Add, Delete, WitUpdate))$ to \mathcal{AM} .

EVAL: On input $(eval, sid, \mathcal{X})$ from \mathcal{AM} the functionality does the following:

1. If sid does not encode the identity of \mathcal{AM} , or if $\mathcal{X} \not\subseteq D_A$, ignore this command. Otherwise, continue.
2. $t \leftarrow t + 1$, and $S[t] \leftarrow \mathcal{X}$.
3. Run $(\Lambda_{\mathcal{X}}, aux) \leftarrow Eval((sk, pk), \mathcal{X})$.
4. Store aux ; add $\Lambda_{\mathcal{X}}$ to A .
5. Send $(eval, sid, \Lambda_{\mathcal{X}}, \mathcal{X})$ to \mathcal{AM} .

WITCREATE: On input $(witcreate, sid, x)$ from \mathcal{AM} the functionality does the following:

1. If sid does not encode the identity of \mathcal{AM} , or if $x \notin D_A$, ignore this command. Otherwise, continue.
2. Run $w \leftarrow WitCreate((sk, pk), \Lambda_{\mathcal{X}}, aux, x)$.
3. If $x \notin S[t]$, send \perp to \mathcal{AM} and halt. Otherwise, continue.
4. If $Verify(pk, \Lambda_{\mathcal{X}}, w, x) = 1$ continue. Otherwise, send \perp to \mathcal{AM} and halt.
5. Send $(witness, sid, x, w)$ to \mathcal{AM} .

VERIFY: On input $(verify, sid, \Lambda, Verify', x, w)$ from party \mathcal{V} the functionality does the following:

1. If $Verify' = Verify \wedge \Lambda \in A$:
 - (a) $t \leftarrow$ largest t such that $S[t]$ corresponds to Λ .
 - (b) If \mathcal{AM} not corrupted $\wedge x \notin S[t] \wedge Verify(pk, \Lambda, x, w) = 1$, send \perp to \mathcal{P} . Otherwise, continue.
 - (c) $b \leftarrow Verify(pk, \Lambda, w, x)$
Otherwise, set $b = Verify'(pk, \Lambda, w, x)$.
2. Send $(verified, sid, \Lambda, Verify', x, w, b)$ to \mathcal{V} .

ADD: On input (add, sid, x) from \mathcal{AM} the functionality does the following:

1. If sid does not encode the identity of \mathcal{AM} , or if $x \notin D_A$, ignore this command. Otherwise, continue.
2. If $x \in S[t]$ send \perp to \mathcal{AM} and halt. Otherwise, continue.
3. $t \leftarrow t + 1$, and $S[t] \leftarrow S[t - 1]$.
4. Run $(\Lambda_{\mathcal{X}'}, aux') \leftarrow Add((sk, pk), \Lambda_{\mathcal{X}}, aux, x)$.
5. Run $w \leftarrow WitCreate((sk, pk), \Lambda_{\mathcal{X}'}, aux', x)$.
6. If $Verify(pk, \Lambda_{\mathcal{X}'}, w, x) = 0$, send \perp to \mathcal{AM} and halt. Otherwise, continue.
7. Store aux' ; add x to $S[t]$ and $\Lambda_{\mathcal{X}'}$ to A .
8. Send $(added, sid, \Lambda_{\mathcal{X}'}, x)$ to \mathcal{AM} .

DELETE: On input $(delete, sid, x)$ from \mathcal{AM} the functionality does the following:

1. If sid does not encode the identity of \mathcal{AM} , or if $x \notin D_A$, ignore this command. Otherwise, continue.
2. If $x \notin S[t]$ send \perp to \mathcal{AM} and halt. Otherwise, continue.
3. $t \leftarrow t + 1$, and $S[t] \leftarrow S[t - 1]$.
4. Run $(\Lambda_{\mathcal{X}'}, aux') \leftarrow Delete((sk, pk), \Lambda_{\mathcal{X}}, aux, x)$.
5. Store aux' ; remove x from $S[t]$ and add $\Lambda_{\mathcal{X}'}$ to A .
6. Send $(deleted, sid, \Lambda_{\mathcal{X}'}, x)$ to \mathcal{AM} .

WITUPDATE: On input $(witupdate, sid, \Lambda_{\mathcal{X}_{old}}, \Lambda_{\mathcal{X}_{new}}, x, w_{old})$ from a party \mathcal{H} , the functionality does the following:

1. If $\Lambda_{\mathcal{X}_{old}} \notin A \vee \Lambda_{\mathcal{X}_{new}} \notin A$, send \perp to \mathcal{H} and halt. Otherwise continue.
2. Run $w_{new} \leftarrow WitUpdate((sk, pk), w_{old}, aux, x)$.
3. If $Verify(pk, \Lambda_{\mathcal{X}_{old}}, w_{old}, x) = 1 \wedge x \in S[t] \wedge Verify(pk, \Lambda_{\mathcal{X}_{new}}, w_{new}, x) = 0$, send \perp to \mathcal{V} and halt. Otherwise, continue.
4. Send $(updatedwit, sid, \Lambda_{\mathcal{X}_{old}}, \Lambda_{\mathcal{X}_{new}}, x, w_{old}, w_{new})$ to \mathcal{H} .

Functionality 1: Ideal Functionality \mathcal{F}_{Acc} for dynamic accumulators

4 Multi-Party Public-Key Accumulators

With the building blocks in place, we are now able to go into the details of our construction. We first present the formal notion of linearly secret-shared accumulators, their ideal functionality, and then present our constructions.

For the syntax of the MPC-based accumulator, which we dub *linear secret-shared accumulator*, we use the bracket notation $\langle s \rangle$ from Section 2.2 to denote a secret shared value. If we want to explicitly highlight the different shares, we write $\langle s \rangle = s_1 + \dots + s_n$, where the share s_i belongs to a party P_i . We base the definition on the framework of Derler et al. [34], where our algorithms behave in the same way, but instead of taking an optional secret trapdoor, the algorithms are given shares of the secret as input. Consequently, Gen outputs shares of the secret trapdoor instead of the secret key. The static version of the accumulator is defined as follows:

Definition 5 (Static Linear Secret-Shared Accumulator). *Let us assume that we have a linear secret sharing-scheme. A static linear secret-shared accumulator for $n \in \mathbb{N}$ parties P_1, \dots, P_n is a tuple of efficient algorithms $(\text{Gen}, \text{Eval}, \text{WitCreate}, \text{Verify})$ which are defined as follows:*

- $\text{Gen}(1^\kappa, t)$: *This algorithm takes a security parameter κ and a parameter t . If $t \neq \infty$, then t is an upper bound on the number of elements to be accumulated. It returns a key pair $(\text{sk}_\Lambda^i, \text{pk}_\Lambda)$ to each party P_i such that $\text{sk}_\Lambda = \text{sk}_\Lambda^1 + \dots + \text{sk}_\Lambda^n$, denoted by $\langle \text{sk}_\Lambda \rangle$. We assume that the accumulator public key pk_Λ implicitly defines the accumulation domain \mathcal{D}_Λ .*
- $\text{Eval}(\langle \text{sk}_\Lambda \rangle, \text{pk}_\Lambda, \mathcal{X})$: *This algorithm takes a secret-shared private key $\langle \text{sk}_\Lambda \rangle$ a public key pk_Λ and a set \mathcal{X} to be accumulated and returns an accumulator $\Lambda_\mathcal{X}$ together with some auxiliary information aux to every party P_i .*
- $\text{WitCreate}(\langle \text{sk}_\Lambda \rangle, \text{pk}_\Lambda, \Lambda_\mathcal{X}, \text{aux}, x)$: *This algorithm takes a secret-shared private key $\langle \text{sk}_\Lambda \rangle$ a public key pk_Λ , an accumulator $\Lambda_\mathcal{X}$, auxiliary information aux and a value x . It returns \perp , if $x \notin \mathcal{X}$, and a witness wit_x for x otherwise to every party P_i .*
- $\text{Verify}(\text{pk}_\Lambda, \Lambda_\mathcal{X}, \text{wit}_x, x)$: *This algorithm takes a public key pk_Λ , an accumulator $\Lambda_\mathcal{X}$, a witness wit_x and a value x . It returns 1 if wit_x is a witness for $x \in \mathcal{X}$ and 0 otherwise.*

In analogy to the non-interactive case, dynamic accumulators provide additional algorithms to add elements to the accumulator and remove elements from it, respectively, and update already existing witnesses accordingly. These algorithms are defined as follows:

Definition 6 (Dynamic Linear Secret-Shared Accumulator). *A dynamic linear secret-shared accumulator is a static linear secret-shared accumulator with an additional tuple of efficient algorithms $(\text{Add}, \text{Delete}, \text{WitUpdate})$ which are defined as follows:*

- $\text{Add}(\langle \text{sk}_\Lambda \rangle, \text{pk}_\Lambda, \Lambda_\mathcal{X}, \text{aux}, x)$: *This algorithm takes a secret-shared private key $\langle \text{sk}_\Lambda \rangle$ a public key pk_Λ , an accumulator $\Lambda_\mathcal{X}$, auxiliary information aux , as*

well as an element x to be added. If $x \in \mathcal{X}$, it returns \perp to every party P_i . Otherwise, it returns the updated accumulator $\Lambda_{\mathcal{X}'}$ with $\mathcal{X}' \leftarrow \mathcal{X} \cup \{x\}$ and updated auxiliary information \mathbf{aux}' to every party P_i .

Delete($(\langle \mathbf{sk}_\Lambda \rangle, \mathbf{pk}_\Lambda), \Lambda_{\mathcal{X}}, \mathbf{aux}, x$): This algorithm takes a secret-shared private key $\langle \mathbf{sk}_\Lambda \rangle$ a public key \mathbf{pk}_Λ , an accumulator $\Lambda_{\mathcal{X}}$, auxiliary information \mathbf{aux} , as well as an element x to be added. If $x \notin \mathcal{X}$, it returns \perp to every party P_i . Otherwise, it returns the updated accumulator $\Lambda_{\mathcal{X}'}$ with $\mathcal{X}' \leftarrow \mathcal{X} \setminus \{x\}$ and updated auxiliary information \mathbf{aux}' to every party P_i .

WitUpdate($(\langle \mathbf{sk}_\Lambda \rangle, \mathbf{pk}_\Lambda), \mathbf{wit}_{x_i}, \mathbf{aux}, x$): This algorithm takes a secret-shared private key $\langle \mathbf{sk}_\Lambda \rangle$ a public key \mathbf{pk}_Λ , a witness \mathbf{wit}_{x_i} to be updated, auxiliary information \mathbf{aux} and an element x which was added to/deleted from the accumulator, where \mathbf{aux} indicates addition or deletion. It returns an updated witness \mathbf{wit}'_{x_i} on success and \perp otherwise to every party P_i .

Correctness and collision-freeness naturally translate from the non-interactive accumulators to the linear secret-shared ones. The work of Baldimtsi et al. also introduced the property creation-correctness. Informally speaking, creation-correctness allows the generation of witnesses during addition. In the above definitions, we see that adding an element to the accumulator and creating a witness are two separate algorithms. Therefore, the notion of creation-correctness does not immediately apply to our accumulators.

For our case, the ideal functionality for linear secret-shared accumulators, dubbed $\mathcal{F}_{\text{MPC-Acc}}$ is more interesting. $\mathcal{F}_{\text{MPC-Acc}}$ is very similar to \mathcal{F}_{Acc} and is depicted in Functionality 2. The only difference in describing the ideal functionality for accumulators in the MPC setting arises from the fact that we now have not only one accumulator manager but n , denoted by $\mathcal{AM}_1, \dots, \mathcal{AM}_n$. More concretely, whenever a sub-functionality of $\mathcal{F}_{\text{MPC-Acc}}$ - that makes use of the secret key - gets a request from a manager identity \mathcal{AM}_i , it now also gets a participation message from the other managers identities \mathcal{AM}_j for $j \neq i$. Furthermore, the accumulator managers take the role of the witness holder. The party \mathcal{V} , however, stays unchanged.

GEN: On input (gen, sid_i) from all parties \mathcal{AM}_i , the functionality does the following:

1. If this is not the first *gen* command, or if for any $i \in [n]$, sid_i does not encode the identity of \mathcal{AM}_i , ignore this command. Otherwise, continue.
2. $t \leftarrow 0$.
3. Initialize an empty list A (keeps track of all accumulator states).
4. Initialize map S , and set $S[0] \leftarrow \emptyset$ (maps operation counters to current accumulated sets).
5. Send $(gen, sid_i)_{i \in [n]}$ to \mathcal{SLM} .
6. Get $(algorithms, (sid_1, \dots, sid_n), (\text{Gen, Eval, WitCreate, Verify, Add, Delete, WitUpdate}))$ from \mathcal{SLM} . Their expected input output behavior is described in Definition 6. All of them should be polynomial-time and Verify should be deterministic.
7. Run $(sk, pk) \leftarrow \text{Gen}(1^\lambda)$. Store sk, pk ; add $\Lambda_\emptyset \leftarrow \emptyset$ to A .
8. Send $(algorithms, sid_i, (\text{Gen, Eval, WitCreate, Verify, Add, Delete, WitUpdate}))$ to \mathcal{AM}_i , for all $i = 1, \dots, n$.

EVAL: On input $(eval, sid_k, \mathcal{X})$ from \mathcal{AM}_k and $(eval, sid_j, ?)$ from all other parties \mathcal{AM}_j , for $j \neq k$, the functionality does the following:

1. If for any $i \in [n]$, sid_i does not encode the identity of \mathcal{AM}_i , or if $\mathcal{X} \not\subseteq D_A$, ignore this command. Otherwise, continue.
2. $t \leftarrow t + 1$, and $S[t] \leftarrow \mathcal{X}$.
3. Run $(\Lambda_{\mathcal{X}}, aux) \leftarrow \text{Eval}((sk, pk), \mathcal{X})$. Store aux ; add $\Lambda_{\mathcal{X}}$ to A .
4. Send $(eval, sid_i, \Lambda_{\mathcal{X}}, \mathcal{X})$ to \mathcal{AM}_i , for all $i = 1, \dots, n$.

WITCREATE: On input $(witcreate, sid_k, x)$ from \mathcal{AM}_k and $(witcreate, sid_j, ?)$ from all parties \mathcal{AM}_j , for $j \neq k$, the functionality does the following:

1. If for any $i \in [n]$, sid_i does not encode the identity of \mathcal{AM}_i , or if $x \notin D_A$, ignore this command. Otherwise, continue.
2. Run $w \leftarrow \text{WitCreate}((sk, pk), \Lambda_{\mathcal{X}}, aux, x)$.
3. If $x \notin S[t]$, send \perp to all \mathcal{AM}_i and halt. Otherwise, continue.
4. If $\text{Verify}(pk, \Lambda_{\mathcal{X}}, w, x) = 1$ continue. Otherwise, send \perp to all \mathcal{AM}_i and halt.
5. Send $(witness, sid_i, x, w)$ to \mathcal{AM}_i , for all $i = 1, \dots, n$.

VERIFY: On input $(verify, sid, \Lambda, \text{Verify}', x, w)$ from party \mathcal{V} the functionality does the following:

1. If $\text{Verify}' = \text{Verify} \wedge \Lambda \in A$:
 - (a) $t \leftarrow$ largest t such that $S[t]$ corresponds to Λ .
 - (b) If at least one \mathcal{AM}_i is not corrupted $\wedge x \notin S[t] \wedge \text{Verify}(pk, \Lambda, x, w) = 1$, send \perp to \mathcal{V} . Otherwise, continue.
 - (c) $b \leftarrow \text{Verify}(pk, \Lambda, w, x)$
- Otherwise, set $b = \text{Verify}'(pk, \Lambda, w, x)$.
2. Send $(verified, sid, \Lambda, \text{Verify}', x, w, b)$ to \mathcal{V} .

ADD: On input (add, sid_k, x) from \mathcal{AM}_k and $(add, sid_j, ?)$ from all parties \mathcal{AM}_j , for $j \neq k$, the functionality does the following:

1. If for any $i \in [n]$, sid_i does not encode the identity of \mathcal{AM}_i , or if $x \notin D_A$, ignore this command. Otherwise, continue.
2. If $x \in S[t]$ send \perp to all \mathcal{AM}_i and halt. Otherwise, continue.
3. $t \leftarrow t + 1$, and $S[t] \leftarrow S[t - 1]$.
4. Run $(\Lambda_{\mathcal{X}'}, aux') \leftarrow \text{Add}((sk, pk), \Lambda_{\mathcal{X}}, aux, x)$. Run $w \leftarrow \text{WitCreate}((sk, pk), \Lambda_{\mathcal{X}'}, aux', x)$.
5. If $\text{Verify}(pk, \Lambda_{\mathcal{X}'}, w, x) = 0$, send \perp to all \mathcal{AM}_i and halt. Otherwise, continue.
6. Store aux' ; add x to $S[t]$ and $\Lambda_{\mathcal{X}'}$ to A .
7. Send $(added, sid_i, \Lambda_{\mathcal{X}'}, x)$ to \mathcal{AM}_i , for all $i = 1, \dots, n$.

DELETE: On input $(delete, sid_k, x)$ from \mathcal{AM}_k and $(delete, sid_j, ?)$ from all parties \mathcal{AM}_j , for $j \neq k$, the functionality does the following:

1. If for any $i \in [n]$, sid_i does not encode the identity of \mathcal{AM}_i , or if $x \notin D_A$, ignore this command. Otherwise, continue.
2. If $x \notin S[t]$ send \perp to all \mathcal{AM}_i and halt. Otherwise, continue.
3. $t \leftarrow t + 1$, and $S[t] \leftarrow S[t - 1]$.
4. Run $(\Lambda_{\mathcal{X}'}, aux') \leftarrow \text{Delete}((sk, pk), \Lambda_{\mathcal{X}}, aux, x)$.
5. Store aux' ; remove x from $S[t]$ and add $\Lambda_{\mathcal{X}'}$ to A .
6. Send $(deleted, sid_i, \Lambda_{\mathcal{X}'}, x)$ to \mathcal{AM}_i , for all $i = 1, \dots, n$.

WITUPDATE: On input $(witupdate, sid_k, \Lambda_{\mathcal{X}_{old}}, \Lambda_{\mathcal{X}_{new}}, x, w_{old})$ from \mathcal{AM}_k and $(witupdate, sid_j, ?, ?, ?, ?)$ from all parties \mathcal{AM}_j , for $j \neq k$, the functionality does the following:

1. If for any $i \in [n]$, sid_i does not encode the identity of \mathcal{AM}_i , or if $x \notin D_A$, ignore this command. Otherwise, continue.
2. If $\Lambda_{\mathcal{X}_{old}} \notin A \vee \Lambda_{\mathcal{X}_{new}} \notin A$, send \perp to all \mathcal{AM}_i and halt. Otherwise continue.
3. Run $w_{new} \leftarrow \text{WitUpdate}((sk, pk), w_{old}, aux, x)$.
4. If $\text{Verify}(pk, \Lambda_{\mathcal{X}_{old}}, w_{old}, x) = 1 \wedge x \in S[t] \wedge \text{Verify}(pk, \Lambda_{\mathcal{X}_{new}}, w_{new}, x) = 0$, send \perp to \mathcal{V} and halt. Otherwise, continue.
5. Send $(updatedwit, sid_i, \Lambda_{\mathcal{X}_{old}}, \Lambda_{\mathcal{X}_{new}}, x, w_{old}, w_{new})$ to \mathcal{AM}_i , for all $i = 1, \dots, n$.

Functionality 2: Ideal Functionality $\mathcal{F}_{\text{Acc-MPC}}$

4.1 Dynamic Linear Secret-Shared Accumulator from the t -SDH Assumption

Let's start with the generation of the public parameters, **Gen**. For this algorithm, we can rely on already established methods to produce ECDSA key pairs and to calculate with secret exponents, respectively. These methods can directly be applied to the accumulators. If we take the t -SDH accumulator as an example, then the first step is to sample the secret scalar $s \in \mathbb{Z}_p$. Intuitively, each party samples its own share s_i of s , which are then combined to $s = \sum s_i$. The next step is the calculation of the basis elements g^{s^j} for $j = 1, \dots, t$. All of these elements can be computed using $\text{Exp}_{\mathbb{G}}$ and the secret-shared s , respectively its powers. Note that producing the powers is necessary to provide public parameters that are useful even to parties without knowledge of s .

For the accumulator evaluation, **Eval**, the idea is that, first, the parties sample their shares of r . Then, they cooperate to compute shares of $r \cdot f(s)$ using their respective shares of r and s . Again, the so-obtained exponent and $\text{Exp}_{\mathbb{G}}$ produce the final result.

For witness creation, **WitCreate**, it gets more interesting. Of course, one could simply run **Eval** again with one element removed from the set. In this case, we can do better, though. The difference between the accumulator and a witness is that in the latter, one factor of the polynomial is canceled. Since s is available, it is thus possible to cancel this factor without recomputing the polynomial from the start. Indeed, to compute the witness for element x , we can compute $(s+x)^{-1}$ and then apply that inverse using $\text{Exp}_{\mathbb{G}}$ to the accumulator to get the witness. Note though, that before the parties perform this step, they need to check if x is actually contained in \mathcal{X} . Otherwise, they would produce a membership witness for a non-member. In that case, the verification would check whether $f(s)(s+x)^{-1}(s+x)$ matches $f(s)$, which of course also holds even if $s+x$ is not a factor of $f(s)$. In contrast, when performing **Eval** with only the publicly available information, this issue does not occur since then the witness will not verify. Similarly, **Add** and **Delete** can be implemented in a similar manner. When adding an element to the accumulator, the polynomial is extended by one factor. Removal of an element requires that one factor is canceled. Both operations can be performed by first computing the factor using the shares of s and then running $\text{Exp}_{\mathbb{G}}$.

Now, we present the MPC version of the t -SDH accumulator in Scheme 2 following the intuition outlined above. Note that we let **Gen** choose the bilinear group BG , but this group can already be fixed a priori.

Theorem 3. *Scheme 2 UC emulates $\mathcal{F}_{\text{Acc-MPC}}$ in the $\mathcal{F}_{\text{SPDZ+}}$ -hybrid model.*

Proof. At this point, we make use of the UC model. Informally speaking, accumulators are UC secure, and SPDZ and the derived operations UC emulate $\mathcal{F}_{\text{SPDZ+}}$. Therefore, according to the universal composition theorem, the use of the SPDZ protocol in the accumulator Scheme 2 can be done without losing UC security. For a better understanding, we elaborate this argumentation. We begin by showing the desired standard accumulator properties for Scheme 2.

<p>Gen($1^\kappa, t$): $\text{BG} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2) \leftarrow \text{BGen}(\kappa)$. Compute $\langle \text{sk}_A \rangle \leftarrow \text{sRand}(\mathbb{Z}_p^*)$. Compute $h \leftarrow \text{Open}(g_2^{\langle \text{sk}_A \rangle})$. Return $\text{pk}_A \leftarrow (\text{BG}, h)$.</p>
<p>Eval($(\langle \text{sk}_A \rangle, \text{pk}_A), \mathcal{X}$): Parse pk_A as (BG, h) and \mathcal{X} as subset of \mathbb{Z}_p^*. Choose $\langle r \rangle \leftarrow \text{sRand}(\mathbb{Z}_p^*)$. Compute $\langle q \rangle \leftarrow \prod_{x \in \mathcal{X}} (x + \langle \text{sk}_A \rangle) \in \mathbb{Z}_p^*$ and $\langle t \rangle \leftarrow \langle q \rangle \cdot \langle r \rangle$. The algorithm returns $\Lambda_{\mathcal{X}} \leftarrow \text{Open}(g_1^{\langle t \rangle})$ and $\text{aux} \leftarrow (\text{add} \leftarrow 0, \mathcal{X})$.</p>
<p>WitCreate($(\langle \text{sk}_A \rangle, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x$): Returns \perp if $x \notin \mathcal{X}$. Otherwise, $\langle z \rangle \leftarrow \langle (x + \langle \text{sk}_A \rangle)^{-1} \rangle$. Return $\text{wit}_x \leftarrow \text{Open}(\Lambda_{\mathcal{X}}^{\langle z \rangle})$.</p>
<p>Verify($\text{pk}_A, \Lambda_{\mathcal{X}}, \text{wit}_x, x$): Parse pk_A as (BG, h). If $e(\Lambda_{\mathcal{X}}, g_2) = e(\text{wit}_x, g_2^x \cdot h)$ holds, return 1, otherwise return 0.</p>
<p>Add($(\langle \text{sk}_A \rangle, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x$): Returns \perp if $x \in \mathcal{X}$. Otherwise set $\mathcal{X}' \leftarrow \mathcal{X} \cup \{x\}$. Return $\Lambda_{\mathcal{X}'} \leftarrow \Lambda_{\mathcal{X}}^x \cdot \text{Open}(\Lambda_{\mathcal{X}}^{\langle \text{sk}_A \rangle})$ and $\text{aux} \leftarrow (\text{add} \leftarrow 1, \mathcal{X}')$.</p>
<p>Delete($(\langle \text{sk}_A \rangle, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x$): If $x \notin \mathcal{X}$, return \perp. Otherwise set $\mathcal{X}' \leftarrow \mathcal{X} \setminus \{x\}$, and compute $\langle y \rangle \leftarrow \langle (x + \langle \text{sk}_A \rangle)^{-1} \rangle$. Return $\Lambda_{\mathcal{X}'} \leftarrow \text{Open}(\Lambda_{\mathcal{X}}^{\langle y \rangle})$ and $\text{aux} \leftarrow (\text{add} \leftarrow -1, \mathcal{X}')$.</p>
<p>WitUpdate($(\langle \text{sk}_A \rangle, \text{pk}_A), \text{wit}_{x_i}, \text{aux}, x$): Parse aux as $(\text{add}, \mathcal{X})$. Return \perp if $\text{add} = 0$ or $x_i \notin \mathcal{X}$. In case $\text{add} = 1$, return $\text{wit}_{x_i} \leftarrow \text{wit}_{x_i}^x \cdot \text{Open}(\text{wit}_{x_i}^{\langle \text{sk}_A \rangle})$ and $\text{aux} \leftarrow (\text{add} \leftarrow 0, \mathcal{X})$. If instead $\text{add} = -1$, it compute $\langle y \rangle \leftarrow \langle (x + \langle \text{sk}_A \rangle)^{-1} \rangle$. Return $\text{wit}_{x_i} \leftarrow \text{Open}(\text{wit}_{x_i}^{\langle y \rangle})$ and $\text{aux} \leftarrow (\text{add} \leftarrow 0, \mathcal{X})$.</p>

Scheme 2: MPC- t -SDH: Dynamic linear secret-shared accumulator from t -SDH for $n \geq 2$ parties.

The proof of the correctness follows directly from the correctness proof from Scheme 1 for the case where the secret key is known. Collision-freeness is also derived from the non-interactive t -SDH accumulator. (It is true that now each party has a share of the trapdoor, but without all the other shares no party can create a valid witness.) Since **Verify** is obviously deterministic, Scheme 2 fulfills all necessary assumption of Theorem 2.

After applying Theorem 2, we get a simulator \mathcal{SIM}_{Acc} interacting with the ideal functionality \mathcal{F}_{Acc} . Since we now also have to simulate the non-interactive sub-protocols, we have to extend \mathcal{SIM}_{Acc} . We construct $\mathcal{SIM}_{Acc\text{-MPC}}$ by building upon \mathcal{SIM}_{Acc} and in addition internally simulate \mathcal{F}_{SPDZ+} .

As described in Section 2.2, the MPC protocols used in the above algorithms are all secure in the UC model. Since we do not open any secret-shared values besides uniformly random elements and the output or values that can be immediately derived from the output, the algorithms are secure due to the universal composition theorem.

Note, that the algorithm for **WitUpdate** is unlikely to be faster than its non-MPC version from Scheme 1. Indeed, the non-MPC version requires only exponentiations in \mathbb{G}_1 and a multiplication without the knowledge of the secret trapdoor. We provide the version using the trapdoor for completeness but will use the non-MPC version of the algorithm in the remainder.

Remark 2. In **Gen** of Scheme 2 we explicitly do not compute $h_i \leftarrow g_1^{s_i}$. Hence, using **Eval** without access to s is not possible. But, on the positive side, the

public key is significantly smaller and so is the runtime of the `Gen` algorithm. If, however, these values are needed to support a non-secret-shared `Eval`, one can modify `Gen` to also compute the following values: $\langle t_1 \rangle \leftarrow \langle s \rangle$, $\langle t_i \rangle \leftarrow \langle t_{i-1} \rangle \cdot \langle s \rangle$, and $h_i \leftarrow \text{Open}(g_1^{\langle t_i \rangle})$ for $i = 1, \dots, t$.

5 Implementation and Performance Evaluation

We implemented the proposed dynamic linear secret-shared accumulator from t -SDH and evaluated it against small to large sets⁸. Our primary implementation is based on MP-SPDZ [77]; however, to demonstrate the usability of our accumulator, we additionally build an implementation in the malicious security setting based on the FRESCO [3] framework. We discuss the benchmarks for the MP-SPDZ implementation in this section; for a discussion of the FRESCO benchmarks we refer the reader to Appendix E. MP-SPDZ implements the SPDZ protocol and various extensions [33,52,53,28]. For pairing and elliptic curve group operations, we rely on the Relic library [5] and integrate $\text{Exp}_{\mathbb{G}}$, $\text{Output}_{\mathbb{G}}$, and the corresponding operations to update the MAC described in [76] into MP-SPDZ. We use the pairing friendly BLS12-381 curve [12], which provides around 120 bit of security following recent estimates [10,62].

For completeness, we also implemented the unmodified versions of the t -SDH accumulator from Scheme 1, and also a Merkle-tree accumulator (cf. Appendix C) using SHA256. This enables us to compare the performance in cases where the secret trapdoors are available in the MPC case and when they are not. In Table 1, we present the numbers for various sizes of the accumulated set.

Table 1. Performance of the accumulator algorithms without access to the secret trapdoors. Time in milliseconds averaged over 100 executions.

Accumulator Operation	t -SDH		Merkle-tree	
	$ \mathcal{X} = 2^{10}$	2^{14}	$ \mathcal{X} = 2^{10}$	2^{14}
Gen	649	9 062	-	-
Eval	1 117	116 031	1.12	15.53
WitCreate	1 116	115 870	0.05 ^a	0.83 ^a
Add	1 116	115 575	1.12	15.53
WitUpdate _{Add}	0.6	0.6	0.05 ^a	0.83 ^a
Delete	1 120	116 154	1.12	15.53
WitUpdate _{Delete}	0.7	0.7	0.05 ^a	0.83 ^a

^a Here we assume that the full Merkle-tree is known as auxiliary data. If this is not the case, the tree has to be rebuilt, which adds the `Eval`-time.

⁸ The source code is available at <https://github.com/IAIK/MPC-Accumulator>.

Table 2. Number of Beaver-triples and shared random values for each MPC- t -SDH operation.

Operation	Beaver-triples	Random values
Gen	0	1
Eval	$ \mathcal{X} $	1
WitCreate	1	1
Add	0	0
WitUpdate _{Add}	0	0
Delete	1	1
WitUpdate _{Delete}	1	1

The evaluation of the MPC protocol was performed on a cluster with a Xeon E5-4669v4 CPU, where each party was assigned only 1 core. The hosts were connected via a 1 Gbit/s LAN network, and an average round-trip time of < 1 ms. For the WAN setting, a network with a round-trip time of 100 ms and a bandwidth of 100 Mbit/s was simulated. For the MPC protocol, we provide benchmarks for both preprocessing and online phases. Note that the cost of the preprocessing phase is determined by the number of shared multiplications, whereas the performance of the online phase is given by the multiplicative depth of the circuit and the number of openings.

5.1 Evaluation of MPC- t -SDH

In the offline phase of the SPDZ protocol, the required Beaver-triples [15] for shared multiplication and the pre-shared random values are generated. A shared inverse operation requires one multiplication and one shared random value. In Table 2, we list the number of triples required for each operation for the MPC- t -SDH accumulator. Each of the operations requires a constant number of multiplications and inverse operations and, therefore, a constant number of Beaver-triples and shared random elements – with the exception of the Eval operation. In Eval, the number of required Beaver-triples is determined by $|\mathcal{X}|$. As discussed in Remark 2, Gen is not producing the public parameters h_i , which enable accumulation of elements without the use of MPC. If this feature is desired, the time and communication of Eval for the respective set sizes should be added to the time and communication of Gen to obtain an estimate of its performance.

Table 3 compares the offline performance of the MPC- t -SDH accumulator in different settings. We give both timings for the accumulation of $|\mathcal{X}|$ elements in Eval and the necessary pre-computation for a single inversion, which is used in several other operations (e.g. WitCreate). Additionally we also give the time for pre-computing a single random element, which is required to generate the authenticated share of the secret-key in Eval. Further note that batching the generation of many triples together like for the Eval phase is more efficient in practice than producing a single triple and as these triples are not dependant on

Table 3. Offline phase performance of different steps of the MPC- t -SDH accumulator with access to the secret trapdoor implemented in MP-SPDZ. Time in milliseconds.

Operation	$ \mathcal{X} $	LAN setting				WAN setting			
		$n = 2$	3	4	5	$n = 2$	3	4	5
BaseOTs	2^{10}	0.03	0.08	0.14	0.23	0.14	0.31	0.56	0.84
	2^{14}	0.03	0.08	0.14	0.23	0.14	0.31	0.56	0.84
Semi-Honest									
Inverse	2^{10}	0.78	1.72	3.06	4.03	209.9	227.5	322.8	331.0
	2^{14}	0.78	1.72	3.06	4.03	209.9	227.5	322.8	331.0
Gen	2^{10}	0.44	1.21	1.76	3.01	207.7	223.6	325.9	332.0
	2^{14}	0.44	1.21	1.76	3.01	207.7	223.6	325.9	332.0
Eval	2^{10}	189	397	706	959	4 695	8 215	13 680	25 725
	2^{14}	4000	8 308	14 380	17 928	55 542	109 720	214 585	356 330
Malicious									
Inverse	2^{10}	4.34	7.93	11.5	15.3	840.5	1 262	1 538	1 914
	2^{14}	4.34	7.93	11.5	15.3	840.5	1 262	1 538	1 914
Gen	2^{10}	2.56	4.23	6.80	9.32	841.3	1 235	1 540	1 856
	2^{14}	2.56	4.23	6.80	9.32	841.3	1 235	1 540	1 856
Eval	2^{10}	1 601	2 849	4 345	6 227	25 737	45 254	87 328	141 181
	2^{14}	31 099	62 978	89 132	145 574	412 747	682 033	1 364 660	2 236 860

the input, all parties can continuously generate triples in the background to fill a triple-buffer for use in the online phase.

In Table 4, we present the online performance of our MPC- t -SDH accumulator for different set sizes, parties, security settings, and network settings. It can clearly be seen, that – except for the **Eval** operation – the runtime of each operation is independent of the set size. In other words, after an initial accumulation of a given set, every other operation has constant time. In comparison, the runtime of the non-MPC accumulators without access to the secret trapdoor, as depicted in Table 1, depends on the size of the accumulated set. Our MPC-accumulator outperforms the non-MPC t -SDH accumulators the larger the accumulated set gets. In the LAN setting MPC- t -SDH’s **Eval** is faster than the non-MPC version for all benchmarked players, even in the WAN settings it outperforms the non-MPC version in the two player case. For 2^{14} elements, it is even faster for all benchmarked players in all settings, including the WAN setting. In any case, the witnesses have constant size contrary to the $\log_2(|\mathcal{X}|)$ sized witnesses of the Merkle-tree accumulator.

The numbers for the evaluation of the online phase in the WAN setting are also presented in Table 4. The overhead that can be observed compared to the LAN setting is influenced by the communication complexity. Since our implementation implements all multiplications in **Eval** in a depth-optimized tree-like fashion, the overhead from switching to a WAN setting is not too severe.

Table 4. Online phase performance of the MPC- t -SDH accumulator with access to the secret trapdoor implemented in MP-SPDZ, for both the LAN and WAN settings with n parties. Time in milliseconds averaged over 50 executions.

Operation	$ \mathcal{X} $	Semi-Honest								Malicious											
		LAN setting				WAN setting				LAN setting				WAN setting							
		$n=$	2	3	4	5	$n=$	2	3	4	5	$n=$	2	3	4	5	$n=$	2	3	4	5
Gen	2^{10}		4	4	7	19		53	110	170	219		11	13	25	37		169	278	395	505
	2^{14}		4	4	9	20		56	111	172	220		11	13	28	48		179	280	396	506
Eval	2^{10}		3	13	58	231		635	1277	1916	2558		10	17	50	131		966	1327	1669	1995
	2^{14}		26	47	117	315		949	1948	3166	4571		89	94	174	225		1297	1979	2830	3872
WitCreate	2^{10}		2	2	32	39		168	320	482	645		5	10	35	75		372	606	823	1050
	2^{14}		2	2	28	51		168	320	473	638		5	6	28	80		365	606	835	1052
Add	2^{10}		2	2	8	17		47	107	166	213		5	5	17	31		170	273	388	499
	2^{14}		2	2	5	14		50	108	170	214		5	5	17	42		173	271	383	491
WitUpdate _{Add}	2^{10}		2	2	5	30		60	108	154	214		5	7	12	34		159	276	390	495
	2^{14}		2	2	3	20		60	107	152	217		5	6	10	54		154	275	390	500
Delete	2^{10}		2	2	21	58		156	319	488	639		5	10	47	78		379	598	818	1034
	2^{14}		2	2	23	55		158	318	489	642		5	6	38	87		385	603	822	1033
WitUpdate _{Delete}	2^{10}		2	2	52	47		165	320	475	643		5	10	26	100		374	604	828	1044
	2^{14}		2	4	43	57		162	323	475	639		5	10	35	74		365	599	827	1048

On the first look, one can observe an irregularity in our benchmarks. More specifically, notice that for four or more parties, the maliciously secure evaluation of the Eval online phase is consistently faster than the semi-honest evaluation of the same phase. However, this is a direct consequence of a difference in how MP-SPDZ handles the communication in those security models, where communication is handled in a non-synchronized send-to-all approach in the malicious setting and a synchronized broadcast approach in the semi-honest setting. The synchronization in the latter case scales worse for more parties and, therefore, introduces some additional delays.

Finally, Table 5 depicts the size of the communication between the parties for both offline and online phases. The communication of Eval has to account for a number of multiplications dependant on \mathcal{X} and therefore scales linearly with its size. As we already observed for the runtime of MPC- t -SDH, also the communication of WitCreate, Add, Delete and WitUpdate is independent of the size of the accumulated set, and additionally less than 200 kB for all algorithms. Combined with the analysis of the runtime, we conclude that the performance of the operations that might be performed multiple times per accumulator is very efficient in both runtime and communication. When compared to the performance of the non-MPC accumulators in Table 1, we see that the performance of operations that benefit from access to the secret trapdoor are multiple orders of magnitude faster in the MPC accumulators and, in the LAN setting, even come close to the performance of the standard Merkle-tree accumulator, for both the semi-honest and malicious variant.

6 Applications

In this section, we discuss the applications of MPC-based accumulators.

Table 5. Communication complexity of the MPC- t -SDH accumulator with access to the secret trapdoor implemented in MP-SPDZ, per party. Communication in kB.

Operation	$ \mathcal{X} $	Semi-Honest		Malicious	
		Offline ^a	Online	Offline ^a	Online
Gen	2^{10}	20	0.10	86	0.24
	2^{14}	20	0.10	86	0.24
Eval	2^{10}	12 571	66	79 549	66
	2^{14}	200 823	1 049	1 271 484	1 049
WitCreate	2^{10}	33	0.15	164	0.37
	2^{14}	33	0.15	164	0.37
Add	2^{10}	4	0.05	4	0.14
	2^{14}	4	0.05	4	0.14
WitUpdate _{Add}	2^{10}	4	0.05	4	0.14
	2^{14}	4	0.05	4	0.14
Delete	2^{10}	33	0.15	164	0.37
	2^{14}	33	0.15	164	0.37
WitUpdate _{Delete}	2^{10}	33	0.15	164	0.37
	2^{14}	33	0.15	164	0.37

^a Includes BaseOTs for a new connection

6.1 Revocation of Credentials in Distributed Credential Systems

As first application of MPC-based accumulators, we focus on distributed credential systems [42], and in particular, on the implementation in Sovrin [54]. In general, anonymous credentials provide a mechanism for making identity assertions while maintaining privacy, yet, in classical, non-distributed systems require a trusted credential issuer. This central issuer, however, is both a single point of failure and a target for compromise and can make it challenging to deploy such a system. In a distributed system credential system, on the other hand, this trusted credential issuer is eliminated, e.g. by using distributed ledgers.

We shortly recall how Sovrin implements revocation. When issuing a credential, every user gets a unique revocation identifier i_R . All valid revocation IDs are accumulated using the t -SDH accumulator and the accumulator gets published. Additionally, the users also obtains a witness certifying membership of its i_R in the accumulator. Whenever a user shows their credential, they have to prove that they know this witness for their i_R with respect to the published accumulator. When a new user joins, the accumulator has to be updated. Consequently, all the witnesses have to be updated as well, as otherwise they would no longer be able to provide a valid proof. Similar, in the case that a user is revoked and thus removed from the accumulator, all other users have to update their witnesses accordingly. Also, the verifiers also always have to check for updated accumulators.

Now, recall the that t -SDH accumulator supports all the required operations without needing access to the trapdoor. Hence, all operations can be performed and, especially, the users can update their witnesses on their own as long as the corresponding i_{RS} are published on the ledger. While functionality-wise all operations are supported, performance-wise a large number of users becomes an issue. With potentially millions to billions of users, adding and deleting members from the accumulator becomes increasingly expensive (cf. Table 1). Hence, at a certain size, having access to the trapdoor would be beneficial. But, on the other side, generating membership witnesses for non-members would become possible in that case.

The latter is also an issue during the setup of the system. Trusting one third party to generate the public parameters of the accumulator might be undesired in a distributed system as in this case. The special structure of the Sovrin ecosystem with their semi-trusted foundation members, however, naturally fits to our multi-party accumulator. First, the foundation members can setup the public parameters in a distributed manner. Secondly, as all of them have shares of the trapdoor, they can also run the updates of the accumulator using the MPC- t -SDH-accumulator. The change to this accumulator is completely transparent to the clients and verifiers and no changes are required there.

6.2 Certificate-Transparency Logs and Privacy-preserving Variants

We finally look at the application of accumulators in the CT ecosystem. Certificate Authorities request the inclusion of certificates in the log whenever they sign a new certificate, and once the certificate was included in the log, auditors can check the consistency of this log. Additionally, TLS clients also verify whether all certificates that they obtain were actually logged, thereby ensuring that log servers do not hand out promises of certificate inclusion without following through. Technically, the CT log is realized as a Merkle-tree accumulator containing all certificates. As certificates need to be added continuously, it is made dynamic by simply recalculating the root hash and all the proofs. While dynamic accumulators would perfectly fit the use-case from a functionality point of view, their real-world performance without secret trapdoors is not good enough – recalculating hash trees is just more efficient. Knowledge of the secret trapdoors would however be catastrophic for this application, as the guarantees of the whole system break down: log servers could produce witnesses for any certificate they get queried on, even if it was never submitted to the log servers for inclusion.

In the CT ecosystem, the clients need to contact the log servers for the inclusion proof, and therefore verifying certificates has negative privacy implications, as querying the inclusion status of a certificate reveals the browsing behavior of the client to the log server. Based on previous work by Lueks and Goldberg [61], Kales et al. [50] proposed to rethink retrieval of the inclusion proofs by employing multi-server private information retrieval (PIR) to query the proofs. Specifically, they build their system using a 2-server PIR from distributed point functions [46]. To further improve performance, the accumulator

is split into sub-accumulator based on, e.g., time periods. The sub-accumulators for a given time period are then accumulated in a top-level accumulator. Consequently, the witnesses with respect to the sub-accumulator stay constant and can be embedded in the server’s certificate and only the membership-proofs of the sub-accumulators need to be updated when new certificates are added to the log. Only these top-level proofs have to be queried using PIR, thus greatly improving the overall performance, as smaller databases are more efficient to query.

However, one drawback of this solution is the increase in certificate size if one were to include this static membership witness for the sub-accumulator in the certificate itself. Kales et al. [50] propose to build sub-accumulators per hour, which would result in sub-accumulators that hold about 2^{16} certificates. A Merkle-tree membership proof for these sub-accumulators is 512 bytes in size when using SHA-256. In contrast, a membership proof for the t-SDH accumulator is only 48 bytes in size (using point compression for the 381-bit BLS curve we use in our implementation). A typical DER-encoded X509 certificate using RSA-2048 as used in TLS is about 1-2 KB in size, meaning inclusion of the Merkle-tree sub-accumulator membership proof would increase the certificate size by 25 – 50%, whereas the t-SDH sub-accumulator membership proof only increases the size by 2.5 – 5%.

We can now leverage the fact that their solution already requires two non-colluding servers for the multi-server PIR. In their original solution, these servers hold copies of the Merkle-tree accumulator and answer private membership queries for the top-level accumulator. We suggest that switching the accumulators from Merkle-trees to our MPC-t-SDH accumulator would give the benefit of small, constant size membership proofs, while still being performant enough to accumulate and produce witnesses for all elements of a sub-accumulator in one hour.

7 Conclusion

In this work, we introduced dynamic linear secret-shared accumulators which remove the need of a trusted third party for public-key accumulators. By replacing the trusted party by a distributed setup algorithm, we achieved even more: since shares of the secret trapdoor are now available, the otherwise expensive algorithms can also be implemented as MPC protocol making use of the trapdoor. Thereby we obtained – especially in the bilinear groups setting – an efficient accumulator even for large sizes of accumulated sets.

Since our constructions are generic in a sense, improvements in the underlying MPC protocols will directly translate to our accumulators. Also, more efficient implementations of those protocols in the respective frameworks will have the same impact.

Acknowledgments

This work was supported by EU’s Horizon 2020 project Safe-DEED, grant agreement n°825225, and KRAKEN, grant agreement n°871473, and EU’s Horizon 2020 ECSEL Joint Undertaking project SECREDAS, grant agreement n°783119, and by the ”DDAI” COMET Module within the COMET – Competence Centers for Excellent Technologies Programme, funded by the Austrian Federal Ministry for Transport, Innovation and Technology (bmvit), the Austrian Federal Ministry for Digital and Economic Affairs (bmdw), the Austrian Research Promotion Agency (FFG), the province of Styria (SFG) and partners from industry and academia. The COMET Programme is managed by FFG.

References

1. Abdolmaleki, B., Bagheri, K., Lipmaa, H., Zajac, M.: A subversion-resistant SNARK. In: ASIACRYPT (3). LNCS, vol. 10626, pp. 3–33. Springer (2017)
2. Abdolmaleki, B., Ramacher, S., Slamanig, D.: Lift-and-shift: Obtaining simulation extractable subversion and updatable snarks generically. Cryptology ePrint Archive, Report 2020/062 (2020), <https://eprint.iacr.org/2020/062>, to appear at CCS’20
3. Alexandra Institute: FRESCO - a FRamework for Efficient Secure COmputation (2019), <https://github.com/aicis/fresco>
4. Aly, A., Keller, M., Rotaru, D., Scholl, P., Smart, N.P., Wood, T.: SCALE-MAMBA (2019), <https://homes.esat.kuleuven.be/~nsmart/SCALE/>
5. Aranha, D.F., et al.: RELIC is an Efficient Library for Cryptography. <https://github.com/relic-toolkit/relic>
6. Au, M.H., Tsang, P.P., Susilo, W., Mu, Y.: Dynamic universal accumulators for DDH groups and their application to attribute-based anonymous credential systems. In: CT-RSA. LNCS, vol. 5473, pp. 295–308. Springer (2009)
7. Baldimtsi, F., Camenisch, J., Dubovitskaya, M., Lysyanskaya, A., Reyzin, L., Samelin, K., Yakoubov, S.: Accumulators with applications to anonymity-preserving revocation. In: EuroS&P. pp. 301–315. IEEE (2017)
8. Baldimtsi, F., Canetti, R., Yakoubov, S.: Universally composable accumulators. In: CT-RSA. LNCS, vol. 12006, pp. 638–666. Springer (2020)
9. Bar-Ilan, J., Beaver, D.: Non-cryptographic fault-tolerant computing in constant number of rounds of interaction. In: PODC. pp. 201–209. ACM (1989)
10. Barbulescu, R., Duquesne, S.: Updating key size estimations for pairings. IACR ePrint **2017**, 334 (2017)
11. Baric, N., Pfitzmann, B.: Collision-free accumulators and fail-stop signature schemes without trees. In: EUROCRYPT. LNCS, vol. 1233, pp. 480–494. Springer (1997)
12. Barreto, P.S.L.M., Lynn, B., Scott, M.: Constructing elliptic curves with prescribed embedding degrees. In: SCN. Lecture Notes in Computer Science, vol. 2576, pp. 257–267. Springer (2002)
13. Barreto, P.S.L.M., Naehrig, M.: Pairing-friendly elliptic curves of prime order. In: Selected Areas in Cryptography. LNCS, vol. 3897, pp. 319–331. Springer (2005)
14. Bauer, A., Herbst, N., Spinner, S., Ali-Eldin, A., Kounev, S.: Chameleon: A hybrid, proactive auto-scaling mechanism on a level-playing field. IEEE Trans. Parallel Distrib. Syst. **30**(4), 800–813 (2019)

15. Beaver, D.: Efficient multiparty protocols using circuit randomization. In: CRYPTO. LNCS, vol. 576, pp. 420–432. Springer (1991)
16. Bellare, M., Fuchsbauer, G., Scafuro, A.: Nizks with an untrusted CRS: security in the face of parameter subversion. In: ASIACRYPT (2). LNCS, vol. 10032, pp. 777–804 (2016)
17. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable zero knowledge with no trusted setup. In: CRYPTO (3). LNCS, vol. 11694, pp. 701–732. Springer (2019)
18. Benaloh, J.C., de Mare, M.: One-way accumulators: A decentralized alternative to digital signatures (extended abstract). In: EUROCRYPT. LNCS, vol. 765, pp. 274–285. Springer (1993)
19. Boneh, D., Boyen, X.: Short signatures without random oracles and the SDH assumption in bilinear groups. *J. Cryptology* **21**(2), 149–177 (2008)
20. Boneh, D., Bünz, B., Fisch, B.: Batching techniques for accumulators with applications to iops and stateless blockchains. In: CRYPTO (1). LNCS, vol. 11692, pp. 561–586. Springer (2019)
21. Bowe, S., Gabizon, A., Miers, I.: Scalable multi-party computation for zk-snark parameters in the random beacon model. *IACR ePrint* **2017**, 1050 (2017)
22. Bresson, E., Stern, J.: Efficient revocation in group signatures. In: PKC. LNCS, vol. 1992, pp. 190–206. Springer (2001)
23. Brickell, E., Li, J.: Enhanced privacy ID from bilinear pairing for hardware authentication and attestation. *IJIPSI* **1**(1), 3–33 (2011)
24. Brickell, E., Li, J.: Enhanced privacy ID: A direct anonymous attestation scheme with enhanced revocation capabilities. *IEEE Trans. Dependable Sec. Comput.* **9**(3), 345–360 (2012)
25. Camenisch, J., Lysyanskaya, A.: Dynamic accumulators and application to efficient revocation of anonymous credentials. In: CRYPTO. LNCS, vol. 2442, pp. 61–76. Springer (2002)
26. Campanelli, M., Gennaro, R., Goldfeder, S., Nizzardo, L.: Zero-knowledge contingent payments revisited: Attacks and payments for services. In: ACM CCS. pp. 229–243. ACM (2017)
27. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: FOCS. pp. 136–145. IEEE (2001)
28. Cramer, R., Damgård, I., Escudero, D., Scholl, P., Xing, C.: SPDZ_{2k}: Efficient MPC mod 2^k for dishonest majority. In: CRYPTO (2). LNCS, vol. 10992, pp. 769–798. Springer (2018)
29. Dalskov, A.P.K., Keller, M., Orlandi, C., Shrishak, K., Shulman, H.: Securing DNSSEC keys via threshold ECDSA from generic MPC. *IACR ePrint* **2019**, 889 (2019)
30. Damgård, I., Keller, M., Larraia, E., Pastro, V., Scholl, P., Smart, N.P.: Practical covertly secure MPC for dishonest majority - or: Breaking the SPDZ limits. In: ESORICS. LNCS, vol. 8134, pp. 1–18. Springer (2013)
31. Damgård, I., Koprowski, M.: Practical threshold RSA signatures without a trusted dealer. In: EUROCRYPT. LNCS, vol. 2045, pp. 152–165. Springer (2001)
32. Damgård, I., Nielsen, J.B.: Universally composable efficient multiparty computation from threshold homomorphic encryption. In: CRYPTO. LNCS, vol. 2729, pp. 247–264. Springer (2003)
33. Damgård, I., Pastro, V., Smart, N.P., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. In: CRYPTO. LNCS, vol. 7417, pp. 643–662. Springer (2012)

34. Derler, D., Hanser, C., Slamanig, D.: Revisiting cryptographic accumulators, additional properties and relations to other primitives. In: CT-RSA. LNCS, vol. 9048, pp. 127–144. Springer (2015)
35. Derler, D., Pöhls, H.C., Samelin, K., Slamanig, D.: A general framework for redactable signatures and new constructions. In: ICISC. LNCS, vol. 9558, pp. 3–19. Springer (2015)
36. Derler, D., Ramacher, S., Slamanig, D.: Post-quantum zero-knowledge proofs for accumulators with applications to ring signatures from symmetric-key primitives. In: PQCrypto. LNCS, vol. 10786, pp. 419–440. Springer (2018)
37. Dobson, S., Galbraith, S.D.: Trustless groups of unknown order with hyperelliptic curves. IACR ePrint **2020**, 196 (2020)
38. Dowling, B., Günther, F., Herath, U., Stebila, D.: Secure logging schemes and certificate transparency. In: ESORICS (2). LNCS, vol. 9879, pp. 140–158. Springer (2016)
39. Fan, C., Hsu, R., Manulis, M.: Group signature with constant revocation costs for signers and verifiers. In: CANS. LNCS, vol. 7092, pp. 214–233. Springer (2011)
40. Frederiksen, T.K., Lindell, Y., Osheter, V., Pinkas, B.: Fast distributed RSA key generation for semi-honest and malicious adversaries. In: CRYPTO (2). LNCS, vol. 10992, pp. 331–361. Springer (2018)
41. Fuchsbauer, G.: Subversion-zero-knowledge snarks. In: PKC (1). LNCS, vol. 10769, pp. 315–347. Springer (2018)
42. Garman, C., Green, M., Miers, I.: Decentralized anonymous credentials. In: NDSS. The Internet Society (2014)
43. Gennaro, R., Goldfeder, S.: Fast multiparty threshold ECDSA with fast trustless setup. In: ACM CCS. pp. 1179–1194. ACM (2018)
44. Gennaro, R., Goldfeder, S., Narayanan, A.: Threshold-optimal DSA/ECDSA signatures and an application to bitcoin wallet security. In: ACNS. LNCS, vol. 9696, pp. 156–174. Springer (2016)
45. Gilboa, N.: Two party RSA key generation. In: CRYPTO. LNCS, vol. 1666, pp. 116–129. Springer (1999)
46. Gilboa, N., Ishai, Y.: Distributed point functions and their applications. In: EUROCRYPT. LNCS, vol. 8441, pp. 640–658. Springer (2014)
47. Groth, J., Kohlweiss, M., Maller, M., Meiklejohn, S., Miers, I.: Updatable and universal common reference strings with applications to zk-snarks. In: CRYPTO (3). LNCS, vol. 10993, pp. 698–728. Springer (2018)
48. Hamdy, S., Möller, B.: Security of cryptosystems based on class groups of imaginary quadratic orders. In: ASIACRYPT. LNCS, vol. 1976, pp. 234–247. Springer (2000)
49. Hanser, C., Ramacher, S.: IAIK ECCelerate (2019), https://jce.iaik.tugraz.at/sic/Products/Core_Crypto_Toolkits/ECCelerate
50. Kales, D., Omolola, O., Ramacher, S.: Revisiting user privacy for certificate transparency. In: EuroS&P. pp. 432–447. IEEE (2019)
51. Katz, J., Kolesnikov, V., Wang, X.: Improved non-interactive zero knowledge with applications to post-quantum signatures. In: ACM CCS. pp. 525–537. ACM (2018)
52. Keller, M., Orsini, E., Scholl, P.: Actively secure OT extension with optimal overhead. In: CRYPTO (1). LNCS, vol. 9215, pp. 724–741. Springer (2015)
53. Keller, M., Orsini, E., Scholl, P.: MASCOT: faster malicious arithmetic secure computation with oblivious transfer. In: ACM CCS. pp. 830–842. ACM (2016)
54. Khovratovich, D., Law, J.: Sovrin: digital signatures in the blockchain area (2016), <https://sovrin.org/wp-content/uploads/AnonCred-RWC.pdf>
55. Laurie, B.: Certificate transparency. ACM Queue **12**(8), 10–19 (2014)

56. Laurie, B., Langley, A., Käsper, E.: Certificate transparency. RFC **6962**, 1–27 (2013)
57. Libert, B., Ling, S., Nguyen, K., Wang, H.: Zero-knowledge arguments for lattice-based accumulators: Logarithmic-size ring signatures and group signatures without trapdoors. In: EUROCRYPT (2). LNCS, vol. 9666, pp. 1–31. Springer (2016)
58. Lindell, Y.: Fast secure two-party ECDSA signing. In: CRYPTO (2). LNCS, vol. 10402, pp. 613–644. Springer (2017)
59. Lindell, Y., Nof, A.: Fast secure multiparty ECDSA with practical distributed key generation and applications to cryptocurrency custody. In: ACM CCS. pp. 1837–1854. ACM (2018)
60. Lipmaa, H.: Secure accumulators from euclidean rings without trusted setup. In: ACNS. LNCS, vol. 7341, pp. 224–240. Springer (2012)
61. Lueks, W., Goldberg, I.: Sublinear scaling for multi-client private information retrieval. In: Financial Cryptography. LNCS, vol. 8975, pp. 168–186. Springer (2015)
62. Menezes, A., Sarkar, P., Singh, S.: Challenges with assessing the impact of NFS advances on the security of pairing-based cryptography. In: Mycrypt. LNCS, vol. 10311, pp. 83–108. Springer (2016)
63. Merkle, R.C.: A certified digital signature. In: CRYPTO. LNCS, vol. 435, pp. 218–238. Springer (1989)
64. Miers, I., Garman, C., Green, M., Rubin, A.D.: Zerocoin: Anonymous distributed e-cash from bitcoin. In: IEEE S&P. pp. 397–411. IEEE (2013)
65. Mohassel, P., Rindal, P.: Aby^3 : A mixed protocol framework for machine learning. In: ACM CCS. pp. 35–52. ACM (2018)
66. Mohassel, P., Zhang, Y.: Secureml: A system for scalable privacy-preserving machine learning. In: IEEE S&P. pp. 19–38. IEEE (2017)
67. Nakanishi, T., Fujii, H., Hira, Y., Funabiki, N.: Revocable group signature schemes with constant costs for signing and verifying. In: PKC. LNCS, vol. 5443, pp. 463–480. Springer (2009)
68. Nakanishi, T., Kubooka, F., Hamada, N., Funabiki, N.: Group signature schemes with membership revocation for large groups. In: ACISP. LNCS, vol. 3574, pp. 443–454. Springer (2005)
69. Nakanishi, T., Sugiyama, Y.: A group signature scheme with efficient membership revocation for reasonable groups. In: ACISP. LNCS, vol. 3108, pp. 336–347. Springer (2004)
70. Nguyen, L.: Accumulators from bilinear pairings and applications. In: CT-RSA. LNCS, vol. 3376, pp. 275–292. Springer (2005)
71. Orsini, E., Smart, N.P., Vercauteren, F.: Overdrive2k: Efficient secure MPC over \mathbb{Z}_2^k from somewhat homomorphic encryption. IACR Cryptology ePrint Archive **2019**, 153 (2019)
72. Papamanthou, C., Tamassia, R., Triandopoulos, N.: Authenticated hash tables. In: ACM CCS. pp. 437–448. ACM (2008)
73. Pöhls, H.C., Samelin, K.: On updatable redactable signatures. In: ACNS. LNCS, vol. 8479, pp. 457–475. Springer (2014)
74. Sander, T.: Efficient accumulators without trapdoor extended abstracts. In: ICICS. LNCS, vol. 1726, pp. 252–262. Springer (1999)
75. Slamanig, D., Spreitzer, R., Unterluggauer, T.: Linking-based revocation for group signatures: A pragmatic approach for efficient revocation checks. In: Mycrypt. LNCS, vol. 10311, pp. 364–388. Springer (2016)
76. Smart, N.P., Alaoui, Y.T.: Distributing any elliptic curve based protocol. In: IMACC. LNCS, vol. 11929, pp. 342–366. Springer (2019)

77. University of Bristol: Multi-Protocol SPDZ (2019), <https://github.com/data61/MP-SPDZ>
78. Verheul, E.R.: Practical backward unlinkable revocation in fido, german e-id, idemix and u-prove. IACR ePrint **2016**, 217 (2016)
79. Wagh, S., Gupta, D., Chandran, N.: Securenn: 3-party secure computation for neural network training. PoPETs **2019**(3), 26–49 (2019)

A Assumptions

Let $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ be cyclic groups of prime order p . A *pairing* $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a map that is *bilinear* (i.e., for all $g_1, g'_1 \in \mathbb{G}_1$ and $g_2, g'_2 \in \mathbb{G}_2$, we have $e(g_1 \cdot g'_1, g_2) = e(g_1, g_2) \cdot e(g'_1, g_2)$ and $e(g_1, g_2 \cdot g'_2) = e(g_1, g_2) \cdot e(g_1, g'_2)$), *non-degenerate* (i.e., for generators $g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2$, we have that $e(g_1, g_2) \in \mathbb{G}_T$ is a generator), and *efficiently computable*. Let BGen be a PPT algorithm that, on input of a security parameter κ , outputs $\text{BG} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g_1, g_2) \leftarrow \text{BGen}(\kappa)$ for generators g_1 and g_2 of \mathbb{G}_1 and \mathbb{G}_2 , respectively, and $\Theta(\kappa)$ -bit prime p . If $\mathbb{G}_1 = \mathbb{G}_2$ the pairing is of Type-1 and if $\mathbb{G}_1 \neq \mathbb{G}_2$ and no non-trivial efficiently computable homomorphism $\mathbb{G}_2 \rightarrow \mathbb{G}_1$ exists, then it is of Type-3.

We recall the t -SDH assumption for Type-3 bilinear groups [19].

Definition 7 (t -SDH assumption). For $t > 0$, we define the advantage of an adversary \mathcal{A} as

$$\text{Adv}_{\text{BGen}, \mathcal{A}}^{t\text{-SDH}}(\kappa) = \Pr \left[\begin{array}{l} x \xleftarrow{R} \mathbb{Z}_p, (c, y) \leftarrow A \left(\text{BG}, (g_1^{x^i})_{i \in [t]}, g_2^x \right) \\ y = g_1^{(x+c)^{-1}} \end{array} \right].$$

The t -SDH assumption holds if $\text{Adv}_{\text{BGen}, \mathcal{A}}^{t\text{-SDH}}$ is a negligible function in the security parameter κ for all PPT adversaries \mathcal{A} .

In the Type-1 setting, the assumption can be simplified to only providing the powers of g_1 to the adversary, since $g_1 = g_2$.

B Omitted Proofs

B.1 Proof of Theorem 1

Proof. Assume that \mathcal{A} is an adversary against the collision freeness of the accumulator. We show that this adversary can be transformed into an efficient adversary \mathcal{B} against the t -SDH assumption. We perform a proof by reduction in the following way:

- When \mathcal{B} is started on a t -SDH instance $(\text{BG}, (g_1^{s^i})_{i \in [t]}, g_2^s)$, set $\text{pk}_{\mathcal{A}} \leftarrow (\text{BG}, (g_1^{s^i})_{i \in [t]}, g_2^s)$ and start \mathcal{A} on $\text{pk}_{\mathcal{A}}$. The oracles for \mathcal{A} are simulated by forwarding the inputs directly to the corresponding algorithms with $(\emptyset, \text{pk}_{\mathcal{A}})$ as argument for the keys.

- $\mathcal{O}^{\text{Eval}}(\mathcal{X})$: Return $\text{Eval}((\emptyset, \text{pk}_A), \mathcal{X})$.
 $\mathcal{O}^{\text{WitCreate}}(\Lambda_{\mathcal{X}}, \text{aux}, x)$: Return $\text{WitCreate}((\emptyset, \text{pk}_A), \mathcal{X})$.
 $\mathcal{O}^{\text{Add}}(\Lambda_{\mathcal{X}}, \text{aux}, x)$: Return $\text{Add}((\emptyset, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x)$.
 $\mathcal{O}^{\text{Delete}}(\Lambda_{\mathcal{X}}, \text{aux}, x)$: Return $\text{Delete}((\emptyset, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x)$.
 $\mathcal{O}^{\text{WitUpdate}}(\text{wit}_{x'}, \text{aux}, x)$: Return $\text{WitUpdate}((\emptyset, \text{pk}_A), \text{wit}_{x'}, \text{aux}, x)$.
- At some point \mathcal{A} outputs a set \mathcal{X}^* , an element $x^* \notin \mathcal{X}^*$, a witness $\text{wit}_{x^*}^*$ for x^* , and the randomizer r^* , such that for $\Lambda_{\mathcal{X}^*}, \text{aux} \leftarrow \text{Eval}((\emptyset, \text{pk}_A), \mathcal{X}^*)$, the verification relation $e(\Lambda_{\mathcal{X}^*}, g_2) = e(\text{wit}_{x^*}^*, g_2^{x^*} \cdot g_2^s)$ holds. Now, compute $h(X) = \prod_{x \in \mathcal{X}^*} (x + X)$ and $\phi(X)$ such that $h(X) = \phi(X)(x^* + X) + d$, which exists since $x^* \notin \mathcal{X}^*$. Then compute $g_1^{r^* \phi(s)}$ by expanding the polynomial $\phi(X)$ and the g_1^s stored in pk_A . Then, \mathcal{B} outputs

$$\begin{aligned}
 \left(\text{wit}_{x^*}^* \cdot \left(g_1^{r^* \phi(s)} \right)^{-1} \right)^{\frac{1}{r^* d}} &= \left(g_1^{\frac{r^* h(s)}{x^* + s}} \cdot g_1^{\frac{-r^* (h(s) - d)}{x^* + s}} \right)^{\frac{1}{r^* d}} \\
 &= \left(g_1^{\frac{r^* d}{x^* + s}} \right)^{\frac{1}{r^* d}} = g_1^{\frac{1}{x^* + s}}
 \end{aligned}$$

and x^* as solution to the t -SDH problem instance.

Hence, \mathcal{B} succeeds with the same probability as \mathcal{A} .

C Merkle-tree Accumulator

In Scheme 3, we cast the Merkle-tree accumulator in the framework of [34] as done in [36,50]. In practical instantiations, the requirement that Eval only works on sets of a size that is a power of 2 can be dropped. It is always possible to repeat the last element until the tree is of the correct size. Correctness can easily be verified. We restate the well-known fact that this accumulator is collision-free.

Lemma 1. *If $\{H_k\}_{k \in \mathbb{K}^\kappa}$ is a family of collision-resistant hash functions, the static accumulator in Scheme 3 is collision-free.*

D Strong-RSA Accumulator

We recall the strong RSA assumption [11].

Definition 8 (Strong RSA assumption). *Given two appropriately chosen primes p and q such that $N = pq$ has bit-length κ , then it holds for all PPT adversaries \mathcal{A} that*

$$\Pr \left[u \xleftarrow{R} \mathbb{Z}_N^*, (v, w) \leftarrow \mathcal{A}(u, N) : v^w \equiv u \pmod{N} \right]$$

is negligible in the security parameter κ .

<p>Gen($1^\kappa, t$): Fix a family of hash functions $\{H_k\}_{k \in \mathbb{K}^\kappa}$ with $H_k: \{0, 1\}^* \rightarrow \{0, 1\}^\kappa \forall k \in \mathbb{K}^\kappa$. Choose $k \xleftarrow{R} \mathbb{K}^\kappa$ and return $(\text{sk}_A, \text{pk}_A) \leftarrow (\emptyset, H_k)$.</p> <p>Eval($(\text{sk}_A, \text{pk}_A), \mathcal{X}$): Parse pk_A as H_k and \mathcal{X} as (x_0, \dots, x_{n-1}). If $\nexists k \in \mathbb{N}$ so that $n = 2^k$ return \perp. Otherwise, let $\ell_{u,v}$ refer to the u-th leaf (the leftmost leaf is indexed by 0) in the v-th layer (the root is indexed by 0) of a perfect binary tree. Return $\Lambda_{\mathcal{X}} \leftarrow \ell_{0,0}$ and $\text{aux} \leftarrow ((\ell_{u,v})_{u \in [n/2^{k-v}]}_{v \in [k]})$, where</p> $\ell_{u,v} \leftarrow \begin{cases} H_k(\ell_{2u,v+1} \ell_{2u+1,v+1}) & \text{if } v < k, \text{ and} \\ H_k(x_i) & \text{if } v = k. \end{cases}$ <p>WitCreate($(\text{sk}_A, \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x_i$): Parse aux as $((\ell_{u,v})_{u \in [n/2^{k-v}]}_{v \in [k]})$ and return wit_{x_i} where</p> $\text{wit}_{x_i} \leftarrow (\ell_{\lfloor i/2^v \rfloor + \eta, k-v})_{0 \leq v \leq k}, \eta = \begin{cases} 1 & \text{if } \lfloor i/2^v \rfloor \pmod{2} = 0 \\ -1 & \text{otherwise.} \end{cases}$ <p>Verify($\text{pk}_A, \Lambda_{\mathcal{X}}, \text{wit}_{x_i}, x_i$): Parse pk_A as H_k, $\Lambda_{\mathcal{X}}$ as $\ell_{0,0}$, set $\ell_{i,k} \leftarrow H_k(x_i)$. Recursively check for all $0 < v < k$ whether the following holds and return 1 if so. Otherwise return 0.</p> $\ell_{\lfloor i/2^{v+1} \rfloor, k-(v+1)} = \begin{cases} H_k(\ell_{\lfloor i/2^v \rfloor, k-v} \ell_{\lfloor i/2^v \rfloor + 1, k-v}) & \text{if } 2 \mid \lfloor i/2^v \rfloor \\ H_k(\ell_{\lfloor i/2^v \rfloor - 1, k-v} \ell_{\lfloor i/2^v \rfloor, k-v}) & \text{otherwise.} \end{cases}$
--

Scheme 3: Merkle-tree accumulator.

Major lines of work investigated accumulators in the hidden order groups, i.e., RSA-based, and the known order groups, i.e., discrete logarithm-based, setting. The first collision-free RSA-based accumulator is due to Baric and Pfitzmann [11]. The accumulator in this construction consists of a generator raised to the product of all elements of the set. Then witnesses essentially consist of the same value skipping the respective elements in the product. Thereby, the witness can easily be verified by raising the power of the witness to the element and checking if the result matches the accumulator. We recall the RSA-based accumulator in Scheme 4.

Note that, whenever the factorization of N is available, the Chinese remainder theorem can be used to speed up the computations. For **WitCreate** and **WitUpdate** we can use the factorization to compute inverses $\pmod{(p-1) \cdot (q-1)}$. Deletion of values from the accumulator is not possible if the factorization is unknown.

Correctness can easily be verified and collision freeness follows from the strong RSA assumption:

Theorem 4 ([11]). *If the strong RSA assumption holds, Scheme 4 is collision-free.*

Again, from Theorems 2 and 4, it follows that Scheme 4 is also secure in the UC model:

Corollary 2. *Scheme 4 emulates \mathcal{F}_{Acc} in the UC model.*

Gen ($1^\kappa, t$):	Choose an RSA modulus $N = p \cdot q$ with two large safe primes p, q , and let g be a random quadratic residue $\pmod N$. Set $\text{sk}_A \leftarrow \emptyset$ and $\text{pk}_A \leftarrow (N, g)$.
Eval ($(\text{sk}_A, \text{pk}_A), \mathcal{X}$):	Parse pk_A as (N, g) and let $\mathcal{X} \subset \mathbb{P}$. Return $A_{\mathcal{X}} \leftarrow g^{\prod_{x \in \mathcal{X}} x} \pmod N$ and $\text{aux} \leftarrow (\text{add} \leftarrow 0, \mathcal{X})$.
WitCreate ($(\text{sk}_A, \text{pk}_A), A_{\mathcal{X}}, \text{aux}, x$):	If $x \notin \mathcal{X}$, return \perp . If $\text{sk}_A \neq \emptyset$, return $\text{wit}_x \leftarrow A_{\mathcal{X}}^{x^{-1}} \pmod N$, otherwise return $\text{wit}_x \leftarrow g^{\prod_{x' \in \mathcal{X} \setminus \{x\}} x'} \pmod N$.
Verify ($\text{pk}_A, A_{\mathcal{X}}, \text{wit}_x, x$):	Parse pk_A as (N, g) . If $\text{wit}_x^x = A_{\mathcal{X}} \pmod N$ holds, return 1, otherwise return 0.
Add ($(\text{sk}_A, \text{pk}_A), A_{\mathcal{X}}, \text{aux}, x$):	Parse pk_A as (N, g) and aux as \mathcal{X} . If $x \in \mathcal{X}$, return \perp . Set $\mathcal{X}' \leftarrow \mathcal{X} \cup \{x\}$, $\text{aux}' \leftarrow (\mathcal{X}', \text{add} \leftarrow 1)$, and $A_{\mathcal{X}'} \leftarrow A_{\mathcal{X}}^x \pmod N$. Return $A_{\mathcal{X}'}$ and aux' .
Delete ($(\text{sk}_A, \text{pk}_A), A_{\mathcal{X}}, \text{aux}, x$):	Parse pk_A as (N, g) and aux as \mathcal{X} . If $x \notin \mathcal{X}$, return \perp . If $\text{sk}_A \neq \emptyset$, set $\mathcal{X}' \leftarrow \mathcal{X} \setminus \{x\}$, $\text{aux}' \leftarrow (\mathcal{X}', \text{add} \leftarrow -1)$, and $A_{\mathcal{X}'} \leftarrow A_{\mathcal{X}}^{x^{-1}} \pmod N$. Otherwise, compute $(A_{\mathcal{X}'}, \text{aux}') \leftarrow \text{Eval}((\emptyset, \text{pk}_A), \mathcal{X} \setminus \{x\})$ with $\text{add} \leftarrow -1$ in aux' . Return $A_{\mathcal{X}'}$ and aux' .
WitUpdate ($(\text{sk}_A, \text{pk}_A), \text{wit}_{x_i}, \text{aux}, x$):	Parse pk_A as (N, g) and aux as $(\mathcal{X}, \text{add})$. If $\text{add} = 0$, return \perp . Return $\text{wit}_{x_i}^x \pmod N$ if $\text{add} = 1$. If instead $\text{add} = -1$ and $\text{sk}_A \neq \emptyset$, then return $\text{wit}_{x_i}^{x^{-1}} \pmod N$. Otherwise, compute $a, b \in \mathbb{Z}$ such that $ax_i + bx = 1$ and return $\text{wit}_{x_i}^b \cdot A_{\mathcal{X}}^a \pmod N$. In the last two cases in addition return $\text{aux} \leftarrow (\text{add} \leftarrow 0)$.

Scheme 4: Strong RSA-based accumulator.

D.1 Dynamic Linear Secret-Shared Accumulator From the Strong RSA Assumption

For our dynamic linear secret-shared accumulator from the strong RSA assumption, observe that the two main operations that have to be performed in the context of an MPC protocol are the sampling of the secret prime factors as well as the computation of the inverses of the public elements in the exponent. Both operations are also performed during RSA key generation with the sole difference that in this case, the public exponent is inverted. Therefore, we can make use of any protocol for distributed RSA key generation. In particular, our construction makes use of the state-of-the-art protocol by Frederiksen et al. [40] for the two-party case.

In the following, we will recap the structure of this UC secure two-party protocol and highlight the most essential parts of the protocol. The key generation in the malicious setting consists of the following four phases:

Candidate Generation: Both parties P_1 and P_2 choose random shares $p_1 \in \mathbb{N}$ respectively $p_2 \in \mathbb{N}$ and commit to it. Based on maliciously secure OT, they do a secure trial division of $p = p_1 + p_2$ with public threshold $B_1 \in \mathbb{N}$.

Construct Modulus: Given shared candidate primes $p = p_1 + p_2, q = q_1 + q_2$ the parties compute $N = pq$ by a custom version of the Gilboa protocol [45]. The candidate modulus N is sent to both parties.

Verify Modulus: This phase consists of three steps:

1. Second trial division with threshold $B_2 > B_1$.

2. Secure biprimality test.
3. Proof of honesty that checks the commitments and whether $\gcd(e, \phi(N)) = 1$, where e is the public exponent.

Construct Keys: Computes the shares of $d = d_1 + d_2$ such that $e(d_1 + d_2) \equiv 1 \pmod{\phi(N)}$ (uses additional output from the proof of honesty).

We will denote this protocol by Π_{RSA} and by \mathcal{F}_{RSA} its ideal functionality. At a later point in the RSA-based accumulator, we need to invert an element $x \in \mathbb{P}$ in \mathbb{Z}_N^* . Since neither P_1 nor P_2 knows the order $\phi(N)$ of this ring, we employ parts of the MPC protocol. For this task, we will perform the 3rd step of **Verify Modulus** with $e = x$, and then immediately run **Construct Keys**. The output of this sub-protocol, which we will denote by $\text{Invert}_{\phi(N)}(x)$, is a secret-shared value $\langle y \rangle = y_1 + y_2$, where y_i is the share of party i , s.t. $x(y_1 + y_2) \equiv 1 \pmod{\phi(N)}$.

<p>Gen($1^\kappa, t$): Generate an RSA modulus $N = (p_1 + p_2) \cdot (q_1 + q_2)$ via the protocol Π_{RSA}, where the party P_i receives (p_i, q_i) for $i = 1, 2$. Further, let g be a random quadratic residue \pmod{N}. Set $\text{pk}_A \leftarrow (N, g)$.</p>
<p>Eval($(\text{sk}_A, \text{pk}_A), \mathcal{X}$): Parse pk_A as (N, g) and $\mathcal{X} \subset \mathbb{P}$. Return $\Lambda_{\mathcal{X}} \leftarrow g^{\prod_{x \in \mathcal{X}} x} \pmod{N}$ and $\text{aux} \leftarrow \mathcal{X}$.</p>
<p>WitCreate($((p_1, q_1), (p_2, q_2), \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x$): If $x \notin \mathcal{X}$, return \perp. Compute $\langle y \rangle \leftarrow \text{Invert}_{\phi(N)}(x)$ and return $\text{wit}_x \leftarrow \text{Open}(\Lambda_{\mathcal{X}}^{\langle y \rangle} \pmod{N})$.</p>
<p>Verify($\text{pk}_A, \Lambda_{\mathcal{X}}, \text{wit}_x, x$): Parse pk_A as (N, g). If $\text{wit}_x^x = \Lambda_{\mathcal{X}} \pmod{N}$ holds, return 1, otherwise return 0.</p>
<p>Add($\text{pk}_A, \Lambda_{\mathcal{X}}, \text{aux}, x$): Parse pk_A as (N, g) and aux as \mathcal{X}. If $x \in \mathcal{X}$, return \perp. Set $\mathcal{X}' \leftarrow \mathcal{X} \cup \{x\}$, $\text{aux}' \leftarrow (\mathcal{X}', \text{add} \leftarrow 1)$, and $\Lambda_{\mathcal{X}'} \leftarrow \Lambda_{\mathcal{X}}^x \pmod{N}$. Return $\Lambda_{\mathcal{X}'}$ and aux'.</p>
<p>Delete($((p_1, q_1), (p_2, q_2), \text{pk}_A), \Lambda_{\mathcal{X}}, \text{aux}, x$): Parse pk_A as (N, g) and aux as \mathcal{X}. If $x \notin \mathcal{X}$, return \perp. Set $\mathcal{X}' \leftarrow \mathcal{X} \setminus \{x\}$, $\text{aux}' \leftarrow (\mathcal{X}', \text{add} \leftarrow -1)$, and $\langle y \rangle \leftarrow \text{Invert}_{\phi(N)}(x)$. Return $\Lambda_{\mathcal{X}'} \leftarrow \text{Open}(\Lambda_{\mathcal{X}}^{\langle y \rangle} \pmod{N})$ and aux'.</p>
<p>WitUpdate($((p_1, q_1), (p_2, q_2), \text{pk}_A), \text{wit}_{x_i}, \text{aux}, x$): Parse pk_A as (N, g) and aux as $(\mathcal{X}, \text{add})$. Return \perp if $\text{add} = 0 \vee x_i \notin \mathcal{X}$. Return $\text{wit}_{x_i}^x \pmod{N}$ if $\text{add} = 1$. If instead $\text{add} = -1$ and compute $\langle y \rangle \leftarrow \text{Invert}_{\phi(N)}(x)$, then return $\text{Open}(\text{wit}_{x_i}^{\langle y \rangle} \pmod{N})$. In the last two cases in addition return $\text{aux} \leftarrow (\text{add} \leftarrow 0)$.</p>

Scheme 5: MPC-RSA: Dynamic linear secret-shared accumulator based on RSA for two parties.

However, we have to note that the used distributed RSA key generation protocols may leak a small amount of secret information. We refer to the full version of [40] for a detailed discussion of this leakage. Since our protocol may invert multiple elements, the parties might need to add bounds on the maximum number of operations to prevent leakage of the secret key.

The last open point during **Gen** is the sampling of the quadratic non-residue \pmod{N} . An option here is to simply sample g at random from \mathbb{Z}_N^* and checking whether the Jacobi symbol satisfies $(\frac{g}{N}) = -1$. Despite its definition as the

product of the Legendre symbols of the prime factors of N , the Jacobi symbol can be computed efficiently without knowledge of the prime factors using an algorithm analogous to the Euclidean algorithm. Hence, we perform this step outside of the MPC protocol once N was generated.

In the security analysis of Scheme 5, we can reuse the arguments of Theorem 3. Therefore, we will omit the proof of the following theorem. We, however, note, that the theorem only holds provided that the parties keep track of the number of potentially leaked bits and abort if this number gets too large.

Theorem 5. *Scheme 5 UC emulates $\mathcal{F}_{Acc-MPC}$ in the $\mathcal{F}_{SPDZ+}, \mathcal{F}_{RSA}$ -hybrid model.*

E FRESCO Benchmarks

In this section, we discuss the benchmarks of our t -SDH implementation in the maliciously secure setting in the FRESCO [3] framework. FRESCO is a Java framework facilitating fast prototyping of MPC-based applications and protocols. FRESCO implements the SPDZ protocol and various extensions [33,52,53,28]. For pairing and elliptic curve group operations, we rely on the ECCelerate library [49] and integrate $\text{Exp}_{\mathbb{G}}$, $\text{Output-}\mathbb{G}$, and the corresponding $\text{MACCheck}_{\text{ECC}}$ algorithms of [76] into FRESCO. As a pairing-friendly curve, a 400-bit Barreto-Naehrig curve [13] is used, which provides around 100 bit of security following recent estimates [10,62].

The offline phase performance of our FRESCO implementation can be seen in Table 6. A comparison to our MP-SPDZ implementation (Table 3) shows that the offline phase in FRESCO is not yet as optimized as the one of MP-SPDZ (as an example, the BaseOTs used in FRESCO are not using elliptic curve arithmetic) and that the performance of the latter is more indicative of an optimized implementation. Further note that batching the generation of many triples together like for the Eval phase is more efficient in practice than producing a single triple and as these triples are not dependant on the input, all parties can continuously generate triples in the background to fill a triple-buffer for use in the online phase.

In Table 7, we present the online phase performance of our t -SDH implementation in the FRESCO framework. Since the FRESCO implementation does not support a depth-optimized tree-like multiplication, the Eval operation scales worse with the number of parties. Compared to our MP-SPDZ implementation (Table 4) it is, therefore, much slower, especially in the WAN setting.

Finally, we present the communication complexity of our t -SDH implementation in the FRESCO framework in Table 8. Again, a comparison to our MP-SPDZ implementation (Table 5) shows, that the FRESCO framework requires more communication to achieve the same result. While some of this overhead can be attributed to the different choice of elliptic curve, the rest is inherent to the implementation of the framework.

Table 6. Offline phase performance of different steps of the MPC- t -SDH accumulator with access to the secret trapdoor implemented in FRESCO. Time in seconds.

Operation	$ \mathcal{X} $	LAN setting				WAN setting			
		$n = 2$	3	4	5	$n = 2$	3	4	5
BaseOTs	2^{10}	21.5	60.3	104.6	148.9	76.5	215.0	364.2	504.3
	2^{14}	21.5	60.3	104.6	148.9	76.5	215.0	364.2	504.3
Eval	2^{10}	54.4	154.0	265.0	409.1	95.7	283.1	465.2	683.3
	2^{14}	825.8	2259.5	4048.8	6150.9	1449.8	3587.1	6578.4	10056.3
Inverse	2^{10}	1.3	15.3	16.8	19.7	3.2	68.2	70.4	72.8
	2^{10}	1.3	15.3	16.8	19.7	3.2	68.2	70.4	72.8

Table 7. Online phase performance of the MPC- t -SDH accumulator with access to the secret trapdoor implemented in FRESCO, for both the LAN and WAN settings with n parties. Time in milliseconds averaged over 50 executions.

Operation	$ \mathcal{X} $	LAN setting				WAN setting			
		$n = 2$	3	4	5	$n = 2$	3	4	5
Gen	2^{10}	78	106	301	772	604	1124	1399	1684
	2^{14}	75	110	246	766	608	1118	1401	1673
Eval	2^{10}	133	354	4062	14193	52316	56802	58728	60043
	2^{14}	1389	4362	61222	196563	828522	892293	918445	935274
WitCreate	2^{10}	40	84	279	906	1109	1858	2156	2444
	2^{14}	41	98	267	992	1120	1853	2152	2442
Add	2^{10}	43	78	231	646	574	1088	1373	1650
	2^{14}	41	74	225	706	571	1087	1363	1642
WitUpdate _{Add}	2^{10}	40	72	259	745	569	1094	1372	1649
	2^{14}	40	85	213	732	564	1088	1369	1644
Delete	2^{10}	47	80	299	958	1075	1849	2150	2454
	2^{14}	42	92	297	857	1071	1847	2152	2446
WitUpdate _{Delete}	2^{10}	40	82	258	977	1071	1852	2151	2447
	2^{14}	38	38	294	880	1077	1852	2149	2441

Table 8. Communication complexity of the MPC- t -SDH accumulator with access to the secret trapdoor implemented in FRESCO, per party. Communication in MiB.

Operation	$ \mathcal{X} $	Offline phase ^a	Online Phase ^b
Gen	2^{10}	0.297	0.103
	2^{14}	0.297	0.103
Eval	2^{10}	220.971	0.176
	2^{14}	3 530.148	3.164
WitCreate	2^{10}	0.634	0.041
	2^{14}	0.634	0.041
Add	2^{10}	0.297	0.039
	2^{14}	0.297	0.039
WitUpdate _{Add}	2^{10}	0.297	0.039
	2^{14}	0.297	0.039
Delete	2^{10}	0.634	0.041
	2^{14}	0.634	0.041
WitUpdate _{Delete}	2^{10}	0.634	0.041
	2^{14}	0.634	0.041

^a Includes BaseOTs for a new connection

^b Includes the setup of a fresh MAC for each share of the secret trapdoor