Accumulators with Applications to Anonymity-Preserving

Foteini Baldimtsi[1], Jan Camenisch[2], Maria Dubovitskaya[2], Anna Lysyanskaya[3],
Leonid Reyzin[4], Kai Samelin[2,5], and Sophia Yakoubov[4]

[1] George Mason University, Fairfax, VA, USA
foteini@gmu.edu
[2] IBM Research – Zurich, Rüschlikon, Switzerland
{jca, mdu, ksa}@zurich.ibm.com
[3] Brown University, Providence, RI, USA
anna@cs.brown.edu
[4] Boston University, Boston, MA, USA
{reyzin, sonka}@bu.edu
[5] TU Darmstadt, Darmstadt, Germany

Abstract. Membership revocation is essential for cryptographic applications, from traditional PKIs to group signatures and anonymous credentials. Of the various solutions for the revocation problem that have been explored, dynamic accumulators are one of the most promising. We propose Braavos, a new, RSA-based, dynamic accumulator. It has optimal communication complexity and, when combined with efficient zero-knowledge proofs, provides an ideal solution for anonymous revocation. For the construction of Braavos we use a modular approach: we show how to build an accumulator with better functionality and security from accumulators with fewer features and weaker security guarantees. We then describe an anonymous revocation component (ARC) that can be instantiated using any dynamic accumulator. ARC can be added to any anonymous system, such as anonymous credentials or group signatures, in order to equip it with a revocation functionality. Finally, we implement ARC with Braavos and plug it into Idemix, the leading implementation of anonymous credentials. This work resolves, for the first time, the problem of practical revocation for anonymous credential systems.

Table of Contents

Accumulators with Applications to Anonymity-Preserving . . . 1
Foteini Baldimtsi, Jan Camenisch, Maria Dubovitskaya, Anna Lysyanskaya, Leonid Reyzin, Kai Samelin, and Sophia Yakoubov
1 Introduction . . . 3
1.1 Our Contributions . . . 3
2 Definitions: A Modular View of Accumulator Functionality . . . 5
2.1 Accumulator Algorithms . . . 6
2.2 Accumulator Security Properties . . . 7
3 Modular Accumulator Constructions . . . 8
3.1 Leveraging Accumulators with Different Functionalities . . . 8
3.2 Leveraging Less Secure Accumulators . . . 10
4 Braavos: A Communication-Optimal Adaptively Sound Dynamic Accumulator . . . 11
4.1 CL-RSA-B: A Communication-Optimal Non-Adaptively Sound Dynamic Accumulator . . . 13
4.2 Braavos Soundness . . . 16
4.3 Comparison with Other Constructions . . . 16
4.4 Adding Zero Knowledge to Braavos . . . 18
5 Anonymous Revocation from Accumulators . . . 18
5.1 Anonymous Revocation Component (ARC) with Accumulators . . . 19
5.2 Revocation for Anonymous Credentials Using Braavos: Performance Evaluation . . . 22
6 Acknowledgements . . . 25
A Lower Bound on Total Communication in Negative Accumulators . . . 27
B BraavosB: Another Communication-Optimal Adaptively Sound Dynamic Accumulator . . . 27
B.1 Range-RSA: A Dynamic Negative Accumulator . . . 28
B.2 Range-RSA Accumulator Algorithms . . . 30
B.3 BraavosB Soundness . . . 30
B.4 Adding Zero Knowledge to BraavosB . . . 30
C Security of ARC . . . 33
C.1 Recalling Security Definitions by Camenisch et al. . . . 33
C.2 Security Definitions for ARC with Join . . . 34

1 Introduction

Authentication of users is vital to most of the electronic systems we use today. It is usually achieved by giving the user a token, or *credential*, that the user must present to prove that she has permission to access a service. An important challenge that such systems face is how to *revoke* a user's privileges in case she misbehaves or her credential gets compromised.

Achieving revocation in practice has been shown to be a very complex problem. The two obvious approaches are *whitelists*, where a user is valid if her public key or identity is on a whitelist, and *blacklists*, where anyone not on the blacklist can be presumed to be a valid user. Both of these solutions are problematic, because the party that manages revocation (hereafter called the *revocation authority*) needs to distribute large lists and update them continuously.

In anonymous settings, where user authentication must not reveal the user's identity, things get even more complicated. In order for a user to show that she hasn't been revoked, she must prove that she is on the whitelist (or is not on the blacklist) in *zero knowledge*, which requires work linear in the size of the list. Having to perform linear work in the number of whitelisted (or blacklisted) users is far from practical; thus, we look into solutions where the users and verifiers in charge of authentication only have to do a *constant* amount of work.

One of the most promising solutions to revocation in anonymous settings has been the use of *cryptographic accumulators* [11,33,9]. An accumulator is a binding (but not necessarily hiding) commitment to a set S of elements. This set can change over time, as elements are added and deleted; accumulators that support both additions and deletions are called *dynamic*. Known dynamic accumulators are based on Merkle hash trees [7], variants of RSA [3,11,26], and bilinear maps [33,18,1]. Upon the addition of a new element x to the set, a *witness* w, or *proof of membership*, is generated for the element. This witness is later used, in conjunction with the accumulator value, to demonstrate that the element is, in fact, in the set. The witnesses may evolve as the accumulated set changes.

A dynamic accumulator can be used as an anonymous revocation mechanism [11,9]. The accumulator contains the whitelist S of users. The revocation authority then only needs to maintain and distribute a short accumulator value, not an entire large whitelist. When a user is issued a credential, she is given a membership witness w for her element x. While she hasn't been revoked, she can prove in zero-knowledge that her x is in the current accumulator, i.e., that she knows (x, w) such that x corresponds to her identity and w is a valid membership witness for x. The complexity of this proof is independent of the numbers of added and revoked users. When the user is revoked, x is removed from the accumulator.

However, there are two downsides to using existing accumulator constructions for revocation in this way. The first is the high communication cost. Each user has to update her membership witness every time the accumulator value changes, which typically happens with every addition or deletion. In order to give the users the information they need for these updates, the revocation authority has to send a broadcast message to all users. The frequency of these broadcasts is unattractive, especially when constrained devices are used.

The second downside is linkability: users who receive these broadcast messages might be able to link revocations of their fellow users with the corresponding additions. That is, when a revocation occurs, the associated broadcast message might contain information related to a previous addition broadcast message, allowing recipients to determine that the revoked user is the same user who was added at a certain previous point. This join-revoke linkability is a breach in anonymity.

1.1 Our Contributions

In this work, we focus on building accumulators that avoid the two downsides described above, namely those of communication overhead and join-revoke linkability. Our contributions are threefold. First, we provide a new modular definitional view of accumulators. Second, we use this modular view to build a new dynamic accumulator with optimal communication cost and join-revoke unlinkability. Finally, we explain how our proposed accumulator can be used to instantiate an anonymous revocation component (ARC) and showcase that our scheme can provide efficient revocation for anonymous credential systems by implementing it for Idemix [16] (IBM's anonymous credential system).

Compositional Accumulator Framework. In Section 2, we give a systematic view of accumulator properties, describing them in a modular way. In Section 3, we leverage our modular view to show simple ways of obtaining new accumulators with more functionality (adding and deleting, proving membership and non-membership) and

better security by combining accumulators that possess only some of the properties. This view of accumulators also sheds light on trade-offs between functionality and efficiency.

Constructions. Utilizing our compositional framework, in Section 4 and Appendix B we introduce two new dynamic accumulators which we call Braavos and BraavosB.[6] All previously known dynamic accumulators with strong security guarantees require that the accumulator and each membership witness be updated *both* every time a new element is added to the accumulator, and every time an element is deleted from the accumulator. In Braavos and BraavosB, the accumulator value and membership witnesses are updated *only* when an element is deleted, achieving optimal communication complexity. In particular, this means that the accumulator manager does not need to broadcast any information to the witness holders when a new element addition occurs, making it practical to support a high element addition rate. Because element addition does not involve the communication of any information to other witness holders, Braavos and BraavosB also ensure that a revocation event is not linkable to a prior addition event, achieving join-revoke unlinkability.

Our Braavos accumulator builds upon the dynamic CL-RSA-B accumulator, which was informally introduced as a brief remark by Camenisch and Lysyanskaya [11]. CL-RSA-B has the communication optimality of Braavos, but has a weaker security guarantee which makes it usable only for the accumulation of random elements. CL-RSA-B is only secure against non-adaptive attacks, as described in Section 2.2.[7] Braavos uses CL-RSA-B to accumulate random elements, and uses digital signatures to bind the random elements to the actual elements to be accumulated. Because neither CL-RSA-B nor digital signatures require witness updates upon additions, Braavos does not either.

BraavosB builds upon another new accumulator which we call range-RSA. Range-RSA is dynamic, and supports proofs of *non*-membership (that is, it is *negative*, as defined in Section 2; we show that it can be modified slightly to be *universal*, i.e., to also support proofs of membership). BraavosB uses digital signatures to accumulate added elements, and range-RSA to accumulate deleted elements. So, to prove the membership of x in BraavosB one must prove possession of a digital signature on x, and *non*-membership in range-RSA. The non-membership witness generation in range-RSA is more efficient than in any existing negative dynamic accumulator, which makes it uniquely suitable for contexts where non-membership witnesses are generated at a high rate. The efficiency of non-membership witness generation in range-RSA makes it possible for BraavosB to support high rates of element addition.

While the asymptotic runtime and communication complexities of Braavos and BraavosB are equally good, in practice Braavos has several advantages over BraavosB. These include more efficient zero knowledge proofs of membership, smaller accumulator manager storage requirements, and the fact that in Braavos the update message associated with a deletion does not reveal the deleted element.

Application and Implementation. In Section 5, we describe an anonymous revocation component (ARC) which uses accumulators to enable revocation in a variety of existing (as well as new) anonymous schemes. Our definition of the ARC is an extension of the revocation component described by Camenisch et. al [10] that allows the addition of users after the initial setup, as well as their re-addition after they have been revoked. When Braavos is the accumulator used to instantiate an ARC, members need to update their witnesses only when another member is revoked (but not when another member is added), which is key since the number of revocations in a typical system is much lower than the number of additions. Consider, for instance, PKI certificates. Around 8% of unexpired certificates are revoked [29]. This number rose from 1% in May 2014 because of the Heartbleed vulnerability, and is expected to eventually return to 1%. However, it is vital to be able to revoke this 1%, since the compromised keys can be used to inflict a lot of damage. Moreover, this has to be done without flooding certificate holders with instructions to update their witnesses roughly 100 times more frequently than necessary.

In Section 5.2, we show how one can directly plug an ARC with Braavos into the leading implementation of anonymous credentials (idemix) [16]) with no modifications on either side. This compatibility with an existing working system resolves, for the first time, the problem of practical revocation for anonymous credentials.

We measure the performance of the resulting system, and demonstrate that, for realistic credential presentation policies, Braavos adds only a ≈10% runtime overhead, both for generation and for verification. In more detail, even for reasonably high security parameter sizes (2,048 bits), credential generation without revocation takes ≈1.6s, while with revocation it takes ≈1.8s. For verification these numbers become even closer: without revocation

[6] The Faceless Men of Braavos, from the Game of Thrones, are able to assume unrecognizable personas, unlinkable to their other personas.

[7] As a side contribution, in Section 4.1 we formalize and prove the security of CL-RSA-B.

we need ≈1.7s, while with revocation we need ≈1.8s. Clearly, this is a small price to pay for the benefits of revocation.

2 Definitions: A Modular View of Accumulator Functionality

λ: The security parameter.

D: The domain of the accumulator (the set of elements that the accumulator can accumulate). Often, D includes all elements (e.g., $\{0,1\}^*$). Sometimes, D is more limited (e.g., primes of a certain size).

sk: The accumulator manager's secret key or trapdoor. (The corresponding public key, if one exists, is not modeled here as it can be considered to be a part of the accumulator itself.)

t: A discrete time / operation counter.

a_t: The accumulator at time t.

m_t: Any auxiliary values which might be necessary for the maintenance of the accumulator. These are typically held by the accumulator manager. Note that while the accumulator itself should be constant (or at least sub-linear) in size, m may be larger.

S_t: The set of elements in the accumulated set at time t. Note that S_0 can be instantiated to be different, based on the initial sets supported by the accumulator in question. Most accumulators assume $S_0 = \emptyset$.

x, y: Elements which might be added to the accumulator.

w_t^x, u_t^x: The witness that element x is (respectively, is not) in accumulator a_t at time t.

$\mathsf{upmsgwh}_t$: A broadcast message sent (by the accumulator manager, if one exists) at time t to all witness holders immediately after the accumulator has been updated. This message is meant to enable all witness holders to update the witnesses they hold for consistency with the new accumulator. It will often contain the new accumulator a_t, and the nature of the update itself (e.g., "x has been added and witness w_t^x has been produced"). It may also contain other information.

$\mathsf{upmsgtp}_t$: A broadcast message sent (by the accumulator manager, if one exists) at time t to all third party verifiers immediately after the accumulator has been updated. This message is meant to provide the third parties with up to date values (the accumulator a_t) to verify proofs of membership or non-membership against. It may also contain other information.

upmsg_t: All of the information broadcast (by the accumulator manager, if one exists) at time t. $\mathsf{upmsg}_t = (\mathsf{upmsgwh}_t, \mathsf{upmsgtp}_t)$. We often use upmsg instead of $\mathsf{upmsgwh}$ and $\mathsf{upmsgtp}$, since these are frequently the same.

Fig. 1. Accumulator algorithm input and output parameters.

In this section, we provide a unified view of accumulator properties. We separately consider accumulators that support additions, deletions or both, and accumulators that support membership proofs, non-membership proofs or both. We have a different focus than Derler et al. [19], who also discuss the unification of accumulators. Namely, unlike Derler et al., we emphasize the modularity of our definitions: subsets of the properties we define can be easily combined, resulting in a broader range of accumulator types than previously described. To limit our scope, we only consider accumulator properties pertaining to functionality and soundness in this section; properties pertaining to privacy are discussed in Section 4.4 and Appendix B.4.

Various flavors of accumulator functionality definitions have been restated in literature a number of times. We leverage the definitions provided by Reyzin and Yakoubov [35], but reformulate them to offer a more complete view of the accumulator space. We begin by introducing four basic kinds of accumulator primitives.

- *static accumulator*: represents a fixed set.
- *additive accumulator*: supports only additions.
- *subtractive accumulator*: supports only deletions.
- *dynamic accumulator*: supports both additions and deletions (as defined in [11]).

Deletions and additions can, of course, be performed simply by re-instantiating the accumulator with the updated set. This takes at most a polynomial amount of time in the number of element additions or deletions which have been performed up until that point, but is not, in general, practical. A dynamic accumulator should support both additions and deletions in time which is either independent of the number of operations performed altogether, or is sublinear in this number.

We also describe accumulators in terms of the kinds of proofs (membership proofs, non-membership proofs, or both) they support.

– *positive accumulator*: supports membership proofs.
– *negative accumulator*: supports non-membership proofs.
– *universal accumulator [26]*: supports both.

2.1 Accumulator Algorithms

Next, we describe the algorithms used by all of these accumulator primitives. For convenience, Figure 1 enumerates and explains all algorithm input and output parameters.

We consider three types of parties: (1) an accumulator manager, (2) an entity responsible for an element and its corresponding witness (from hereon-out referred to as *witness holder*), and (3) a third party (e.g. a verifier, who is given any relevant witnesses by the witness holder at the time of verification). Parameters which are omitted in some schemes have a bar on top.

The following are algorithms performed by the accumulator manager:

– $\mathsf{Gen}(1^\lambda, S_0) \rightarrow (\overline{sk}, a_0, \overline{m_0})$ instantiates the accumulator manager's secret key sk, the accumulator a_0 (representing the initial set $S_0 \subseteq D$ of elements in the accumulator, where D is the domain of the accumulator), and the auxiliary value m_0 necessary for the maintenance of the accumulator. m can be thought of as the accumulator manager's memory or storage. The allowable S_0 sets vary from accumulator to accumulator. There are accumulators that support only $S_0 = \emptyset$; others support any polynomial-size S_0, and yet others support any S_0 that can be expressed as a polynomial number of ranges.
– $\mathsf{Add}(\overline{sk}, a_t, \overline{m_t}, x) \rightarrow (a_{t+1}, \overline{m_{t+1}}, w^x_{t+1}, \mathsf{upmsg}_{t+1})$ (for additive and dynamic accumulators) adds the element $x \in D$ to the accumulator, producing the updated accumulator value a_{t+1} and the membership witness w^x_{t+1} for x. It also produces upmsg_{t+1}, which enables witness holders to bring their witnesses up to date, and an updated auxiliary value m_{t+1}.
– $\mathsf{Del}(\overline{sk}, a_t, \overline{m_t}, x) \rightarrow (a_{t+1}, \overline{m_{t+1}}, \overline{u^x_{t+1}}, \mathsf{upmsg}_{t+1})$ (for subtractive and dynamic accumulators) deletes the element $x \in D$ from the accumulator, producing the updated accumulator value a_{t+1} and (optionally) the non-membership witness u^x_{t+1} for x. It also produces an update message upmsg_{t+1}, which enables witness holders to bring their witnesses up to date, and an updated auxiliary value m_{t+1}.

A positive or universal accumulator additionally has the following algorithms:

– $\mathsf{VerMem}(a_t, x, w^x_t) \rightarrow \{0, 1\}$ (executed by any third party) verifies the membership of x in the accumulator using its membership witness w^x_t.
– $\mathsf{MemWitUpOnAdd}(x, w^x_t, \mathsf{upmsg}_{t+1}) \rightarrow w^x_{t+1}$ (executed by a witness holder for additive and dynamic accumulators) updates the membership witness for element x after y is added to the accumulator.

The membership witness update algorithm MemWitUpOnDel (for subtractive and dynamic accumulators) is defined analogously to MemWitUpOnAdd. Notice that the accumulator manager can eliminate the need for these algorithms by sending each witness holder a fresh witness on demand; however, this is unreasonable to ask of the accumulator manager, who might then have to do additional work *per witness holder* per addition or deletion. Instead, the accumulator manager broadcasts a single update message upmsg which witness holders use to bring their witnesses up to date.[8]

If an element x is in S_0, there might be a need to create a membership witness for x independently of the Add algorithm. To this end, there exists a membership witness creation algorithm MemWitCreate. (In this paper, we focus on witness generation during Add, so MemWitCreate will not appear in future sections.)

– $\mathsf{MemWitCreate}(\overline{sk}, \overline{a_t}, \overline{m_t}, x, \overline{(\mathsf{upmsg}_1, \ldots, \mathsf{upmsg}_t)}) \rightarrow w^x_t$ (executed by the accumulator manager or any third party) generates a membership witness w^x_t for x outside the element addition protocol Add.

For negative and universal accumulators, the non-membership witness creation algorithm NonMemWitCreate, the non-membership verification algorithm VerNonMem, and the non-membership witness update algorithms NonMemWitUpOnAdd and NonMemWitUpOnDel are defined analogously to MemWitCreate, VerMem, MemWitUpOnAdd, and MemWitUpOnDel, respectively.

The presence or absence of all of these algorithms is simple to infer from the accumulator type. For convenience, Figure 2 describes the algorithms corresponding to each accumulator type.

[8] The witness holders can also process the update messages in batches if they choose.

	Static			Additive			Subtractive			Dynamic		
	Pos	Neg	Uni	Pos	Neg	Uni	Pos	Neg	Uni	Pos	Neg	Uni
Accumulator Manager Algorithms												
Gen	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Add				✓	✓	✓				✓	✓	✓
Del							✓	✓	✓	✓	✓	✓
MemWitCreate	✓		✓	✓		✓	✓		✓	✓		✓
NonMemWitCreate		✓	✓		✓	✓		✓	✓		✓	✓
Witness Holder Algorithms												
MemWitUpOnAdd				✓		✓				✓		✓
MemWitUpOnDel							✓		✓	✓		✓
NonMemWitUpOnAdd					✓	✓					✓	✓
NonMemWitUpOnDel								✓	✓		✓	✓
Third Party Algorithms												
VerMem	✓		✓	✓		✓	✓		✓	✓		✓
VerNonMem		✓	✓		✓	✓		✓	✓		✓	✓

Fig. 2. Accumulator Algorithms.

2.2 Accumulator Security Properties

In this section, we describe accumulator security properties. Though in Section 2.1 we described many different types of accumulators, here we limit ourselves to positive dynamic accumulators, because our focal constructions Braavos and BraavosB are positive and dynamic. Additionally, for simplicity, we assume that the initial set S_0 is empty. If desired, our definitions can be extended to other initial sets, and other accumulator types, in a straightforward way.

Naturally, accumulators must be both correct and sound. We refer to Reyzin and Yakoubov [35] for formal definitions of correctness. Informally, in a positive accumulator, *correctness* requires that for every element in the accumulator, an honest membership witness holder can always prove membership.

In a positive accumulator, *soundness* (also referred to as *security* or *collision-freeness*) requires that for every element not in the accumulator it is infeasible to prove membership. We present two definitions of soundness for positive dynamic accumulators with $S_0 = \emptyset$ (that is, positive dynamic accumulators which start empty): *adaptive* soundness (Definition 1), and *non-adaptive* soundness (Definition 2). Adaptive soundness is the standard definition of accumulator soundness; it is referred to simply as soundness. We introduce non-adaptive soundness as a building block for our new accumulator, Braavos, as described in Section 3. Informally, in non-adaptive soundness, the adversary must commit to his choice of elements to add in advance, whereas in adaptive soundness the adversary can choose these elements on the fly.

Definition 1. *A positive dynamic accumulator is* adaptively sound *(or simply* sound*) if for all security parameters λ, for all stateful probabilistic polynomial-time adversaries $\mathcal{A}$ with black-box access to* Add *and* Del *oracles (which take elements x) on accumulator a, it holds that:*

$$\Pr\begin{bmatrix} (\overline{sk}, a_0, \overline{m_0}) \leftarrow \mathsf{Gen}(1^\lambda, \emptyset); \\ (x, w) \leftarrow \mathcal{A}^{\mathsf{Add},\mathsf{Del}}(a_0, \overline{m_0}); \\ x \notin S : \\ \mathsf{VerMem}(a, x, w) = 1 \end{bmatrix} \leq \nu(\lambda)\,,$$

where ν is a negligible function in the security parameter, x is an element that is not currently a member ($x \notin S$, where S started out empty and was updated with every call to Add *and* Del*), and a is the accumulator after the adversary made all of his calls to* Add *and* Del*.*

Our definition assumes that inputs to Add and Del are in D and that Del does nothing when called on an element that is not in the accumulated set. Some external mechanisms must ensure that this is indeed the case, or else soundness is not guaranteed. This notion of soundness is sufficient for many scenarios, including anonymity-preserving revocation described in this paper.

We also define a weaker notion of soundness in which the elements are picked in advance (though the ordering of their additions and deletions may still be adaptive).

Definition 2. *A positive dynamic accumulator is* non-adaptively sound *(NA-sound) if the conditions in Definition 1 hold with the following modification: before* Gen *is run, the adversary* $\mathcal{A}$ *produces a set of elements* $x_1, \ldots, x_q \in D$*, and queries to* Add *and* Del *must come from this set.*

This weaker notion of soundness suffices to ensure soundness for randomly chosen elements, because it does not matter when they are chosen.

3 Modular Accumulator Constructions

In this section, we introduce the idea of combining different accumulators to obtain new accumulators with different properties. This technique can lead to the creation of more efficient accumulators, such as the Braavos accumulator described in Section 4 and the BraavosB accumulator described in Appendix B.

3.1 Leveraging Accumulators with Different Functionalities

Notice that, though the notion of a subtractive accumulator helps us draw a more complete mental picture of the accumulator space, there are conceptual equivalences that allow us to ignore subtractive accumulators from hereon out. Let S be the accumulated set of elements, and $\overline{S}$ be the complement of that set ($\overline{S} = D \backslash S$, where D is the domain of all accumulatable elements).[9]

Notice that, conceptually,

- a positive additive accumulator of S is the same as a negative subtractive accumulator of $\overline{S}$.
- a negative additive accumulator of S is the same as a positive subtractive accumulator of $\overline{S}$.
- a universal additive accumulator of S is the same as a universal subtractive accumulator of $\overline{S}$.

Next, we discuss combining simple positive and negative accumulators to obtain universal and dynamic accumulators. An example of a simple positive accumulator is any digital signature scheme; the accumulator value is the verification key, and a membership witness for x is a signature on x [34]. Simple negative accumulators include the Merkle tree construction over ranges [7] and the range-RSA accumulator introduced in Appendix B.1.

For the purpose of this discussion, we assume that all accumulators under consideration have $D = \{0, 1\}^*$ (that is, they can can hold arbitrary elements). We also assume that all of these accumulators are used to accumulate *sets*, not *multi-sets* (that is, an element in the accumulated set is not added again unless it was previously deleted).

Figure 3 gives an illustration of our derivations. The proofs of the correctness and soundness of these constructions are easy exercises, and are therefore omitted. We include only the proof of soundness of construction B given its relevance to one of our constructions(it is used in Appendix B).

Adding Dynamism. We can build a *dynamic positive accumulator* ACC (construction A in Figure 3) out of a positive accumulator $\mathsf{ACC_P}$ and a negative accumulator $\mathsf{ACC_N}$ by adding indexing to the elements. When a new element x is added to ACC, the pair $(x, 1)$ is added to $\mathsf{ACC_P}$. Then, when the element is deleted, the pair $(x, 1)$ is added to $\mathsf{ACC_N}$. Next time x is added, it is added as $(x, 2)$; each time the element is added and deleted, the index is incremented. (Notice that this requires the accumulator manager to maintain an auxiliary storage m the size of which is linear in $|\mathsf{a}|$, where a is the set of all elements ever added.) Proving the membership of x then consists of producing an index i and proving that $((x, i) \in \mathsf{ACC_P}) \wedge ((x, i) \notin \mathsf{ACC_N})$.

Similarly, we can build a *dynamic negative accumulator* ACC (construction C in Figure 3) out of a positive accumulator $\mathsf{ACC_P}$ and a negative accumulator $\mathsf{ACC_N}$. However, the roles of the two accumulators are reversed; when an element x is added to ACC, (x, i) is added to $\mathsf{ACC_N}$ for the appropriate index i. When the element is deleted, (x, i) is added to $\mathsf{ACC_P}$. Proving the non-membership of x then consists of producing an index i and proving that $((x, 1) \notin \mathsf{ACC_N}) \vee (((x, i-1) \in \mathsf{ACC_P}) \wedge ((x, i) \notin \mathsf{ACC_N}))$.

Flipping the Sign of a Dynamic Accumulator. There is an alternative way to build dynamic positive (or negative) accumulators; however, it assumes the existence of a dynamic negative (or positive, respectively) accumulator. We can build a dynamic positive accumulator ACC (construction B in Figure 3) out of a positive accumulator

[9] Note that as long as S is polynomial in size, $\overline{S}$ can be expressed as a polynomial number of ranges.

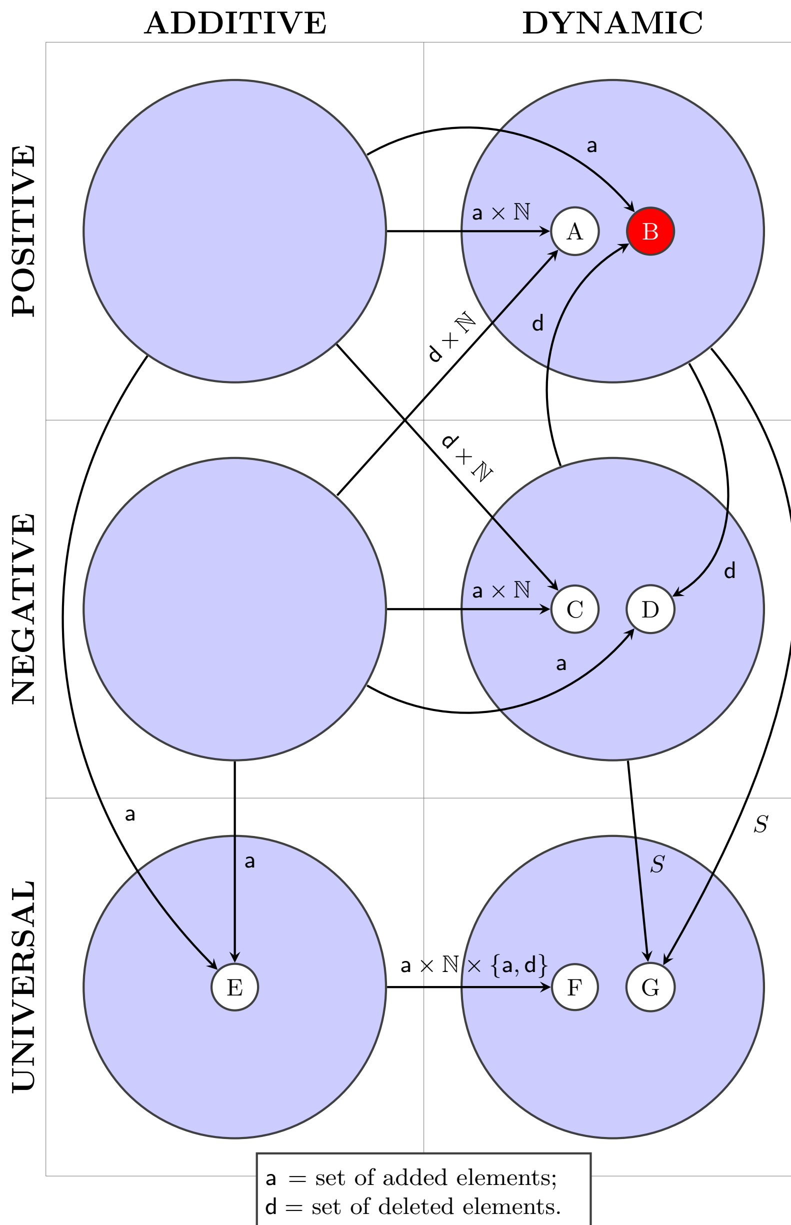

Fig. 3. Modular Accumulator Derivations. Large circles represent a space of accumulator constructions (e.g., the top left-most large circle contains all positive accumulators). Each dot (labeled A-G) within a large circle represents a construction of the type corresponding to the large circle. Arrows denote the modular usage of accumulators of the type corresponding to their start-point to build an accumulator of the type corresponding to their end-point. Arrow labels denote the types of objects being accumulated by the start-point accumulator. S denotes the set of current members, $\mathbb{N}$ the set of natural numbers (used for indexing), and $\{\mathsf{a}, \mathsf{d}\}$ the set of possible actions ('add' or 'delete').

$\mathsf{ACC_P}$ and a dynamic negative accumulator $\mathsf{ACC_N}$ by adding all added elements to $\mathsf{ACC_P}$, and adding all deleted elements to $\mathsf{ACC_N}$. When an element that has previously been deleted is re-added, it is removed from $\mathsf{ACC_N}$. Proving the membership of x consists of proving that $(x \in \mathsf{ACC_P}) \wedge (x \notin \mathsf{ACC_N})$. The construction of a dynamic negative accumulator (construction D in Figure 3) out of a simple negative accumulator and a dynamic positive accumulator mirrors this one, and we will not discuss it further.

Theorem 1. *Construction B is an adaptively sound positive dynamic accumulator if* $\mathsf{ACC_P}$ *is a adaptively sound positive accumulator and* $\mathsf{ACC_N}$ *is an adaptively sound negative dynamic accumulator.*

Proof. The proof consists of a reduction to the adaptive soundness of either $\mathsf{ACC_P}$ or $\mathsf{ACC_N}$. If an adversary produces a witness for an element that is not a member of the accumulated set, then if that element was never added the adversary has succeeded in breaking the adaptive soundness of $\mathsf{ACC_P}$, and if that element was deleted the adversary has succeeded in breaking the adaptive soundness of $\mathsf{ACC_N}$.

Instantiating Construction B. In Appendix B, we describe a communication-optimal instantiation of construction B which we call BraavosB. BraavosB uses digital signatures as the positive accumulator $\mathsf{ACC_P}$, and another new accumulator called range-RSA (described in Appendix B.1) as the dynamic negative accumulator $\mathsf{ACC_N}$. The communication complexity of BraavosB rivals that of Braavos, which is the focal construction of this paper and is described in detail in Section 4. However, though (asymptotically) the communication costs of Braavos and BraavosB are equally small, Braavos has several advantages over BraavosB. One of these is that Braavos supports more efficient zero knowledge proofs of member knowledge. Another is that BraavosB requires the accumulator manager to store an amount of information linear in the number of deleted elements, while Braavos only requires the accumulator manager to store a constant amount of information.

Adding Universality. Additive (construction E in Figure 3) and dynamic (construction G in Figure 3) universal accumulators can be built by combining a positive and negative accumulator of the same type in a straightforward way; both the positive and negative accumulators are used to accumulate the elements in the set.

A dynamic universal accumulator (construction F in Figure 3) can also be built out of an additive universal accumulator $\mathsf{ACC_U}$ in a manner similar to those used to produce constructions A and C. Each element x, when seen for the first time, is assigned a counter $i = 1$. When x is added, a tuple of the form (x, i, a) is added to the additive universal accumulator. When x is deleted, a tuple of the form (x, i, d) is added to the additive universal accumulator, and the counter i is incremented. Proving the membership of x then consists of producing the counter i and proving that $((x, i, \mathsf{a}) \in \mathsf{ACC_U}) \wedge ((x, i, \mathsf{d}) \notin \mathsf{ACC_U})$. Proving the non-membership of x consists of proving that $((x, 1, \mathsf{a}) \notin \mathsf{ACC_U}) \vee (((x, i-1, \mathsf{d}) \in \mathsf{ACC_U}) \wedge ((x, i, \mathsf{a}) \notin \mathsf{ACC_U}))$.

3.2 Leveraging Less Secure Accumulators

In addition to considering combining accumulators with different functionalities, we can consider combining accumulators with different security properties. Given an (adaptively) sound positive additive accumulator $\mathsf{ACC_A}$ and non-adaptively sound (NA-sound) positive dynamic accumulator $\mathsf{ACC_{NA}}$, we can build an adaptively sound dynamic accumulator ACC, as shown in Figure 4.[10] We call this construction "Construction H". Its algorithms are described in Figures 5 through 7. When an element x is added, the accumulator manager selects a random element r from the domain D of $\mathsf{ACC_{NA}}$. She then adds r to $\mathsf{ACC_{NA}}$, and (x, r) to $\mathsf{ACC_A}$. (Recall that random elements can always be safely accumulated in non-adaptively sound accumulators, since those random elements can be chosen without using any information about the accumulator.) When deleting x, the accumulator manager removes r from $\mathsf{ACC_{NA}}$. Proving the membership of x in ACC consists of producing an r and proving that $(r \in \mathsf{ACC_{NA}}) \wedge ((x, r) \in \mathsf{ACC_A})$.

Note that, in order to support deletions, the accumulator manager must store a mapping from every element x to the corresponding r. This can be avoided by having the accumulator manager use a pseudorandom function F_s (where s is the secret pseudorandom function seed) to select an r corresponding to a given x: $r = F_s(x)$. Even though this causes elements added to $\mathsf{ACC_{NA}}$ to be computed rather than chosen at random (therefore seemingly requiring adaptive soundness rather than non-adaptive soundness), non-adaptive soundness is still sufficient because of the indstinguishability of the pseudorandom and random cases.

The correctness of Construction H follows by inspection.

[10] Shamir and Tauman [36] achieve a similar goal of construct chosen message unforgeable signatures from random message unforgeable ones by using a different technique.

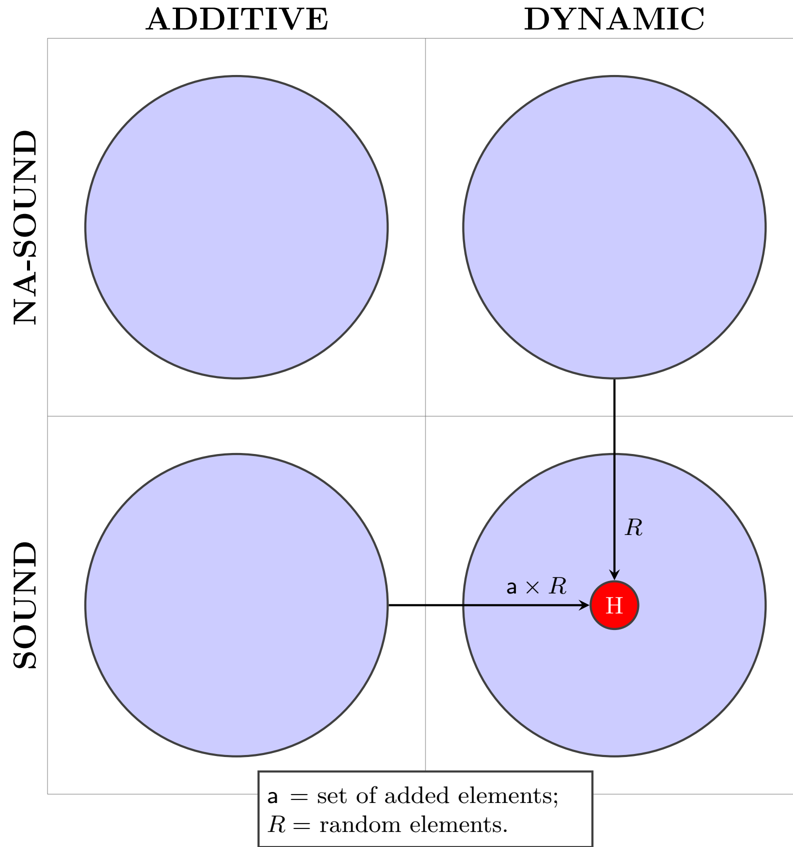

Fig. 4. Modular Accumulator Derivations. All accumulators in this diagram are positive. The notation is the same as that used in Figure 3; additionally, R denotes random elements. A specific (particularly efficient) instantiation of construction H is discussed further in Section 4.

Theorem 2. *Construction H is an adaptively sound positive dynamic accumulator if* $\mathsf{ACC_A}$ *is a positive additive adaptively sound accumulator,* $\mathsf{ACC_{NA}}$ *is a positive dynamic non-adaptively sound accumulator, and* F_s *is a pseudorandom function.*

Proof. The proof consists of a reduction to the adaptive soundness of $\mathsf{ACC_A}$, or the non-adaptive soundness of $\mathsf{ACC_{NA}}$, or the pseudorandomness of F_s. We give only the outline here, because the formal reductions are simple exercises. If an adversary produces a witness for an element that is not a member of the accumulated set, then there are two cases: either the element was never added, or it was deleted. In the first case, the adversary has succeeded in breaking the adaptive soundness of $\mathsf{ACC_A}$. In the second case, the adversary has succeeded in forging a witness for a pseudo-random element in $\mathsf{ACC_{NA}}$. We know that if truly random elements were used in $\mathsf{ACC_{NA}}$, then this would break the non-adaptive soundness of $\mathsf{ACC_{NA}}$, since random elements could have been chosen non-adaptively. Thus, if the adversary is able to forge a witness for $\mathsf{ACC_{NA}}$ with pseudorandom elements, the adversary breaks either the non-adaptive soundness of $\mathsf{ACC_{NA}}$ or the pseudorandomness of F_s.

4 Braavos: A Communication-Optimal Adaptively Sound Dynamic Accumulator

In this section we introduce the Braavos accumulator, which is an instantiation of construction H from Figure 4. Braavos is an adaptively sound positive dynamic accumulator derived from an adaptively sound positive additive accumulator $\mathsf{ACC_A}$ and an non-adaptively sound positive dynamic accumulator $\mathsf{ACC_{NA}}$.

We aim for Braavos to have two properties: communication optimality [6] and efficient zero knowledge proofs, as described in Section 4.4. Our choice of underlying adaptively sound accumulator $\mathsf{ACC_A}$ in Braavos is the

$\mathsf{Gen}(1^\lambda, \emptyset)$:

1. $(\mathsf{ACC_A}.a, \mathsf{ACC_A}.m, \mathsf{ACC_A}.sk) \leftarrow \mathsf{ACC_A}.\mathsf{Gen}(1^\lambda, \emptyset)$
2. $(\mathsf{ACC_{NA}}.a, \mathsf{ACC_{NA}}.m, \mathsf{ACC_{NA}}.sk) \leftarrow \mathsf{ACC_{NA}}.\mathsf{Gen}(1^\lambda, \emptyset)$
3. Let $\{F_s\}_{s\in\{0,1\}^\lambda}$ be a pseudorandom function family whose range is the domain of $\mathsf{ACC_{NA}}$; pick $s \in \{0,1\}^\lambda$ at random.
4. Let $sk = (\mathsf{ACC_A}.sk, \mathsf{ACC_{NA}}.sk, s)$.
5. Let $a = (\mathsf{ACC_A}.a, \mathsf{ACC_{NA}}.a)$.
6. Let $m = (\mathsf{ACC_A}.m, \mathsf{ACC_{NA}}.m)$.
7. Return $(sk,\ a,\ m)$.

$\mathsf{Add}(sk, a, m, x)$:

1. Set $r = F_s(x)$.
2. $(\mathsf{ACC_A}.a, \mathsf{ACC_A}.m, \mathsf{ACC_A}.w, \mathsf{ACC_A}.\mathsf{upmsg}) \leftarrow \mathsf{ACC_A}.\mathsf{Add}(\mathsf{ACC_A}.sk, \mathsf{ACC_A}.a, \mathsf{ACC_A}.m, (x, r))$.
3. $(\mathsf{ACC_{NA}}.a, \mathsf{ACC_{NA}}.m, \mathsf{ACC_{NA}}.w, \mathsf{ACC_{NA}}.\mathsf{upmsg}) \leftarrow \mathsf{ACC_{NA}}.\mathsf{Add}(\mathsf{ACC_{NA}}.sk, \mathsf{ACC_{NA}}.a, \mathsf{ACC_{NA}}.m, r)$.
4. Let $w = (\mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w)$.
5. Let $\mathsf{upmsg} = (\mathsf{ACC_A}.\mathsf{upmsg}, \mathsf{ACC_{NA}}.\mathsf{upmsg})$.
6. Return $(a, m, w, \mathsf{upmsg})$.

$\mathsf{Del}(sk, a, m, x)$:

1. Set $r = F_s(x)$.
2. Let $(\mathsf{ACC_{NA}}.a, \mathsf{ACC_{NA}}.m, \mathsf{ACC_{NA}}.\mathsf{upmsg}) \leftarrow \mathsf{ACC_{NA}}.\mathsf{Del}(\mathsf{ACC_{NA}}.sk, \mathsf{ACC_{NA}}.a, \mathsf{ACC_{NA}}.m, r)$.
3. Return $(a, m, \mathsf{ACC_{NA}}.\mathsf{upmsg})$.

Fig. 5. Construction H from Figure 4 accumulator manager algorithms (Gen, Add, Del and MemWitCreate), in terms of the underlying positive adaptively sound accumulator $\mathsf{ACC_A}$ and positive dynamic non-adaptively sound accumulator $\mathsf{ACC_{NA}}$. In all of these algorithms, we assume that $sk = (\mathsf{ACC_A}.sk, \mathsf{ACC_{NA}}.sk, s)$, $a = (\mathsf{ACC_A}.a, \mathsf{ACC_{NA}}.a)$, and $m = (\mathsf{ACC_A}.m, \mathsf{ACC_{NA}}.m)$.

$\mathsf{MemWitUpOnAdd}(a, x, w, \mathsf{upmsg})$:

1. Parse $(\mathsf{ACC_A}.\mathsf{upmsg}, \mathsf{ACC_{NA}}.\mathsf{upmsg}) = \mathsf{upmsg}$.
2. Parse $(\mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w) = w$.
3. $\mathsf{ACC_A}.w \leftarrow \mathsf{ACC_A}.\mathsf{MemWitUpOnAdd}(x, \mathsf{ACC_A}.w, \mathsf{ACC_A}.\mathsf{upmsg})$.
4. $\mathsf{ACC_{NA}}.w \leftarrow \mathsf{ACC_{NA}}.\mathsf{MemWitUpOnAdd}(x, \mathsf{ACC_{NA}}.w, \mathsf{ACC_{NA}}.\mathsf{upmsg})$.
5. Return $w = (\mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w)$.

$\mathsf{MemWitUpOnDel}(a, x, w, \mathsf{upmsg})$:

1. Parse $(\mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w) = w$.
2. $\mathsf{ACC_{NA}}.w \leftarrow \mathsf{ACC_{NA}}.\mathsf{MemWitUpOnAdd}(x, \mathsf{ACC_{NA}}.w, \mathsf{upmsg})$.
3. Return $w = (\mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w)$.

The witness holder can run BatchMemWitUpOnDel immediately before producing a proof.

Fig. 6. Construction H from Figure 4 witness holder algorithms (MemWitUpOnAdd and MemWitUpOnDel), in terms of the underlying positive adaptively sound accumulator $\mathsf{ACC_A}$ and positive dynamic non-adaptively sound accumulator $\mathsf{ACC_{NA}}$. Note that when Construction H is instantiated with a digital signature scheme as the positive accumulator $\mathsf{ACC_A}$, there is no need for a MemWitUpOnAdd algorithm.

$\mathsf{VerMem}(a, x, w)$:

1. Parse $(\mathsf{ACC_A}.a, \mathsf{ACC_{NA}}.a) = a$.
2. Parse $(\mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w) = w$.
3. Let $b_1 \leftarrow \mathsf{ACC_A}.\mathsf{VerMem}(\mathsf{ACC_A}.a, x, \mathsf{ACC_A}.w)$.
4. Let $b_2 \leftarrow \mathsf{ACC_{NA}}.\mathsf{VerMem}(\mathsf{ACC_{NA}}.a, x, \mathsf{ACC_{NA}}.w)$.
5. Return 1 if $b_1 = b_2 = 1$, and return 0 otherwise.

Fig. 7. Construction H from Figure 4 third party algorithms (VerMem), in terms of the underlying positive adaptively sound accumulator $\mathsf{ACC_A}$ and positive dynamic non-adaptively sound accumulator $\mathsf{ACC_{NA}}$.

CL signature scheme [12], because it supports efficient zero knowledge proofs of knowledge of a signature on a committed value. Note that though construction H has a MemWitUpOnAdd algorithm (described in Figure 6), this algorithm is not used by Braavos, since signatures do not require witness updates when additions take place.

The challenge that remains is finding a communication-optimal, dynamic, non-adaptively sound accumulator $\mathsf{ACC_{NA}}$. $\mathsf{ACC_{NA}}$ should only require membership witness updates upon element deletions, not element additions. In Section 4.1, we describe CL-RSA-B, which is exactly such an accumulator.

4.1 CL-RSA-B: A Communication-Optimal Non-Adaptively Sound Dynamic Accumulator

In this section, we formally describe the CL-RSA-B accumulator, which was informally introduced by Camenisch and Lysyanskaya [11] in a remark on page 12. The CL-RSA-B accumulator is similar to the standard RSA accumulator [11], which evolves the accumulator value (as well as all membership witnesses) with every addition and deletion. The CL-RSA-B accumulator, unlike the RSA accumulator, evolves the accumulator value with every deletion only. However, the price is that, as far as we can tell, the CL-RSA-B accumulator is only non-adaptively sound.

The RSA Accumulator. In order to understand the CL-RSA-B accumulator, it helps to understand the RSA accumulator first. Its value is a quadratic residue a modulo n, where n is an RSA modulus: $n = pq$, where $p = 2p' + 1$ and $q = 2q' + 1$ for prime p, p', q, and q'. The domain D of the RSA accumulator consists of all odd positive prime integers x.[11]

During the addition of x to the accumulator, the new accumulator value is computed as $a_{t+1} = a_t^x \bmod n$. The membership witness w for x is then defined to be the old accumulator value a_t. A membership verification consists of checking that $w^x \equiv a \bmod n$. When another element y is added to the accumulator, the membership witness for x is updated by taking $w_{t+1} = w_t^y \bmod n$.

When an element y is deleted, the accumulator manager (who knows the trapdoor $p'q'$) computes the new accumulator as $a_{t+1} = a_t^{y^{-1} \bmod p'q'} \bmod n$. The membership witness w for x can then be updated using the Bezout coefficients b and c such that $bx + cy = 1$. (Recall that the domain D of the accumulator contains only odd prime numbers, so such b and c are guaranteed to exist.) The new witness is computed as $w_{t+1} = w_t^c a_{t+1}^b \bmod n$.

The CL-RSA-B *Accumulator.* The CL-RSA-B accumulator preserves the relationship between the accumulator value and the witnesses, but avoids computing a new accumulator value and updating witnesses during each addition. Instead, during the addition of odd prime x the accumulator manager keeps the accumulator constant, and computes the membership witness w for x as $w = a^{x^{-1} \bmod p'q'} \bmod n$. Notice that this eliminates the need for updating existing membership witnesses during additions. The process for proving membership and for deletions is the same as in the RSA accumulator. The algorithms of the CL-RSA-B accumulator are detailed in Figure 8.

CL-RSA-B *Soundness.* The RSA accumulator is adaptively sound, meaning that an adversary cannot find a membership witness for an element that is not a member even if she chooses which elements should be added, optionally based on accumulator and witness values she has previously seen.

The CL-RSA-B accumulator is non-adaptively sound, meaning that an adversary cannot find a membership witness for an element that is not a member if she chooses all elements to add prior to seeing any accumulator information. In particular, the CL-RSA-B accumulator is sound when only random elements are added to the accumulator, since those can be chosen prior to seeing any accumulator or witness values.[12] This holds under the *strong RSA assumption* [2].

Assumption 1 (Strong RSA) *For any probabilistic polynomial-time adversary $\mathcal{A}$,*

$$\Pr[p, q \leftarrow \{\lambda\text{-}bit\ safe\ primes\}; n = pq; t \leftarrow \mathbb{Z}_n^*; (r, e) \leftarrow \mathcal{A}(n, t) : r^e = t \bmod n \wedge e\ is\ prime] = \nu(\lambda)$$

For some negligible function ν.

[11] Note that p' or q' cannot themselves be accumulated, since $(p')^{-1} \bmod p'q'$ and $(q')^{-1} \bmod p'q'$ do not exist; however, that only happens with negligible probability in the adaptive soundness game in Definition 1, since if the adversary finds p' or q', he or she has succeeded in factoring n.

[12] We are not certain whether CL-RSA-B is also adaptively sound. Proving that it is or is not is an open problem. It is adaptively sound when the when a polynomial-size subset of D is used as the domain; however, this is a very limiting restriction.

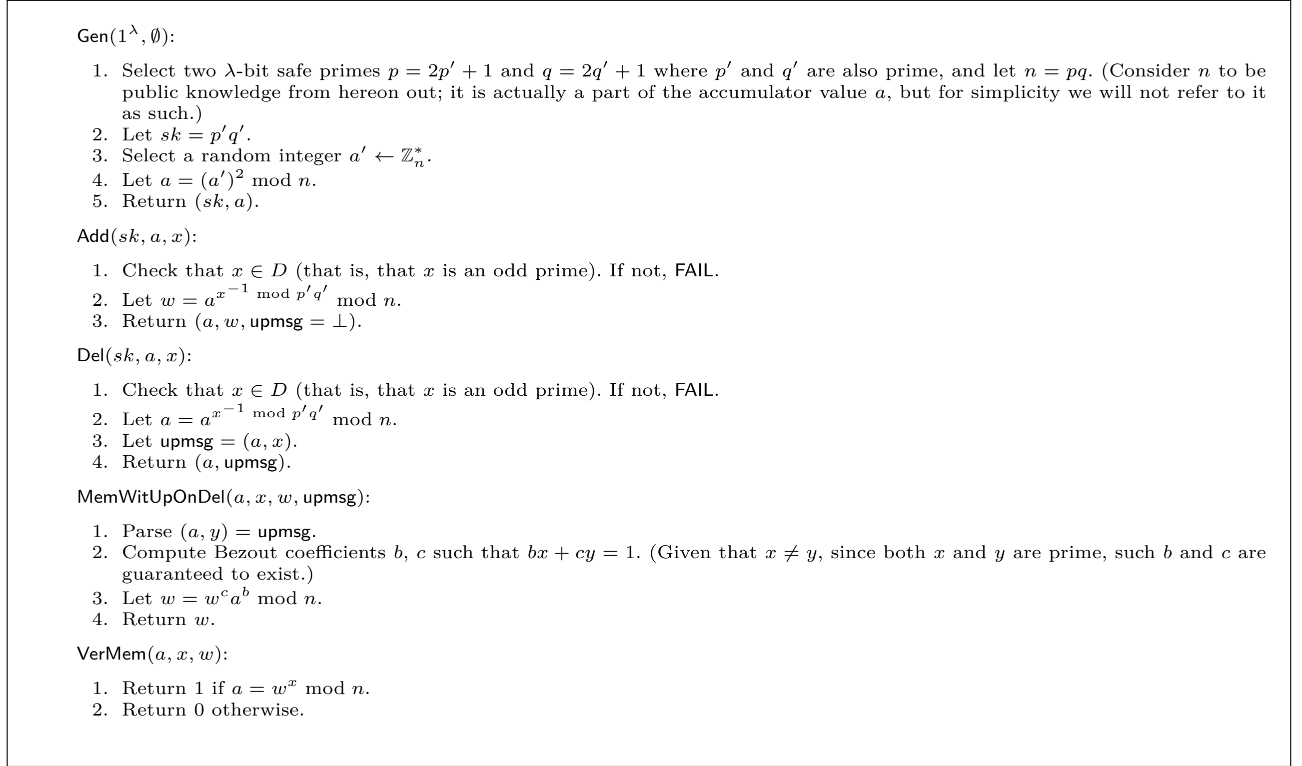

$\mathsf{Gen}(1^\lambda, \emptyset)$:

1. Select two λ-bit safe primes $p = 2p' + 1$ and $q = 2q' + 1$ where p' and q' are also prime, and let $n = pq$. (Consider n to be public knowledge from hereon out; it is actually a part of the accumulator value a, but for simplicity we will not refer to it as such.)
2. Let $sk = p'q'$.
3. Select a random integer $a' \leftarrow \mathbb{Z}_n^*$.
4. Let $a = (a')^2 \bmod n$.
5. Return (sk, a).

$\mathsf{Add}(sk, a, x)$:

1. Check that $x \in D$ (that is, that x is an odd prime). If not, FAIL.
2. Let $w = a^{x^{-1} \bmod p'q'} \bmod n$.
3. Return $(a, w, \mathsf{upmsg} = \bot)$.

$\mathsf{Del}(sk, a, x)$:

1. Check that $x \in D$ (that is, that x is an odd prime). If not, FAIL.
2. Let $a = a^{x^{-1} \bmod p'q'} \bmod n$.
3. Let $\mathsf{upmsg} = (a, x)$.
4. Return (a, upmsg).

$\mathsf{MemWitUpOnDel}(a, x, w, \mathsf{upmsg})$:

1. Parse $(a, y) = \mathsf{upmsg}$.
2. Compute Bezout coefficients b, c such that $bx + cy = 1$. (Given that $x \neq y$, since both x and y are prime, such b and c are guaranteed to exist.)
3. Let $w = w^c a^b \bmod n$.
4. Return w.

$\mathsf{VerMem}(a, x, w)$:

1. Return 1 if $a = w^x \bmod n$.
2. Return 0 otherwise.

Fig. 8. CL-RSA-B algorithms.

Theorem 3. *The* CL-RSA-B *accumulator with a domain D consisting of odd primes is non-adaptively sound under the strong RSA assumption.*

Proof. In Figure 9, we reduce the non-adaptive soundness of the CL-RSA-B accumulator to the strong RSA assumption. The reduction $\mathcal{R}$ takes in an RSA integer n and a random value $t \in \mathbb{Z}_n^*$, and returns r, e such that $t = r^e \bmod n$. $\mathcal{R}$ leverages an adversary $\mathcal{A}$ which can break the non-adaptive soundness of the CL-RSA-B accumulator; that is, after making addition (Add) and deletion (Del) queries on elements chosen before seeing the initial state of the accumulator, $\mathcal{A}$ can produce an odd prime x and a witness w such that $a = w^x \bmod n$, and x is not in the accumulator.

$\mathcal{R}$ must be able to answer two types of queries from $\mathcal{A}$: Add queries on the non-adaptively chosen elements, and Del queries on the same elements. Let q_{Add} be an upper bound on the number of Add queries, and q_{Del} be an upper bound on the number of Del queries $\mathcal{A}$ can make. During the setup phase, having received the elements $x_1, \ldots, x_{q_{\mathsf{Add}}}$ from $\mathcal{A}$, the reduction $\mathcal{R}$ creates an accumulator for which it can answer Add and Del queries on elements $x_1, \ldots, x_{q_{\mathsf{Add}}}$. It does so by starting with $a = t^2 \bmod n$, and raising a to the power of the elements. By raising a to the power of $x_j^{q_{\mathsf{Del}}}$, $\mathcal{R}$ creates an accumulator value for which it is able to answer Del and Add queries on x_j even if $\mathcal{A}$ spends all of its Del queries on that one element. However, if $\mathcal{A}$ forges a witness w for x_j (after having added and deleted it fewer than q_{Del} times), the reduction won't be able to use w to break the strong RSA assumption, since it already knows w! For that reason, $\mathcal{R}$ guesses a "target" element x_j from among $x_1, \ldots, x_{q_{\mathsf{Add}}}$, and the number e_j of times that x_j will be added and deleted before the forgery (which can be anywhere from 0 to q_{Del}), and only raises t to the power of $x_j^{e_j}$, not $x_j^{q_{\mathsf{Del}}}$. Figure 9 shows the details of how the reduction picks an accumulator value based on $x_1, \ldots, x_{q_{\mathsf{Add}}}$, how it answers Add and Del queries, and how it then uses the output of $\mathcal{A}$ to break the strong RSA assumption.

This reduction succeeds as long as:

1. During the query phase, $\mathcal{R}$ does not output FAIL. $\mathcal{R}$ does not output FAIL if the target exponent e_i was chosen correctly, which happens with probability $\frac{1}{q_{\mathsf{Del}}+1}$.

$\mathsf{Setup}(n, t, q_{\mathsf{Add}}, q_{\mathsf{Del}})$:

1. Let $x_1, \ldots, x_{q_{\mathsf{Add}}}$ be the distinct odd primes provided by the adversary $\mathcal{A}$.
2. Let $a = t^2$ (so as to make a a quadratic residue).
3. Let $e_c = 2$ be the current exponent linking t to a. (So, $a = t^{e_c} \bmod n$ is an invariant.)
4. Pick a random index $j \leftarrow \{1, \ldots, q_{\mathsf{Add}}\}$.
5. For $i \in [1, \ldots, q_{\mathsf{Add}}]$:
   (a) If $i = j$: pick a random $e_i \leftarrow \{0, \ldots, q_{\mathsf{Del}}\}$.
   (b) Else: $e_i = q_{\mathsf{Del}}$.
   (c) Let $a = a^{(x_i^{e_i})} \bmod n$.
   (d) Let $e_c = e_c x_i^{e_i}$.
6. Return $a = e_i^{e_c} \bmod n$ to the adversary $\mathcal{A}$.

$\mathsf{Add}(x_i)$:

1. If $e_i = 0$: FAIL.
2. Let $w = t^{e_c/x_i} \bmod n$. (Note that e_c must be divisible by x_i, since e_c has a factor of $x_i^{e_i}$, and $e_i > 0$.)
3. Return w.

$\mathsf{Del}(x_i)$:

1. If $e_i = 0$: FAIL.
2. Let $e_c = e_c/x_i$. (Note that e_c must be divisible by x_i, since e_c has a factor of $x_i^{e_i}$, and $e_i > 0$.)
3. Let $e_i = e_i - 1$.
4. Let $a = t^{e_c} \bmod n$.
5. Let $\mathsf{upmsg} = (a, x_i)$.
6. Return (a, upmsg).

$\mathsf{Output}(e, w)$:

1. Check that $a = w^e \bmod n$. If not, FAIL.
2. If $e = x_i$ for $i \in \{1, \ldots, q_{\mathsf{Add}}\}$ and ($i \neq j$ or $e_i > 0$): FAIL.
3. We know that e and e_c must be relatively prime; e_c has no factors outside of $x_1, \ldots, x_{q_{\mathsf{Add}}}$, and those factors have powers e_i. So, $\mathcal{R}$ can compute Bezout coefficients b, c such that $be + ce_c = 1$.
4. Let $r = t^b w^c \bmod n$.
   (Let $y = t^{e^{-1} \bmod p'q'} \bmod n$; equivalently, $y^e \bmod n = t$. Since $w^e \bmod n = a$ and $t^{e_c} \bmod n = (y^e)^{e_c} \bmod n = a$, it follows that $w = y^{e_c} \bmod n$. So, $r = t^b w^c = (y^e)^b (y^{e_c})^c = y^1 = y$.)
5. Return (r, e).

Fig. 9. Reduction $\mathcal{R}$ from the non-adaptive soundness of CL-RSA-B to the strong RSA assumption.

2. During the output phase, $\mathcal{R}$ does not output FAIL. If $\mathcal{A}$ outputs a witness for an element $x_i \in \{x_1, \ldots, x_{q_{\mathsf{Add}}}\}$, $\mathcal{R}$ does not output FAIL as long as:
   (a) $\mathcal{R}$ makes x the "target" prime (that is, $j = i$). This happens with probability $\frac{1}{q_{\mathsf{Add}}}$.
   (b) $\mathcal{R}$ correctly chooses the target exponent e_i for x. This happens with probability $\frac{1}{q_{\mathsf{Del}}+1}$. However, this is already accounted for in item 1.
3. $\mathcal{A}$ succeeds in breaking the security of the CL-RSA-B accumulator, which we assume happens with non-negligible probability ϵ.

As long as $\mathcal{R}$ does not output FAIL, $\mathcal{A}$ sees the same transcript it would when interacting with a real accumulator manager. The probability of the reduction $\mathcal{R}$ succeeding is $\frac{1}{q_{\mathsf{Add}}} \frac{1}{q_{\mathsf{Del}}+1} \epsilon$, which is non-negligible.

Other Approaches for Expanding the Domain and Getting Adaptive Soundness for CL-RSA-B. The domain D of CL-RSA-B consists of odd primes. Such a limited domain is not a problem for our main application, because the Braavos accumulator manager can choose a (psuedo)random prime r when a new element is added to the accumulator for the first time, as described in Construction H. In fact, Construction H can be viewed as one approach to expanding the domain of CL-RSA-B and obtaining adaptive security for it. Here we briefly mention other approaches. Let D' be the desired domain. Let f be a mapping from D' to λ-bit odd primes. To add $x \in D'$ to the accumulator, add $f(x)$ instead. We can obtain adaptive soundness in the following ways:

- We can model f as a random oracle (the proof is straightforward).
- We can avoid the random oracle by making a different strong assumption instead: namely, the assumption that f is collision-resistant, and the very strong "adaptive strong-RSA assumption". Informally, the adaptive strong-RSA assumption states that even given an oracle that can take roots modulo n, it is difficult to find new roots whose power is relatively prime to those of the roots produced by the oracle.

- We can get somewhat better assumptions by having f be a randomized mapping, and include the randomness R as part of the witness. Then, assuming that for every two elements x_1 and x_2, the distributions $f(x_1; R)$ and $f(x_2; R)$ (over random choices of R) are statistically close, we can use the technique from [23]. To do so, we need to assume that the strong RSA assumption (Assumption 1) also holds in a model where there exists an oracle $\mathcal{O}$ that on input x, p returns a random R such that $f(R; x) = p$.
- Alternatively, we can use the strong-RSA assumption without modification if f is a trapdoor hash function, following the technique of [36].

All of these approaches require f that maps to primes. A way to build such f is described in [5, Section 3.2] (see also [30, Section 7]).

4.2 Braavos Soundness

The Braavos accumulator uses Camenisch-Lysyanskaya (CL) signatures [12] as the underlying positive accumulator ACC_A, and the CL-RSA-B accumulator as the underlying dynamic positive non-adaptively sound accumulator ACC_NA. CL signatures are existentially unforgeable under the strong RSA assumption. Recall that according to Theorem 3, the CL-RSA-B accumulator is non-adaptively sound under the same assumption. By Theorem 2, this implies that the Braavos accumulator is an adaptively sound positive dynamic accumulator under the strong RSA assumption.

Properties other than adaptive soundness (such as correctness) are self-evident.

4.3 Comparison with Other Constructions

The Braavos accumulator is a positive, dynamic accumulator with efficient (constant-time) membership witness generation, and no membership witness updates upon element additions — only upon element deletions. In particular, for a fixed security parameter λ, Braavos achieves the total communication lower bound shown by Camacho [6]. (Total communication refers to the sum of the sizes of all upmsg messages sent by the accumulator manager to the witness holders after $|\mathsf{a}|$ additions and $|\mathsf{d}|$ deletions.) In Appendix A, we prove that adding universality would necessarily degrade the total communication of Braavos.

In Figure 10, we compare Braavos (and BraavosB, described in Appendix B) to prior constructions in terms of the properties introduced in Section 2. We compare them to digital signatures, and to the three other primary lines of work on accumulators: the RSA construction [3,11,26], the bilinear map constructions [33,18,1], and the Merkle tree constructions [7].[15] In our comparison we also include two building blocks: the CL-RSA-B accumulator (used in Braavos, and described in Section 4.1), and the range-RSA accumulator (used in BraavosB, and described in Appendix B.1).

As mentioned earlier, though the properties and asymptotic costs of Braavos and BraavosB appear similar, Braavos has several advantages over BraavosB. One of these is more efficient zero knowledge proofs of member knowledge; another is that the accumulator manager storage in Braavos is only $O(1)$, whereas in BraavosB it is $O(|\mathsf{d}|)$.

Though Figure 10 includes some of the most well known accumulator constructions to compare with Braavos and BraavosB, we would like to note that there exists a large number of other dynamic accumulator constructions in the literature [26,18,9,1,17,20]. To the best of our knowledge, these constructions do not achieve the efficiency we aim for.

Non-accumulator constructions. Finally, although cryptographic accumulators are the most prominent way to achieve anonymous revocation, other methods have been suggested in the literature. One such method is based on the Naor-Naor-Lotspiech (NNL) broadcast encryption scheme [32]. Libert, Peters and Yung [27] adapt the NNL method to achieve anonymous revocation for group signatures. At a high level, the NNL algorithm works as

[13] Here range-RSA is presented as a negative accumulator, because that is how it is used in the BraavosB accumulator. However, range-RSA can be easily modified to be universal.

[14] $|\mathsf{a}|$ and $|\mathsf{d}|$ refer to the number of elements added and deleted *after* the addition of the element whose witness updates are being discussed.

[15] For those interested in a more concrete comparison, Lapon et al. [25] provide concrete running time measurements of bilinear map accumulator constructions and the RSA construction. We discuss the concrete running times of Braavos in Section 5.2.

Protocol Runtimes								
Accumulator	**Sigs**	**RSA**	**BM**	**Merkle**	***range-RSA*** [13]	BraavosB	CL-RSA-B	Braavos
Add	1	1	1	$\log \lvert a \rvert$	1	1	1	1
Del	–	1	1	$\log \lvert a \rvert$	1	1	1	1
NonMemWitCreate	–	$\lvert S \rvert$	$\lvert S \rvert$	$\log \lvert a \rvert$	1	–	–	–
MemWitUpOnAdd	0	1	1	$\log \lvert a \rvert$	–	0	0	0
MemWitUpOnDel	–	1	1	$\log \lvert a \rvert$	–	1	1	1
NonMemWitUpOnAdd	–	1	1	$\log \lvert a \rvert$	1	–	–	–
NonMemWitUpOnDel	–	1	1	$\log \lvert a \rvert$	1	–	–	–
VerMem	1	1	1	$\log \lvert a \rvert$	–	1	1	1
VerNonMem	–	1	1	$\log \lvert a \rvert$	1	–	–	–
Storage								
Accumulator size	1	1	1	1	1	1	1	1
Witness size	1	1	1	$\log \lvert a \rvert$	1	1	1	1
Manager storage ($\lvert m \rvert$)	1	$\lvert S \rvert$	$\lvert S \rvert$	$\lvert a \rvert$	$\lvert S \rvert$	$\lvert d \rvert$	1	1
Properties								
Additive?	✓	✓	✓	✓	✓	✓	✓	✓
Subtractive?		✓	✓	✓	✓	✓	✓	✓
Positive?	✓	✓	✓	✓		✓	✓	✓
Negative?		✓	✓	✓	✓			
Total comm. to Verifier [14]	0	$\lvert a \rvert + \lvert d \rvert$	$\lvert a \rvert + \lvert d \rvert$	$(\lvert a \rvert + \lvert d \rvert) \log \lvert a \rvert$	$\lvert a \rvert + \lvert d \rvert$	$\lvert d \rvert$	$\lvert d \rvert$	$\lvert d \rvert$
Total comm. to Member	0	$\lvert a \rvert + \lvert d \rvert$	$\lvert a \rvert + \lvert d \rvert$	$(\lvert a \rvert + \lvert d \rvert) \log \lvert a \rvert$	$\lvert a \rvert + \lvert d \rvert$	$\lvert d \rvert$	$\lvert d \rvert$	$\lvert d \rvert$
Efficient ZKPs?	✓	✓	✓		✓	✓	✓	✓
Adaptively sound?	✓	✓	✓	✓	✓	✓		✓

Fig. 10. Various Accumulators and their Protocol Runtimes, Storage Requirements, and Properties. We let $|a|$ denote the number of elements added to the accumulator, $|d|$ denote the number of elements deleted from the accumulator, and $|S|$ denote the total number of member elements in the accumulator. (Note that $|S|$ is $|a| - |d|$.) The Braavos and BraavosB accumulators are the first adaptively sound dynamic (additive and subtractive) accumulators to have the optimal total communication of $O(|d|)$. Sigs represents any digital signature scheme. The RSA Construction is due to [3,11,26]. The BM (bilinear map) construction is due to [33,18,1]. The Merkle tree construction is due to [7]. Range-RSA and BraavosB are described in Appendix B. CL-RSA-B and Braavos were described earlier in this section. A logarithmic factor is omitted everywhere; it is implicit as the size of our elements. Big-O notation is omitted from this table in the interest of brevity.

follows: it uses a binary tree with $2^{\ell-1}$ leaves which represent the members of the group. If r leaves are revoked, the NLL algorithm partitions the set of unrevoked leaves into m disjoint sets. Each such set can be represented using nodes (a, b) of the tree, such that b is a descendant of a and a leaf i is in the partition that corresponds to nodes (a, b) if i is a descendant of a but not a descendant of b.

In the bilinear map setting, in the standard model, Libert, Peters and Yung construct a revocable group signature scheme in which the group manager organizes the unrevoked users according to the NNL partition, and a valid user, as part of her group signature, demonstrates that she is in a valid partition, i.e. that her membership certificate i corresponds to a valid NNL subset. Their main tools are vector-commitments and proof friendly signature schemes. In contrast to the accumulator-based solution, the amount of work each user has to do is independent of the total number of users revoked. This comes at a price: (a) the group manager will have to perform $O(r)$ work in order to publish the updated revocation tree in each time period, where r is the number of users revoked so far (note that in Braavos the accumulator manager only performs constant work); (b) their solution is tailored for the bilinear-map setting; finally, (c) it is not clear how to extend the their techniques from revocable group signatures to revocable anonymous credentials.

4.4 Adding Zero Knowledge to Braavos

So far, we have only discussed the functionality of accumulators, ignoring potential privacy concerns. There typically exist three primary privacy goals in the context of accumulators: hiding the membership (or non-membership) witness, hiding the element whose membership (or non-membership) is being demonstrated as well as the witness, and hiding all information about the accumulated set [24]. For our application of anonymous credential revocation (discussed in Section 5), we mostly care about zero knowledge proofs of member knowledge, which hide not only the witness, but the member element itself.

The Braavos accumulator supports efficient zero-knowledge proofs of member knowledge. Given that Braavos is composed of two accumulators $\mathsf{ACC_A}$ and $\mathsf{ACC_{NA}}$, in order for a witness holder to produce a zero-knowledge proof of member knowledge in Braavos, she would have to produce a conjunction of proofs of member knowledge in both $\mathsf{ACC_A}$ and $\mathsf{ACC_{NA}}$ and a proof that those members have the correct relationship. More concretely, she would have to compute the following zero-knowledge proof (described using Camenisch-Stadler [15] notation):

$$\begin{aligned}&\mathsf{ZKP}[(x, r, \mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w) :\\&\quad \wedge\, \mathsf{ACC_A}.\mathsf{VerMem}(\mathsf{ACC_A}.a, (x, r), \mathsf{ACC_A}.w)\\&\quad \wedge\, \mathsf{ACC_{NA}}.\mathsf{VerMem}(\mathsf{ACC_{NA}}.a, r, \mathsf{ACC_{NA}}.w)\\&](\mathsf{ACC_{NA}}.a, \mathsf{ACC_A}.a)\end{aligned}$$

Where $\mathsf{ACC_A}$ is the signature scheme $\mathsf{SIG}_{CL} = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ due to Camenisch and Lysyanskaya [12], and $\mathsf{ACC_{NA}}$ is the CL-RSA-B accumulator.

For integration into larger systems, it might be important to be able to link the witnesses used in the proof to other statements, while still keeping the elements and witnesses private. To this end, commitments to the witnesses can be used. Let $\mathsf{Com} = (\mathsf{Commit}, \mathsf{Verify})$ be a commitment scheme; to integrate commitments into the zero knowledge proof, a witness holder computes commitments to the membership witnesses $\mathsf{ACC_A}.w$ and $\mathsf{ACC_{NA}}.w$: $(C_1, o_1) = \mathsf{Com.Commit}(\mathsf{ACC_A}.w)$ and $(C_2, o_2) = \mathsf{Com.Commit}(\mathsf{ACC_{NA}}.w)$, where o_1 and o_2 are decommitment values. The proof is then enhanced, as follows:

$$\begin{aligned}&\mathsf{ZKP}[(x, r, \mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w, o_1, o_2) :\\&\quad \mathsf{Com.Verify}(C_1, \mathsf{ACC_A}.w, o_1)\\&\quad \wedge\, \mathsf{Com.Verify}(C_2, \mathsf{ACC_{NA}}.w, o_2)\\&\quad \wedge\, \mathsf{ACC_A}.\mathsf{VerMem}(\mathsf{ACC_{NA}}.a, (x, r), \mathsf{ACC_A}.w)\\&\quad \wedge\, \mathsf{ACC_{NA}}.\mathsf{VerMem}(\mathsf{ACC_{NA}}.a, r, \mathsf{ACC_{NA}}.w)\\&](\mathsf{ACC_{NA}}.a, \mathsf{ACC_A}.a, C_1, C_2)\end{aligned}$$

For concrete descriptions of the individual clauses of this proof using the commitment scheme due to Fujisaki and Okamoto [22], please refer to Fujisaki and Okamoto [22] and Camenisch and Lysyanskaya [12,11].

5 Anonymous Revocation from Accumulators

In this section, we show how accumulators can be used in practice. Accumulators combined with zero knowledge proofs are a perfect solution for providing revocation in a system where preserving users' privacy is crucial.

As an example, consider an anonymous credential system where transactions involving the same credential need to be unlinkable. An anonymous credential system is comprised of users, issuers, and verifiers. An issuer certifies a user's attributes in the form of a credential. To authenticate a user, a verifier first sends her a presentation policy that describes which statements she should prove about her credentials. Based on the policy, the user derives a fresh unlinkable proof (or *token*) from her credentials and sends it to the verifier. The verifier then determines whether the token is valid with respect to the policy.

To make credentials *revocable*, the user needs to be able to prove that the credential, on which the token is based, was not revoked. This must be done in a privacy-preserving fashion, i.e., without destroying unlinkability. Camenisch et al. [10] describe a generic revocation component, which can be added to any anonymous system, including anonymous credentials and group signature schemes. We refer to this component as the anonymous revocation component (ARC).

An ARC requires an additional entity called a revocation authority (RA). The RA assists the issuer with adding new users to the system, maintains the necessary revocation information, and changes the revocation status of any user in the system. (The role of the RA can optionally be played by an issuer or a verifier.) Camenisch et al. [10] describe the necessary interfaces and definitions for an ARC, and show how to instantiate it with the revocation scheme of Nakanishi et al. [31].

In Section 5.1, we show how to instantiate an ARC with accumulators and zero knowledge proofs, and extend the definition of an ARC to include the addition of users after the initial setup and the re-addition of users after they have been revoked. In Section 5.2, we provide performance measurements of our ARC with Braavos in the Identity Mixer anonymous credential system (idemix) [16].

5.1 Anonymous Revocation Component (ARC) with Accumulators

ARC.SPGen($spar_g$): Using global system parameters $spar_g$ (group descriptions, parameters for ZKP, etc.), generate revocation system parameters $spar_r = (spar_g, \mathcal{RS})$, where $\mathcal{RS}$ specifies the set of supported revocation handles. The revocation system parameters can be given to any algorithm of the revocation framework.

ARC.RKGen($spar_r$):

- Generate the initial accumulator value:
  [WL]: $(sk, a, m) \leftarrow \mathsf{ACC_P.Gen}(1^\lambda, \emptyset)$.
  [BL]: $(sk, a, m) \leftarrow \mathsf{ACC_N.Gen}(1^\lambda, \emptyset)$.
- Generate witnesses for every $rh \in \mathcal{RS}$:
  [WL]: $(a, m, w^{rh}, \mathsf{upmsg}) \leftarrow \mathsf{ACC_P.Add}(sk, a, m, rh)$. Let $\mathcal{M} = \mathcal{M} \cup \{\mathsf{upmsg}\}, \mathcal{W} = \mathcal{W} \cup \{w^{rh}\}$.
  [BL]: $w^{rh} \leftarrow \mathsf{ACC_N.NonMemWitCreate}(\overline{sk}, \overline{a}, \overline{m}, rh)$. Let $\mathcal{W} = \mathcal{W} \cup \{w^{rh}\}, \mathcal{M} = \emptyset$.
- Generate signing keys $(sgk, vk) \leftarrow \mathsf{SIG.KeyGen}(spar_g)$.
- Sign the revocation information: $\sigma \leftarrow \mathsf{SIG.Sign}(sgk, (a, \mathcal{W}, \mathcal{M}))$.
- Output $rpk = vk$, $rsk = (sgk, sk, m)$, and $RI = (a, \mathcal{W}, \mathcal{M}, \sigma)$.

ARC.Revoke(rh, rsk, RI):

- Parse RI as $(a, \mathcal{W}, \mathcal{M}, \sigma)$ and rsk as (sgk, sk, m).
  Abort if $\mathsf{SIG.Verify}(vk, \sigma, (a, \mathcal{W}, \mathcal{M})) = 0$.
- Update the accumulator value:
  [WL]: $(a, m, \mathsf{upmsg}) \leftarrow \mathsf{ACC_P.Del}(sk, a, m, rh)$.
  [BL]: $(a, m, w, \mathsf{upmsg}) \leftarrow \mathsf{ACC_N.Add}(sk, a, m, rh)$.
- Let $\mathcal{M} = \mathcal{M} \cup \{\mathsf{upmsg}\}$.
- Remove the corresponding witness from the revocation information: $\mathcal{W} = \mathcal{W} \backslash \{w^{rh}\}$.
- Sign the updated revocation information: $\sigma \leftarrow \mathsf{SIG.sign}(sgk, (a, \mathcal{W}, \mathcal{M})$ and append the signature to $RI : RI = (a, \mathcal{W}, \mathcal{M}, \sigma)$.
- Output the updated RI.

ARC.RevTokenGen(rh, C, o, RI, rpk):

- Parse RI as $(a, \mathcal{W}, \mathcal{M}, \sigma)$ and rpk as vk. Abort if $\mathsf{SIG.Verify}(vk, \sigma, (a, \mathcal{W}, \mathcal{M})) = 0$.
- Update the accumulator witness w^{rh}:
  For every new $\mathsf{upmsg} \in \mathcal{M}$ since the last witness update:
  [WL]: $w^{rh} \leftarrow \mathsf{ACC_P.MemWitUpOnDel}(rh, w^{rh}, \mathsf{upmsg})$
  [BL]: $w^{rh} \leftarrow \mathsf{ACC_N.NonMemWitUpOnAdd}(rh, w^{rh}, \mathsf{upmsg})$.
- Prove knowledge of w^{rh} and rh such that w^{rh} is a witness that rh is (not) in the accumulator a:
  [WL]: $rt \leftarrow \mathsf{ZKP.Prove}[(rh, o, w^{rh}) : \mathsf{Com.Verify}(rh, C, o) = 1 \wedge \mathsf{ACC_P.VerMem}(a, rh, w^{rh}) = 1](C, a)$.
  [BL]: $rt \leftarrow \mathsf{ZKP.Prove}[(rh, o, w^{rh}) : \mathsf{Com.Verify}(rh, C, o) = 1 \wedge \mathsf{ACC_N.VerNonMem}(a, rh, w^{rh}) = 1](C, a)$.
- Output rt.

ARC.RevTokenVer(rt, C, RI, rpk):

- Parse RI as $(a, \mathcal{W}, \mathcal{M}, \sigma)$ and rpk as vk. Abort if $\mathsf{SIG.Verify}(vk, \sigma, (a, \mathcal{W}, \mathcal{M})) = 0$.
- Verify the proof:
  [WL]: $b \leftarrow \mathsf{ZKP.Verify}[(rh, o, w^{rh}) : \mathsf{Com.Verify}(rh, C, o) = 1 \wedge \mathsf{ACC_P.VerMem}(a, rh, w^{rh}) = 1](C, a, rt)$.
  [BL]: $b \leftarrow \mathsf{ZKP.Verify}[(rh, o, w^{rh}) : \mathsf{Com.Verify}(rh, C, o) = 1 \wedge \mathsf{ACC_N.VerNonMem}(a, rh, w^{rh}) = 1](C, a, rt)$.
- Output b.

Fig. 11. ARC algorithms using accumulators, for both the whitelisting and blacklisting approaches.

ARC *Syntax.* In an ARC, revocation is done via a special value called a *revocation handle* (rh) that can be embedded into the revocable object, i.e., as a special attribute in a credential. rh is bound to the revocable object with a signature. By using a commitment to rh, a proof that rh has not been revoked can be easily combined with any other proof about rh– for example, that rh was signed in a credential, as shown in Section 4.4.

As described by Camenisch et al. [10], an ARC consists of the following algorithms: ARC.SPGen, ARC.RKGen, ARC.RevTokenGen, ARC.Revoke, and ARC.RevTokenVer.

- Revocation parameters are generated using ARC.SPGen($spar_g$) $\rightarrow$ $spar_r$, and then added to the global system parameters $spar_g$.
- The revocation authority RA runs ARC.RKGen($spar_r$) $\rightarrow$ (rsk, rpk, RI) to generate the RA's secret and public keys (rsk, rpk) and the initial revocation information RI. RI contains all public data that parties need in order to generate and verify proofs of non-revocation. RI can also be supplemented by privately held witnesses.
- The RA can revoke a user based on her revocation handle rh by updating the revocation information RI: ARC.Revoke(rh, rsk, RI) $\rightarrow$ RI'.
- A user who has a valid credential can generate a publicly verifiable token rt proving that her revocation handle rh has not been revoked and that C is a commitment to rh. ARC.RevTokenGen(rh, C, o, RI, rpk) $\rightarrow$ rt. (For each new revocation token rt, the user generates a fresh commitment to rh in order to avoid making her tokens linkable. (C, o, rh) can also be used in other proofs - for example, the user should also prove that rh is an attribute of her credential.)
- A verifier can check such a token by running ARC.RevTokenVer(rt, C, RI, rpk) $\rightarrow$ $\{0, 1\}$.

Note that ARC.RevTokenGen and ARC.RevTokenVer can be integrated into an interactive protocol if interactive zero knowledge proofs are used instead of non-interactive ones.

Accumulator-Based ARC. Let $\mathsf{ACC_P}$ and $\mathsf{ACC_N}$ be a positive and a negative dynamic accumulator, respectively. Let ZKP = (Prove, Verify) be a zero knowledge proof of knowledge system (as described in Section 4.4), SIG = (KeyGen, Sign, Verify) be a signature scheme, and Com = (Commit, Verify) be a commitment scheme.

Let $\mathcal{RS}$ be a set of supported revocation handles, $\mathcal{W}$ be a list of all witnesses, and $\mathcal{M}$ be a list of all update messages that are contained in the revocation information RI and are necessary for the witness updates.

In Figure 11, we describe our accumulator-based ARC. For clarity we describe two approaches in parallel: the blacklist [**BL**] approach and the whitelist [**WL**] one. In the blacklist approach, only the revoked users' revocation handles are added to the accumulator. To prove that her credential has not been revoked, a user proves that her revocation handle rh is not in the accumulator (by means of a non-membership witness). In the whitelist approach all users' revocation handles are added to the accumulator when their credentials are issued, and are removed from the accumulator upon revocation. To prove that her credential has not been revoked, a user proves that her revocation handle rh is in the accumulator (by means of a membership witness). Naturally, a positive accumulator realizes the whitelist approach and a negative accumulator realizes the blacklist one.

Security Analysis of ARC With Accumulators. We now analyze the security of the accumulator-based ARC constructions described in Figure 11. Here we only provide an intuition for the security properties and recall the formal revocation security definitions from [10] in Appendix C.1 (Definitions 3–5).

Correctness. Correctness requires that whenever an honestly computed revocation information RI is used, an honest user is able to successfully generate valid tokens.

Theorem 4. *The* ARC *described in Figure 11 is correct (Definition 3 ([10]: Definition 4.4.1)).*

The proof follows immediately from the correctness of the accumulator scheme and the properties of ZKP.

Soundness. Revocation soundness captures the following: to make the verifier accept, the user must know the revocation handle contained in the commitment it computes a revocation token for. Further, nobody except for the revocation authority can come up with a new valid revocation information, i.e., the revocation information is always authentic. Finally, this revocation handle must not have been revoked in an earlier revocation step.

Theorem 5. *Assuming that the accumulators* $\mathsf{ACC_P}$ *and* $\mathsf{ACC_N}$ *are adaptively sound, the signature scheme* SIG *is existentially unforgeable, and the zero knowledge proof system* ZKP *is sound, then the* ARC *described in Figure 11 is sound (Definition 4 ([10]: Definition 4.4.2)) in the random oracle model.*

Proof. We prove this theorem by showing that a prover can only convince the verifier with negligible probability for each of the three given winning conditions.

Case a: Assume that the adversary outputs $(RI_{\mathcal{A}}, rt, C)$ such that the verifier accepts, but commitment verification fails. By construction, a revocation token rt is a zero knowledge proof of knowledge of rh, o such that $\mathsf{Com.Verify}(rh, C, o) = 1$. Therefore, it follows from the soundness of ZKP that $\mathsf{Com.Verify}(rh, C, o) = 0 \wedge \mathsf{ARC.RevTokenVer}(rt, C, RI, rpk) = 1$ can only happen with negligible probability.

We note that here we consider the revocation component in isolation, without bridging it with any signature or credential scheme. Therefore, it is not necessary to require the binding property from the commitment scheme. When the revocation handle is used in any other proof (of a valid signature, etc.) the binding property is required for the security of the overall system.

Case b: Assume that the adversary outputs $(RI_{\mathcal{A}}, rt, C)$ such that the verifier accepts although the given revocation information was never generated by the revocation authority. Then the following algorithm $\mathcal{R}$ can be used to break the unforgeability of the underlying signature scheme SIG. Briefly, $\mathcal{R}$ behaves as follows.

- It runs ARC.SPGen and $\mathsf{ARC.RKGen}(spar_r)$ as described in Figure 11, obtaining values rpk, RI and par_r.
- It then calls $\mathcal{A}$ on input (rpk, RI, par_r).
- For every call (Revoke, rh) to $\mathcal{O}^{\mathsf{Revoke}}$, $\mathcal{R}$ computes the updated revocation information by requesting the required signatures from the signing oracle.
- When $\mathcal{A}$ outputs $(RI_{\mathcal{A}}, rt, C)$, $\mathcal{R}$ extracts the signature component $\sigma_{\mathcal{A}}$ of $RI_{\mathcal{A}}$ (and removes the signature from the $RI_{\mathcal{A}}$ value itself.. If the signature $\sigma_{\mathcal{A}}$ does not verify on the signature-less $RI_{\mathcal{A}}$ or if the signature $\sigma_{\mathcal{A}}$ was previously returned by the signing oracle, $\mathcal{R}$ returns $\perp$. Otherwise, $\mathcal{R}$ returns $(RI_{\mathcal{A}}, \sigma_{\mathcal{A}})$.

It is easy to see that $\mathcal{R}$ succeeds in forging a signature as long as it does not output $\perp$, which occurs with non-negligible probability if $\mathcal{A}$ succeeds with non-negligible probability.

Case c: Assume that the adversary outputs $(RI_{\mathcal{A}}, rt, C)$ such that the verifier accepts although the revocation handle was previously revoked. Then through a sequence of games we can build a reduction to the adaptive soundness property of the underlying accumulator scheme and the soundness of ZKP. Briefly, $\mathcal{R}$ uses the adaptive soundness oracles to add and delete revocation handles from the accumulator, and then uses the ZKPK extractor to extract the revocation handle and the corresponding membership witness. It then outputs these revocation handle and witness as a forgery to the adaptive soundness challenger.

Privacy. Revocation privacy ensures that no adversary can tell which of two unrevoked revocation handles rh_0, rh_1 underlies a revocation token.

Theorem 6. *Assuming that the commitment scheme* Com *is hiding and that the zero knowledge proof system* ZKP *is zero knowledge, the* ARC *described in Figure 11 is private (Definition 5 ([10]: Definition 4.4.3)).*

As revocation tokens are zero knowledge, they do not leak any information about the revocation handle, and the claim follows immediately.

Enabling Dynamic Revocation. In most large systems, it is insufficient to start with a fixed set of users, and revoke users over time. Other users might need to be added to the system, or users' revocation status might need to be changed back to "un-revoked". We extend the syntax described by Camenisch et al. [10] to include this additional functionality; that is, we add a Join algorithm that allows the RA to add users to the system after the initial setup has been performed.

For the whitelist approach, this reduces the computational complexity of the RA key generation algorithm ARC.RKGen and the size of the revocation information RI, since the RA no longer needs to include all potential users' revocation handles in the initial whitelist accumulator. Users' witnesses are no longer a part of RI, since they are given to their intended holders when a Join is executed. Additionally, RevTokenGen algorithm takes a witness as a separate input.

Let $\mathcal{RU}$ be a set of the revocation handles that are already in use. $\mathcal{RU}$ is initialized in ARC.RKGen, at which time it is empty since we assume that initially there are no users in the system – they are now only added via the Join algorithm. We describe the details of the Join algorithm in Figure 12.

The RA is usually asked to execute the Join algorithm by the issuer. When issuing a revocable credential, the issuer requests a membership witness from the RA. If the user in question has previously joined - and subsequently been revoked from - the system, the issuer gives the RA her revocation handle rh. Otherwise, the RA generates a fresh revocation handle for the user. The issuer signs the revocation handle inside the credential, and sends the credential together with rh and the witness to the user.

$\mathsf{ARC.Join}(rsk, rpk, RI, rh') \rightarrow (w^{rh}, RI', rh)$: RA optionally receives a revocation handle rh' (for a previously revoked user) as input. (If the user is joining for the first time, the input rh' is $\perp$ and RA picks a fresh rh, which is returned to the user.) RA then does the following:

- Parse RI as $(a, \mathcal{M}, \sigma)$, rpk as vk, and rsk as (sgk, sk, m). Abort if $\mathsf{SIG.Verify}(vk, \sigma, (a, \mathcal{M})) = 0$ or if $rh' \notin \mathcal{RS}$.
- If $rh' = \perp$, pick a fresh revocation handle $rh' \leftarrow \mathcal{RS} \backslash \mathcal{RU}$.
- Let $rh = rh'$.
- Add rh to $\mathcal{RU}$.
- Update the accumulator value:
  [**WL**]: $(a, m, w^{rh}, \mathsf{upmsg}) \leftarrow \mathsf{ACC_P.Add}(sk, a, m, rh)$.
  [**BL**]: $(a, m, u^{rh}, \mathsf{upmsg}) \leftarrow \mathsf{ACC_N.Del}(sk, a, m, rh)$.
- Let $\mathcal{M} = \mathcal{M} \cup \{\mathsf{upmsg}\}$. $w^{rh} \leftarrow \mathsf{ACC_N.NonMemWitCreate}(\overline{sk}, \overline{a}, \overline{m}, rh)$.
- Sign the updated revocation information: $\sigma \leftarrow \mathsf{SIG.sign}(sgk, (a, \mathcal{M}))$.
- Let $RI = (a, \mathcal{M}, \sigma)$.
- Output w^{rh} and rh privately to the party that invoked the algorithm, and output RI publicly.

Fig. 12. Join algorithm.

If the Join algorithm changes the accumulator value, users need to update their witnesses to account for the additions. Therefore, the $\mathsf{ACC_P}$.MemWitUpOnAdd (in the whitelist approach) or $\mathsf{ACC_N}$.NonMemWitUpOnDel (in the blacklist approach) algorithms are used in ARC.RevTokenGen to bring the witness up to date.

Theorem 7. *The* ARC *extended with the* Join *algorithm (Figure 12) is correct (Definition 6), sound (Definition 7), and private (Definition 8) if the commitment scheme* Com *is hiding, the zero knowledge proof system* ZKP *is sound and zero knowledge, the accumulators* $\mathsf{ACC_P}$ *and* $\mathsf{ACC_N}$ *are adaptively sound, and the signature scheme* SIG *is existentially unforgeable.*

The proof is very similar to the one from the previous section and, therefore, omitted.

Join-Revoke Unlinkability. Before adding the Join algorithm, we did not have to worry about a user addition being linkable with a user revocation, because we had no user additions. Now that we do have a Join algorithm, though, this is a real concern; the revocation information could allow others to determine that the user revoked just now was the user who joined two hours ago, and not the user who joined four hours ago.

More formally, join-revoke unlinkability ensures that no adversary can determine which joining session a revocation corresponds to. This is similar to the blindness of blind signature schemes. The adversary should be unable to guess which user out of two has joined, even if it can choose the revocation handles itself, can arbitrarily join and revoke users, and can generate revocation tokens for all participants. Note, however, the revocation authority parameters must be generated honestly, and thus this definition does not imply privacy. We provide a formal definition for the join-revoke unlinkability in Appendix C.2 (Definition 9).

Theorem 8. *The* ARC *extended with* Join *algorithm and instantiated with* Braavos *is join-revoke unlinkable (Definition 9).*

Using ARC with Braavos precludes join-revoke linkability, because joins are not reflected at all in the revocation information.

5.2 Revocation for Anonymous Credentials Using Braavos: Performance Evaluation

In the previous section we discussed how one can construct an ARC using any dynamic accumulator that supports zero-knowledge proofs of member knowledge (as well as any secure commitment scheme and any existentially unforgeable signature scheme). In particular, the Braavos accumulator can be used, together with the zero knowledge mechanisms described in Section 4.4 and Fujisaki-Okamoto commitments. Using Braavos is especially efficient because during ARC.RevTokenGen, users only need to run MemWitUpOnDel, not MemWitUpOnAdd.

We now discuss the performance of a real anonymous credential system that uses an ARC with Braavos to support revocation. The system we have tested is idemix [16].[16] We left out key generation, as this is only relevant at setup once.

[16] `https://abc4trust.eu/idemix`

Idemix binds multiple user attributes (e.g., name, age, citizenship status, or employer) into a single credential. It allows a credential to be used in a number of ways, which are described in detail by Camenisch et al. [8]. Briefly, during credential presentation, a user can either simply prove that she has a valid credential ("proof of possession"), she can reveal one or more attributes ("opening", i.e., she can disclose that she is a citizen), or she can describe attributes in a range ("range proof", i.e., she can disclose that she is between 21 and 65 years old, without revealing her actual age).

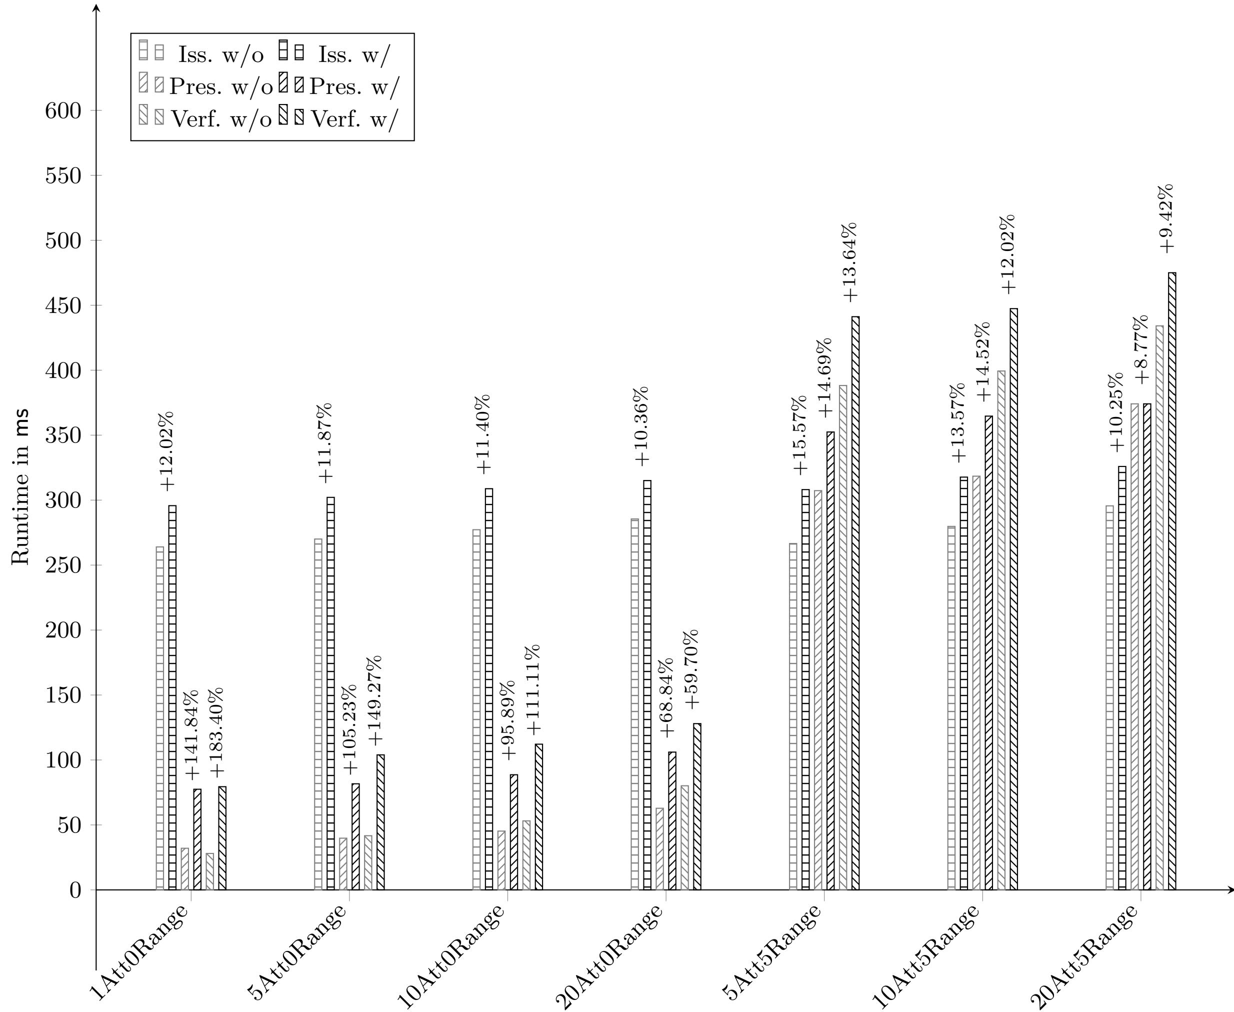

Fig. 13. 1,024Bit Measurements

In Figures 13 and 14, we show the timings of credential issuance ("Iss."), presentation ("Pres.") and verification ("Verf."), with ("w/") and without ("w/o") the anonymous revocation component (ARC) with a 1024-bit and 2048-bit RSA modulus, respectively. We use credentials with 1, 5, 10, and 20 attributes. For all of those credentials, we measure presentation consisting of a simple proof of possession. For credentials with 5, 10, and 20 attributes, we also measure presentation consisting of 5 range proofs, to show that using our ARC in conjunction with more realistic presentation policies only adds a marginal amount of run-time. On the x-axis, xAtt$y$Range means that the credential has x attributes and y range proofs are performed during presentation.

The measurements were performed locally on a machine with a Intel Quad-Core CPU with 2.70GHz, 16GB of RAM and Java8u77. In total, 1,000 runs where taken and the figures represent the average run-time. The precise numbers are given in Table 1 and Table 2.

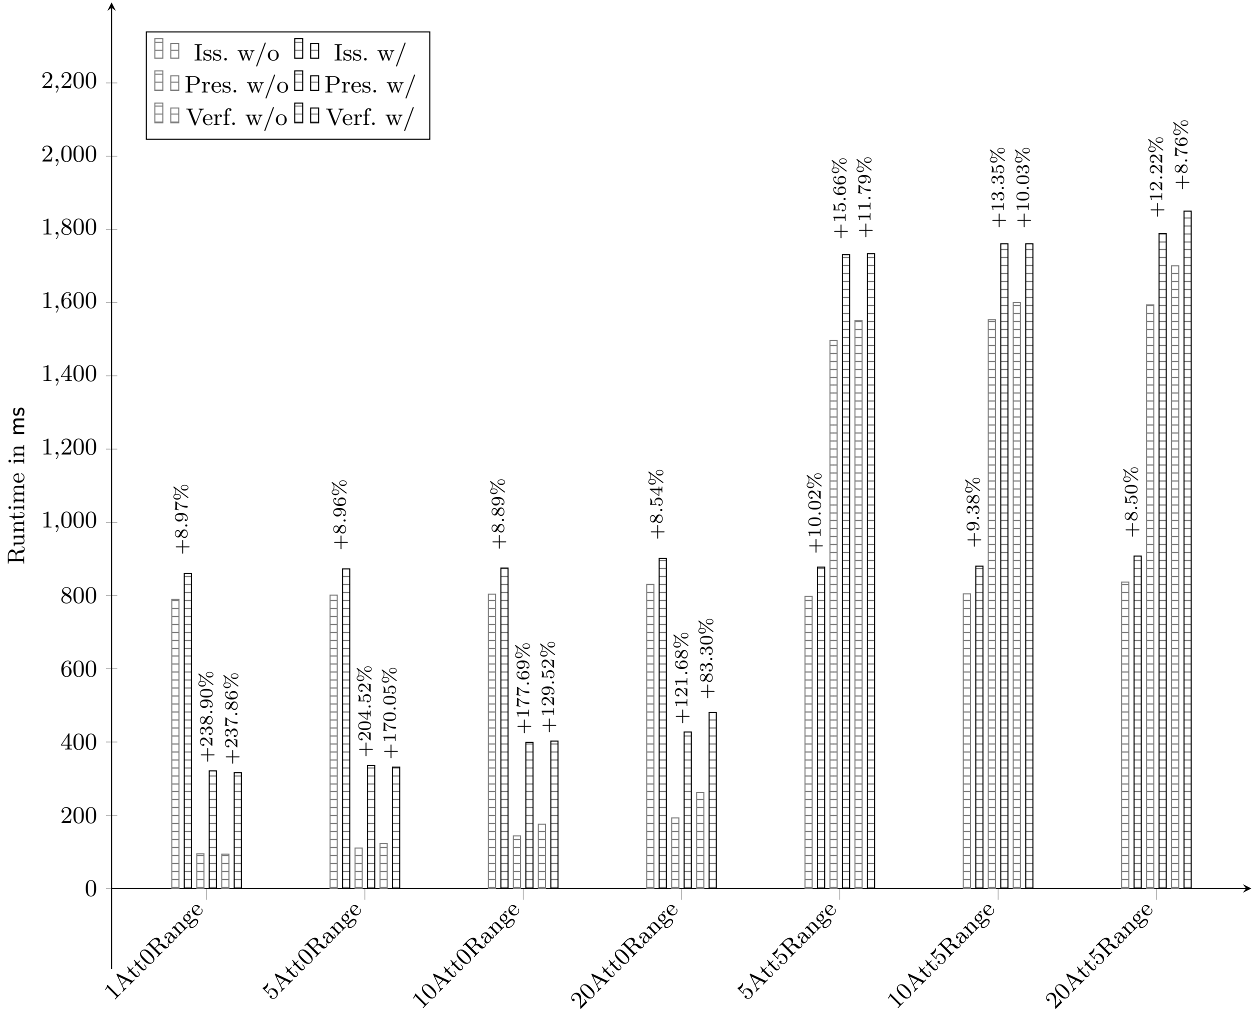

Fig. 14. 2,048Bit Measurements

Table 1. Measurements 1,024 Bit

Spec. / Action	1A0R	5A0R	10A0R	20A0R	5A5R	10A5R	20A5R
Iss. w/o	263.99	270.05	277.21	285.51	266.59	279.75	295.58
Iss. w/	295.73	302.11	308.80	315.08	308.10	317.71	325.90
Overhead	12.02%	11.87%	11.40%	10.36%	15.57%	13.57%	10.25%
Gen. w/o	32.03	39.79	45.25	62.81	307.29	318.42	343.93
Gen. w/	77.46	81.66	88.62	106.05	352.44	364.65	374.10
Overhead	141.84%	105.23%	95.85%	68.84%	14.69%	14.52%	8.77%
Verf. w/o	28.02	41.69	53.12	80.13	388.21	399.38	434.13
Verf. w/	79.42	103.92	112.14	127.97	441.18	447.39	475.02
Overhead	183.40%	149.27%	111.11%	59.70%	13.64%	12.02%	9.42%

Conclusions. The revocation overhead for issuance is the same for all credential types, and is very small. For presentation and verification, the revocation overhead is only significant for simple proofs of possession with few attributes. In these cases, the absolute numbers are very small anyway, so the overhead is not a practical problem. We think that having more attributes inside a credential and more complex presentations (e.g., involving range proofs) is more realistic, and in the case of 1,024 Bit RSA-moduli and credentials with 20 attributes and 5 range

Table 2. Measurements 2,048 Bit

Action \ Spec.	1A0R	5A0R	10A0R	20A0R	5A5R	10A5R	20A5R
Iss. w/o	789.47	801.01	803.61	830.46	797.68	804.73	836.67
Iss. w/	860.33	872.86	875.06	901.41	877.58	880.23	907.78
Overhead	8.97%	8.96%	8.89%	8.54%	10.02%	9.38%	8.50%
Gen. w/o	94.82	110.32	143.68	192.77	1,496.70	1,553.48	1,593.98
Gen. w/	321.35	335.95	398.99	427.34	1,731.06	1,760.80	1,788.75
Overhead	238.90%	204.52%	177.69%	121.68%	15.66%	13.35%	12.22%
Verf. w/o	93.62	122.68	175.35	262.16	1,550.83	1,600.33	1,700.77
Verf. w/	316.31	331.30	402.46	480.71	1,733.73	1,760.80	1,849.84
Overhead	237.86%	170.05%	129.52%	83.30%	11.79%	10.03%	8.76%

proofs, the overhead for a presentation and verification is less than 10%. Thus, anonymity-preserving revocation using ARC with Braavos is practical, as it does not impact performance significantly. The advantages one gains by achieving revocability is clearly worth the small price of 10% runtime overhead.

6 Acknowledgements

Jan Camenisch, Maria Dubovitskaya, and Kai Samelin were supported by the European Research Council under grant agreement number 321310 (PERCY). The work of Foteini Baldimtsi, Leonid Reyzin, and Sophia Yakoubov was supported, in part, by US NSF grants 1012798, 1012910, and 1422965. Foteini Baldimtsi performed this work while at Boston University. Leonid Reyzin is grateful for the hospitality and support of IST Austria, where part of this work was performed. Anna Lysyanskaya's work was supported by US NSF grant 1422361.

References

1. M. H. Au, P. P. Tsang, W. Susilo, and Y. Mu. Dynamic universal accumulators for DDH groups and their application to attribute-based anonymous credential systems. In M. Fischlin, editor, *Topics in Cryptology – CT-RSA 2009*, volume 5473 of *Lecture Notes in Computer Science*, pages 295–308. Springer, Heidelberg, Apr. 2009.
2. N. Bari and B. Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In W. Fumy, editor, *Advances in Cryptology – EUROCRYPT'97*, volume 1233 of *Lecture Notes in Computer Science*, pages 480–494. Springer, Heidelberg, May 1997.
3. J. C. Benaloh and M. de Mare. One-way accumulators: A decentralized alternative to digital sinatures (extended abstract). In T. Helleseth, editor, *Advances in Cryptology – EUROCRYPT'93*, volume 765 of *Lecture Notes in Computer Science*, pages 274–285. Springer, Heidelberg, May 1994.
4. F. Boudot. Efficient proofs that a committed number lies in an interval. In B. Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 431–444. Springer, Heidelberg, May 2000.
5. C. Cachin, S. Micali, and M. Stadler. Computationally private information retrieval with polylogarithmic communication. In J. Stern, editor, *Advances in Cryptology – EUROCRYPT'99*, volume 1592 of *Lecture Notes in Computer Science*, pages 402–414. Springer, Heidelberg, May 1999.
6. P. Camacho. On the impossibility of batch update for cryptographic accumulators. Cryptology ePrint Archive, Report 2009/612, 2009. `http://eprint.iacr.org/2009/612`.
7. P. Camacho, A. Hevia, M. A. Kiwi, and R. Opazo. Strong accumulators from collision-resistant hashing. In T.-C. Wu, C.-L. Lei, V. Rijmen, and D.-T. Lee, editors, *ISC 2008: 11th International Conference on Information Security*, volume 5222 of *Lecture Notes in Computer Science*, pages 471–486. Springer, Heidelberg, Sept. 2008.
8. J. Camenisch, M. Dubovitskaya, R. R. Enderlein, A. Lehmann, G. Neven, C. Paquin, and F. Preiss. Concepts and languages for privacy-preserving attribute-based authentication. *J. Inf. Sec. Appl.*, 19(1):25–44, 2014.
9. J. Camenisch, M. Kohlweiss, and C. Soriente. An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In S. Jarecki and G. Tsudik, editors, *PKC 2009: 12th International Conference on Theory and Practice of Public Key Cryptography*, volume 5443 of *Lecture Notes in Computer Science*, pages 481–500. Springer, Heidelberg, Mar. 2009.

10. J. Camenisch, S. Krenn, A. Lehmann, G. L. Mikkelsen, G. Neven, and M. Ø. Pedersen. Formal treatment of privacy-enhancing credential systems. In *Selected Areas in Cryptography - SAC 2015 - 22nd International Conference, Sackville, NB, Canada, August 12-14, 2015, Revised Selected Papers*, volume 9566 of *Lecture Notes in Computer Science*, pages 3–24. Springer, 2015.
11. J. Camenisch and A. Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In M. Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 61–76. Springer, Heidelberg, Aug. 2002.
12. J. Camenisch and A. Lysyanskaya. A signature scheme with efficient protocols. In S. Cimato, C. Galdi, and G. Persiano, editors, *SCN 02: 3rd International Conference on Security in Communication Networks*, volume 2576 of *Lecture Notes in Computer Science*, pages 268–289. Springer, Heidelberg, Sept. 2003.
13. J. Camenisch and M. Michels. Proving in zero-knowledge that a number is the product of two safe primes. In J. Stern, editor, *Advances in Cryptology – EUROCRYPT'99*, volume 1592 of *Lecture Notes in Computer Science*, pages 107–122. Springer, Heidelberg, May 1999.
14. J. Camenisch and M. Michels. Separability and efficiency for generic group signature schemes. In M. J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 413–430. Springer, Heidelberg, Aug. 1999.
15. J. Camenisch and M. Stadler. Efficient group signature schemes for large groups (extended abstract). In B. S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO'97*, volume 1294 of *Lecture Notes in Computer Science*, pages 410–424. Springer, Heidelberg, Aug. 1997.
16. J. Camenisch and E. Van Herreweghen. Design and implementation of the idemix anonymous credential system. In V. Atluri, editor, *ACM CCS 02: 9th Conference on Computer and Communications Security*, pages 21–30. ACM Press, Nov. 2002.
17. D. Catalano and D. Fiore. Vector commitments and their applications. In K. Kurosawa and G. Hanaoka, editors, *PKC 2013: 16th International Conference on Theory and Practice of Public Key Cryptography*, volume 7778 of *Lecture Notes in Computer Science*, pages 55–72. Springer, Heidelberg, Feb. / Mar. 2013.
18. I. Damgård and N. Triandopoulos. Supporting non-membership proofs with bilinear-map accumulators. Cryptology ePrint Archive, Report 2008/538, 2008. `http://eprint.iacr.org/2008/538`.
19. D. Derler, C. Hanser, and D. Slamanig. Revisiting cryptographic accumulators, additional properties and relations to other primitives. In K. Nyberg, editor, *Topics in Cryptology – CT-RSA 2015*, volume 9048 of *Lecture Notes in Computer Science*, pages 127–144. Springer, Heidelberg, Apr. 2015.
20. D. Derler, C. Hanser, and D. Slamanig. Revisiting cryptographic accumulators, additional properties and relations to other primitives. Cryptology ePrint Archive, Report 2015/087, 2015. `http://eprint.iacr.org/2015/087`.
21. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *Advances in Cryptology – CRYPTO'86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, Heidelberg, Aug. 1987.
22. E. Fujisaki and T. Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In B. S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO'97*, volume 1294 of *Lecture Notes in Computer Science*, pages 16–30. Springer, Heidelberg, Aug. 1997.
23. R. Gennaro, S. Halevi, and T. Rabin. Secure hash-and-sign signatures without the random oracle. In J. Stern, editor, *Advances in Cryptology – EUROCRYPT'99*, volume 1592 of *Lecture Notes in Computer Science*, pages 123–139. Springer, Heidelberg, May 1999.
24. E. Ghosh, O. Ohrimenko, D. Papadopoulos, R. Tamassia, and N. Triandopoulos. Zero-knowledge accumulators and set operations. Cryptology ePrint Archive, Report 2015/404, 2015. `http://eprint.iacr.org/2015/404`.
25. J. Lapon, M. Kohlweiss, B. D. Decker, and V. Naessens. Performance analysis of accumulator-based revocation mechanisms. In *Security and Privacy - Silver Linings in the Cloud - 25th IFIP TC-11 International Information Security Conference, SEC 2010, Held as Part of WCC 2010, Brisbane, Australia, September 20-23, 2010. Proceedings*, volume 330 of *IFIP Advances in Information and Communication Technology*, pages 289–301. Springer, 2010.
26. J. Li, N. Li, and R. Xue. Universal accumulators with efficient nonmembership proofs. In J. Katz and M. Yung, editors, *ACNS 07: 5th International Conference on Applied Cryptography and Network Security*, volume 4521 of *Lecture Notes in Computer Science*, pages 253–269. Springer, Heidelberg, June 2007.
27. B. Libert, T. Peters, and M. Yung. Group signatures with almost-for-free revocation. In R. Safavi-Naini and R. Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 571–589. Springer, Heidelberg, Aug. 2012.
28. H. Lipmaa. On diophantine complexity and statistical zero-knowledge arguments. In C.-S. Laih, editor, *Advances in Cryptology – ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, pages 398–415. Springer, Heidelberg, Nov. / Dec. 2003.
29. Y. Liu, W. Tome, L. Zhang, D. R. Choffnes, D. Levin, B. M. Maggs, A. Mislove, A. Schulman, and C. Wilson. An end-to-end measurement of certificate revocation in the web's PKI. In *Proceedings of the 2015 ACM Internet Measurement Conference, IMC 2015, Tokyo, Japan, October 28-30, 2015*, pages 183–196. ACM, 2015.

30. S. Micali, M. O. Rabin, and S. P. Vadhan. Verifiable random functions. In *40th Annual Symposium on Foundations of Computer Science*, pages 120–130. IEEE Computer Society Press, Oct. 1999.
31. T. Nakanishi, H. Fujii, Y. Hira, and N. Funabiki. Revocable group signature schemes with constant costs for signing and verifying. In S. Jarecki and G. Tsudik, editors, *PKC 2009: 12th International Conference on Theory and Practice of Public Key Cryptography*, volume 5443 of *Lecture Notes in Computer Science*, pages 463–480. Springer, Heidelberg, Mar. 2009.
32. D. Naor, M. Naor, and J. Lotspiech. Revocation and tracing schemes for stateless receivers. In J. Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 41–62. Springer, Heidelberg, Aug. 2001.
33. L. Nguyen. Accumulators from bilinear pairings and applications. In A. Menezes, editor, *Topics in Cryptology – CT-RSA 2005*, volume 3376 of *Lecture Notes in Computer Science*, pages 275–292. Springer, Heidelberg, Feb. 2005.
34. H. C. Pöhls and K. Samelin. On updatable redactable signatures. In I. Boureanu, P. Owesarski, and S. Vaudenay, editors, *ACNS 14: 12th International Conference on Applied Cryptography and Network Security*, volume 8479 of *Lecture Notes in Computer Science*, pages 457–475. Springer, Heidelberg, June 2014.
35. L. Reyzin and S. Yakoubov. Efficient asynchronous accumulators for distributed PKI. Cryptology ePrint Archive, Report 2015/718, 2015. `http://eprint.iacr.org/2015/718`.
36. A. Shamir and Y. Tauman. Improved online/offline signature schemes. In J. Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 355–367. Springer, Heidelberg, Aug. 2001.

A Lower Bound on Total Communication in Negative Accumulators

Very importantly, note that unlike prior schemes in the standard model, in Braavos and BraavosB no witness updates need to be performed when new elements are added to the accumulator. We achieve the lower bound given by Camacho [6] which states that $|\mathsf{d}|$ deletions requires the total size of update messages upmsg to be of size $\Omega(|\mathsf{d}|)$ (we are ignoring the implicit factor of $\log |S|$ used in their proof). We prove that while achieving this lower bound for dynamic accumulators, we cannot also support universality.

Theorem 9. *In a negative (or universal) accumulator,* $|\mathsf{a}|$ *additions require the total size of update messages to be of size* $\Omega(|\mathsf{a}|)$.

We prove this theorem in the style of Camacho [6].

Proof. Assume a witness holder has l non-membership witnesses. The accumulator manager then adds a set of elements a to the accumulator. The witness holder must be able to determine which of the elements for which he holds non-membership witnesses have been added to the accumulator, simply by bringing his non-membership witnesses up to date and determining which of them are still valid. This must be true even if the witness holder holds non-membership witnesses for a superset of a. Thus, the update messages must specify all of a. Specifying a set a requires at least $\log\left(\binom{l}{|\mathsf{a}|}\right) \geq |\mathsf{a}| \log \frac{l}{|\mathsf{a}|}$ bits of information. (The factor of $\log \frac{l}{|\mathsf{a}|}$ can be ignored, as it is implicit in the size of the elements.)

B BraavosB: Another Communication-Optimal Adaptively Sound Dynamic Accumulator

In this section, we introduce the BraavosB accumulator, which is an instantiation of construction B from Figure 3. BraavosB is a dynamic positive accumulator derived from a positive accumulator and a dynamic negative accumulator. The positive accumulator is used to accumulate added elements, and the negative accumulator is used to accumulate deleted elements. A proof of membership in BraavosB consists of a proof of membership in the positive accumulator (that is, a proof that the element in question has been added) and a proof of non-membership in the negative accumulator (that is, a proof that the element in question has not been deleted). In Figures 15, 16 and 17, we provide a detailed description of the construction B/ BraavosB algorithms (in terms of the two accumulators of which it is composed, the positive accumulator ACC_P and the negative dynamic accumulator ACC_N).

Just like Braavos, we aim for BraavosB to be communication optimal, and to support of zero knowledge. It might seem strange to be building a dynamic accumulator out of a negative accumulator which is already dynamic; however, using the dynamic accumulator to accumulate deleted elements instead of added elements is the key to saving on communication.

Our choice of underlying positive accumulator $\mathsf{ACC_P}$ in $\mathsf{BraavosB}$ is any existentially unforgeable digital signature scheme (as long as it supports zero knowledge proofs of knowledge of a signature on a committed value). Note that though construction B has a $\mathsf{MemWitUpOnAdd}$ algorithm, this algorithm is not used by $\mathsf{BraavosB}$, since digital signatures do not require membership updates.

The challenge that remains is instantiating $\mathsf{ACC_N}$; that is, building a negative dynamic accumulator that supports efficient zero knowledge proofs and the efficient generation of non-membership witnesses. The Merkle tree accumulator [7] does not fit the criteria, because it does not support efficient zero-knowledge proofs. The RSA [3,11,26] and Bilinear Map [33,18,1] accumulators do not fit these criteria either, because for each of them, generating a non-membership witness takes time linear in the number of accumulated elements. If the number of accumulated elements is large, and the demand for non-membership witnesses is high, this can be prohibitive; ideally, non-membership witnesses should be generated in constant time. To fill this need, we construct the range-RSA accumulator, described in Appendix B.1. The range-RSA accumulator is the core technical piece of the $\mathsf{BraavosB}$ accumulator.

$\mathsf{Gen}(1^\lambda, \emptyset)$:

1. $(\mathsf{ACC_P}.a, \mathsf{ACC_P}.m, \mathsf{ACC_P}.sk) \leftarrow \mathsf{ACC_P}.\mathsf{Gen}(1^\lambda, \emptyset)$.
2. $(\mathsf{ACC_N}.a, \mathsf{ACC_P}.m, \mathsf{ACC_N}.sk) \leftarrow \mathsf{ACC_N}.\mathsf{Gen}(1^\lambda, \emptyset)$.
3. Let $sk = (\mathsf{ACC_P}.sk, \mathsf{ACC_N}.sk)$.
4. Let $a = (\mathsf{ACC_P}.a, \mathsf{ACC_N}.a)$.
5. Let $m = (\mathsf{ACC_P}.m, \mathsf{ACC_N}.m)$.
6. Return $(sk,\ a,\ m)$.

$\mathsf{Add}(sk, a, m, x)$:

1. $(\mathsf{ACC_P}.a,\ \mathsf{ACC_P}.m,\ \mathsf{ACC_P}.w,\ \mathsf{ACC_P}.\mathsf{upmsg}) \leftarrow \mathsf{ACC_P}.\mathsf{Add}(\mathsf{ACC_P}.sk,\ \mathsf{ACC_P}.a,\ \mathsf{ACC_P}.m,\ x)$.
2. $(\mathsf{ACC_N}.a,\ \mathsf{ACC_N}.m,\ \mathsf{ACC_N}.\mathsf{upmsg}) \leftarrow \mathsf{ACC_N}.\mathsf{Del}(\mathsf{ACC_N}.sk,\ \mathsf{ACC_N}.a,\ \mathsf{ACC_N}.m,\ x)$. (This should do nothing if x is not in $\mathsf{ACC_N}$ already.)
3. $\mathsf{ACC_N}.u \leftarrow \mathsf{ACC_N}.\mathsf{NonMemWitCreate}(\mathsf{ACC_N}.sk,\ \mathsf{ACC_N}.a,\ \mathsf{ACC_N}.m,\ x)$.
4. Let $w = (\mathsf{ACC_P}.w,\ \mathsf{ACC_N}.u)$.
5. Let $\mathsf{upmsg} = (\mathsf{ACC_P}.\mathsf{upmsg},\ \mathsf{ACC_N}.\mathsf{upmsg})$.
6. Return $(a,\ m,\ w,\ \mathsf{upmsg})$.

$\mathsf{Del}(sk, a, m, x)$:

1. $(\mathsf{ACC_N}.a,\ \mathsf{ACC_N}.m,\ \mathsf{ACC_N}.\mathsf{upmsg}) \leftarrow \mathsf{ACC_N}.\mathsf{Add}(\mathsf{ACC_N}.sk,\ \mathsf{ACC_N}.a,\ \mathsf{ACC_N}.m,\ x)$.
2. Return $(a, m, \mathsf{ACC_N}.\mathsf{upmsg})$.

$\mathsf{MemWitCreate}(sk, a, m, x, (\mathsf{upmsg}_1, \ldots, \mathsf{upmsg}_t))$:

1. $\mathsf{ACC_P}.w \leftarrow \mathsf{ACC_P}.\mathsf{MemWitCreate}((\mathsf{ACC_P}.sk,\ \mathsf{ACC_P}.a,\ \mathsf{ACC_P}.m,\ x,\ (\mathsf{ACC_P}.\mathsf{upmsg}_1, \ldots, \mathsf{ACC_P}.\mathsf{upmsg}_t))$.
2. $\mathsf{ACC_N}.u \leftarrow \mathsf{ACC_N}.\mathsf{NonMemWitCreate}((\mathsf{ACC_N}.sk,\ \mathsf{ACC_N}.a,\ \mathsf{ACC_N}.m,\ x,\ (\mathsf{ACC_N}.\mathsf{upmsg}_1, \ldots, \mathsf{ACC_N}.\mathsf{upmsg}_t))$.
3. Return $w = (\mathsf{ACC_P}.w, \mathsf{ACC_N}.u)$.

Fig. 15. Construction B from Figure 3 accumulator manager algorithms (Gen, Add and Del), in terms of the underlying positive accumulator $\mathsf{ACC_P}$ and negative dynamic accumulator $\mathsf{ACC_N}$. In all of these algorithms, we assume that $a = (\mathsf{ACC_P}.a, \mathsf{ACC_N}.a)$, and $m = (\mathsf{ACC_P}.m, \mathsf{ACC_N}.m)$.

B.1 Range-RSA: A Dynamic Negative Accumulator

In this section, we present the range-RSA accumulator. This accumulator is a modified version of the RSA accumulator. Like the Merkle tree accumulator of Camacho et. al [7], the range-RSA accumulator is based on ranges; it accumulates ranges in a positive RSA accumulator [11]. All elements belonging to such a range are considered to be *non-members*; so, a proof of non-membership of x in the range-RSA accumulator would just be a proof that some range $(\mathsf{low}, \mathsf{high})$ is in the underlying positive RSA accumulator, and that $\mathsf{low} \leq x \leq \mathsf{high}$. [17]

[17] Note that range-RSA can be made universal by using open ranges instead of closed ones. A proof of non-membership of x in the range-RSA accumulator would be a proof that some range $(\mathsf{low}, \mathsf{high})$ is in the underlying positive RSA accumulator, and $\mathsf{low} < x < \mathsf{high}$. Then, the range-RSA accumulator would support proofs of membership as well as proofs of non-membership. A proof of membership of x in the range-RSA accumulator would be a proof that some

$\mathsf{MemWitUpOnAdd}(a, x, w, \mathsf{upmsg})$:

1. Parse $(\mathsf{ACC_P.upmsg}, \mathsf{ACC_N.upmsg}) = \mathsf{upmsg}$.
2. Parse $(\mathsf{ACC_P}.w, \mathsf{ACC_N}.u) = w$.
3. $\mathsf{ACC_N}.w \leftarrow \mathsf{ACC_N.MemWitUpOnDel}(x,\ \mathsf{ACC_N}.u,\ \mathsf{ACC_N.upmsg})$.
4. $\mathsf{ACC_P}.w \leftarrow \mathsf{ACC_P.MemWitUpOnAdd}(x,\ \mathsf{ACC_P}.w,\ \mathsf{ACC_P.upmsg})$.
5. Return $w = (\mathsf{ACC_P}.w, \mathsf{ACC_N}.u)$.

$\mathsf{MemWitUpOnDel}(a, x, w, \mathsf{upmsg})$:

1. Parse $(\mathsf{ACC_P}.w, \mathsf{ACC_N}.u) = w$.
2. $\mathsf{ACC_N}.u \leftarrow \mathsf{ACC_N.NonMemWitUpOnAdd}(x,\ \mathsf{ACC_N}.u,\ \mathsf{upmsg})$.
3. Return $w = (\mathsf{ACC_P}.w, \mathsf{ACC_N}.u)$.

The witness holder can run BatchMemWitUpOnDel immediately before producing a proof.

Fig. 16. Construction B from Figure 3 witness holder algorithms (MemWitUpOnAdd and MemWitUpOnDel), in terms of the underlying positive accumulator $\mathsf{ACC_P}$ and negative dynamic accumulator $\mathsf{ACC_N}$. Note that when Construction B is instantiated with a digital signature scheme as the positive accumulator $\mathsf{ACC_P}$ (as in the BraavosB construction), there is no need for a MemWitUpOnAdd algorithm.

$\mathsf{VerMem}(a, x, w)$:

1. Parse $(\mathsf{ACC_P}.a, \mathsf{ACC_N}.a) = a$.
2. Parse $(\mathsf{ACC_P}.w, \mathsf{ACC_N}.u) = w$.
3. Let $b_1 \leftarrow \mathsf{ACC_P.VerMem}(\mathsf{ACC_P}.a, x, \mathsf{ACC_P}.w)$
4. Let $b_2 \leftarrow \mathsf{ACC_N.VerNonMem}(\mathsf{ACC_N}.a, x, \mathsf{ACC_N}.u)$
5. Return 1 if $b_1 = b_2 = 1$, and return 0 otherwise.

Fig. 17. Construction B from Figure 3third party algorithms (VerMem), in terms of the underlying positive accumulator $\mathsf{ACC_P}$ and negative dynamic accumulator $\mathsf{ACC_N}$.

A range-RSA accumulator can be instantiated empty; the Gen algorithm then creates an empty positive RSA accumulator, and adds $(\mathsf{low}, \mathsf{high})$ to it (where low is smaller than the smallest supported element and high is the highest supported element). Informally, when an element x is added to the range-RSA accumulator, the range containing x is deleted from the underlying positive RSA accumulator. To replace that range, at most two new ranges are added, covering all of the other elements in the deleted range. When an element x is deleted from the range-RSA accumulator, at most two ranges containing x's direct neighbors are deleted from the underlying positive RSA accumulator. A new range, which covers x together with the deleted ranges, is added. The accumulator manager stores all of the range membership witnesses. Each range membership witness functions as the non-membership witness for all of the elements in that range; so, producing a non-membership witness is a simple matter of a look-up. In Appendix B.2, we spell out all of the details of range-RSA accumulator algorithms.

What remains is the question of how one accumulates ranges in the underlying positive RSA accumulator. This accumulator requires that all accumulated elements (in our case, ranges) be mapped to prime numbers using some canonical function f. We choose a function f that is particularly well-suited for use with efficient zero-knowledge proofs. We define $f(\mathsf{low}, \mathsf{high})$ to choose an integer suffix such that $p = \mathsf{low}||\mathsf{high}||\mathsf{suffix}$ is a prime number, where $||$ denotes concatenation, and each of $\mathsf{low}, \mathsf{high}$ and suffix use a fixed number l of bits. (Assuming that prime numbers are dense, for a sufficiently large l such a suffix will always exist.) Our choice of f allows $f(\mathsf{low}, \mathsf{high})$ to be expressed using arithmetic operations: $f(\mathsf{low}, \mathsf{high}) = 2^{2l}\mathsf{low} + 2^{l}\mathsf{high} + \mathsf{suffix}$.

Notice that the range-RSA accumulator can only accumulate elements x such that $0 \leq x \leq 2^l - 1$. This can be avoided in one of two ways: by allowing l to depend on the range $(\mathsf{low}, \mathsf{high})$ in question and encoding l as part of the f output, or by using a collision-resistant hash function to map all elements to l-bit binary strings.

range $(\mathsf{low}, \mathsf{high})$ is in the underlying positive RSA accumulator, with $x = \mathsf{low}$ or $x = \mathsf{high}$ or both. However, we do not make this simple modification in this paper, because the zero-knowledge proofs described in Section B.4 are more efficient for closed ranges. We do not discuss proofs of membership in the range-RSA accumulator any further, as we only use range-RSA to prove non-membership.

For the rest of this paper, we only consider the accumulation of elements x such that $0 \leq x \leq 2^l - 1$, since that is sufficient for the anonymous revocation application described in Section 5.

Theorem 10. *The range-RSA accumulator is a adaptively sound negative dynamic accumulator under the strong RSA assumption.*

Proof (Sketch). In order to break the adaptive soundness of the range-RSA accumulator, an adversary would need to compute a non-membership witness u for an element x that *is* actually in the accumulator. This is equivalent to computing a membership witness w for $r = f(\mathsf{low}, \mathsf{high})$ such that $\mathsf{low} \leq x \leq \mathsf{high}$ in the underlying positive RSA accumulator. No such r is actually in the underlying positive RSA accumulator, so this would require breaking the security of the positive RSA accumulator, which is hard under the strong RSA assumption.

Properties other than security (such as completeness) are self-evident.

B.2 Range-RSA Accumulator Algorithms

In this section, we describe the algorithms of the range-RSA accumulator introduced in Appendix B.1. Recall that $f(\mathsf{low}, \mathsf{high})$ is a function that can be applied to ranges to get a prime integer of the form $p = \mathsf{low}||\mathsf{high}||\mathsf{suffix}$, where $||$ denotes concatenation, and each of $\mathsf{low}, \mathsf{high}$ and suffix have a fixed number l of bits assigned to them. suffix can be any value which makes p prime.

Let RSA.Gen, RSA.Add, RSA.MemWitUpOnAdd, RSA.Del, RSA.MemWitUpOnDel and RSA.VerMem be the protocols the dynamic RSA accumulator [11]. The protocols of the range-RSA accumulator are described in Figures 18, 19 and 20.

B.3 BraavosB Soundness

The BraavosB accumulator uses Camenisch-Lysyanskaya (CL) signatures [12] as the underlying positive accumulator $\mathsf{ACC_P}$, and the range-RSA accumulator as the underlying dynamic negative accumulator $\mathsf{ACC_N}$. CL signatures are existentially unforgeable under the strong RSA assumption. Recall that the range-RSA accumulator is secure under the same assumption. By Theorem 1, this implies that the BraavosB accumulator is an adaptively sound positive dynamic accumulator under the strong RSA assumption.

Properties other than adaptive soundness (such as completeness) are self-evident.

B.4 Adding Zero Knowledge to BraavosB

Zero-Knowledge Proofs of Non-Member Knowledge in Range-RSA.

The range-RSA accumulator supports efficient zero-knowledge proofs of non-member knowledge. To prove knowledge of a non-member, a witness holder proves knowledge of values $(x, w, r, \mathsf{low}, \mathsf{high}, \mathsf{suffix})$ such that

1. $\mathsf{low} \leq x \leq \mathsf{high}$,
2. $r = \mathsf{low} * 2^{2l} + \mathsf{high} * 2^l + \mathsf{suffix}$, and
3. $w^r \equiv a \bmod n$ (where n is the appropriate RSA integer).

Given that the witness holder (also referred to as the prover) wants to keep all of x, low and high secret, she will first commit to all of $x, w, r, \mathsf{low}, \mathsf{high}$ and suffix, and then use those commitments to prove statements about the underlying values. A suitable commitment scheme is the Fujisaki-Okamoto (FO) commitment scheme [22].

For all the statements that the witness holder needs to prove, there exist standard techniques in the literature. These techniques, together with their conjunctions, come from a standard zero knowledge proof toolbox:

- proofs of knowledge of a committed value (i.e. knowledge of discrete logarithm representation modulo a composite [22]),
- proofs of equality of committed values (i.e. proof of knowledge of equality of discrete logarithms modulo a composite (or two different composites) [14]), and
- proofs that a committed value is the product of two other committed values [13].

$\mathsf{Gen}(1^\lambda)$:

1. $(\mathsf{RSA}.sk, \mathsf{RSA}.a) \leftarrow \mathsf{RSA}.\mathsf{Gen}(1^\lambda, \emptyset)$
2. Let $r = f(\mathsf{low}, \mathsf{high})$, where low denotes 0, and high denotes 2^l.
3. $(\mathsf{RSA}.a, \mathsf{RSA}.w, \mathsf{RSA}.\mathsf{upmsg}) \leftarrow \mathsf{RSA}.\mathsf{Add}(\mathsf{RSA}.a, r)$
4. Let $sk = \mathsf{RSA}.sk$.
5. Let $a = \mathsf{RSA}.a$.
6. Let $m = [((\mathsf{low}, \mathsf{high}), r, \mathsf{RSA}.w)]$.
7. Return $(sk,\ a,\ m)$.

$\mathsf{NonMemWitCreate}(x, m)$:

1. Find the entry $((\mathsf{low}, \mathsf{high}), r, \mathsf{RSA}.w)$ in m such that $\mathsf{low} \leq x \leq \mathsf{high}$.
2. If no such element exists, then x is a member of the accumulator: return $\perp$.
3. Otherwise, return $u = (r, \mathsf{RSA}.w)$.

$\mathsf{Add}(sk, a, m, x)$ (executed by the accumulator manager):

1. Let $=\leftarrow []$.
2. Find $((\mathsf{low}, \mathsf{high}), r, w)$ in m such that $\mathsf{low} \leq x \leq \mathsf{high}$. If no such element exists, then x is already a member and no work remains to be done. Otherwise, continue.
3. Remove $((\mathsf{low}, \mathsf{high}), r, w)$ from m.
4. $(\mathsf{RSA}.a, \mathsf{RSA}.m, \mathsf{RSA}.\mathsf{upmsg}) \leftarrow \mathsf{RSA}.\mathsf{Del}(\mathsf{RSA}.sk, \mathsf{RSA}.a, r)$.
5. Add $\mathsf{RSA}.\mathsf{upmsg}$ to upmsg.
6. If $x \neq \mathsf{low}$:
   (a) Let $r^l = f(\mathsf{low}, x - 1)$.
   (b) $(\mathsf{RSA}.a, \mathsf{RSA}.w^l, \mathsf{RSA}.\mathsf{upmsg}^l) \leftarrow \mathsf{RSA}.\mathsf{Add}(\mathsf{RSA}.a, r^l)$.
   (c) Add $\mathsf{RSA}.\mathsf{upmsg}^l$, $\mathsf{RSA}.w^l$ and r^l to upmsg.
   (d) Add $((\mathsf{low}, x - 1), r^l, \mathsf{RSA}.w^l)$ to m.
7. If $x \neq \mathsf{high}$:
   (a) Let $r^h = f(x + 1, \mathsf{high})$.
   (b) $(\mathsf{RSA}.a, \mathsf{RSA}.w^h, \mathsf{RSA}.\mathsf{upmsg}^h) \leftarrow \mathsf{RSA}.\mathsf{Add}(\mathsf{RSA}.a, r^h)$.
   (c) Add $\mathsf{RSA}.\mathsf{upmsg}^h$, $\mathsf{RSA}.w^h$ and r^{high} to upmsg.
   (d) Add $((x + 1, \mathsf{high}), r^h, \mathsf{RSA}.w^h)$ to m.
8. Let $a = \mathsf{RSA}.a$.
9. Return (a, m, upmsg).

$\mathsf{Del}(sk, a, m, x)$ (executed by the accumulator manager):

1. If there exists $((\mathsf{low}, \mathsf{high}), r, w)$ in m such that $\mathsf{low} \leq x \leq \mathsf{high}$, then x is already a non-member: return $\perp$. Otherwise, continue.
2. Let $\mathsf{upmsg} = []$.
3. Let $\mathsf{low} = x - 1$ and $\mathsf{high} = x + 1$.
4. If there exists $((l, \mathsf{low}),\ r^l,\ \mathsf{RSA}.w^l)$ for some l in m:
   (a) $(\mathsf{RSA}.a, \mathsf{RSA}.m, \mathsf{RSA}.\mathsf{upmsg}^l) \leftarrow \mathsf{RSA}.\mathsf{Del}(\mathsf{RSA}.sk, \mathsf{RSA}.a, r^l)$.
   (b) Add $\mathsf{RSA}.\mathsf{upmsg}^l$ to upmsg.
   (c) Delete $((l, \mathsf{low}), r^l, \mathsf{RSA}.w^l)$ from m.
   (d) Let $\mathsf{low} = l$.
5. If there exists $((\mathsf{high}, h),\ r^h,\ \mathsf{RSA}.w^h)$ for some h in m:
   (a) $(\mathsf{RSA}.a, \mathsf{RSA}.m, \mathsf{RSA}.\mathsf{upmsg}^h) \leftarrow \mathsf{RSA}.\mathsf{Del}(\mathsf{RSA}.sk, \mathsf{RSA}.a, r^h)$.
   (b) Add $\mathsf{RSA}.\mathsf{upmsg}^h$ to upmsg.
   (c) Delete $((\mathsf{high}, h), r^h, \mathsf{RSA}.w^h)$ from m.
   (d) Let $\mathsf{high} = h$.
6. $r = f(\mathsf{low}, \mathsf{high})$.
7. $(\mathsf{RSA}.a, \mathsf{RSA}.m, \mathsf{RSA}.w, \mathsf{RSA}.\mathsf{upmsg}) \leftarrow \mathsf{RSA}.\mathsf{Add}(\mathsf{RSA}.sk, \mathsf{RSA}.a, r)$.
8. Add r, $\mathsf{RSA}.w$ and $\mathsf{RSA}.\mathsf{upmsg}$ to upmsg.
9. Add $((\mathsf{low}, \mathsf{high}), r, \mathsf{RSA}.w)$ to m.
10. Let $a = \mathsf{RSA}.a$.
11. Return (a, m, upmsg).

Fig. 18. Range-RSA accumulator manager algorithms (Gen, NonMemWitCreate, Add and Del).

All of the above-mentioned proofs are sound under the strong RSA assumption.

More specifically, to prove item 1 in the list above, the witness holder will use a range proof [4,28]. A range proof involves showing that $x - \mathsf{low} \geq 0$ and that $\mathsf{high} - x \geq 0$. To do so, one can use the observation that any non-negative number can be represented as a sum of four squares. The prover would have to find these four squares for each of $x - \mathsf{low}$ and $\mathsf{high} - x$, commit to them, and use commitment equality and product proofs to demonstrate that each of $x - \mathsf{low}$ and $\mathsf{high} - x$ is, in fact, a sum of four squares.

To prove item 2, the witness holder will again have to use commitments to all of the elements in question: r, low, high and suffix. She will then use the homomorphic properties of FO commitments to obtain a single

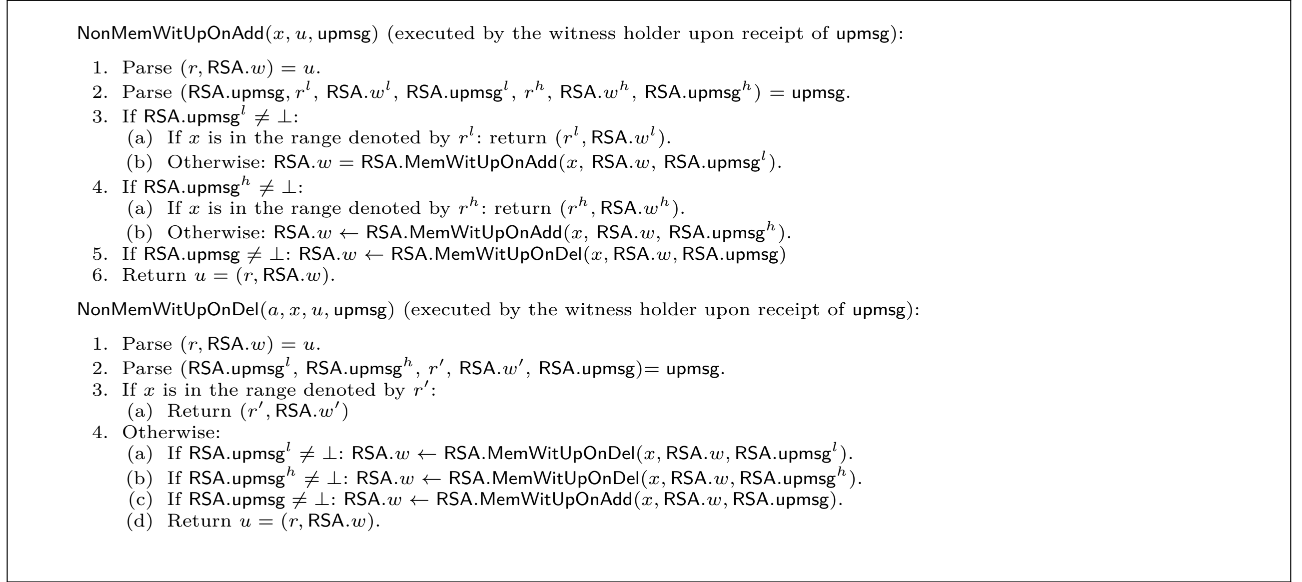

$\mathsf{NonMemWitUpOnAdd}(x, u, \mathsf{upmsg})$ (executed by the witness holder upon receipt of upmsg):

1. Parse $(r, \mathsf{RSA}.w) = u$.
2. Parse $(\mathsf{RSA.upmsg}, r^l, \mathsf{RSA}.w^l, \mathsf{RSA.upmsg}^l, r^h, \mathsf{RSA}.w^h, \mathsf{RSA.upmsg}^h) = \mathsf{upmsg}$.
3. If $\mathsf{RSA.upmsg}^l \neq \perp$:
   (a) If x is in the range denoted by r^l: return $(r^l, \mathsf{RSA}.w^l)$.
   (b) Otherwise: $\mathsf{RSA}.w = \mathsf{RSA.MemWitUpOnAdd}(x, \mathsf{RSA}.w, \mathsf{RSA.upmsg}^l)$.
4. If $\mathsf{RSA.upmsg}^h \neq \perp$:
   (a) If x is in the range denoted by r^h: return $(r^h, \mathsf{RSA}.w^h)$.
   (b) Otherwise: $\mathsf{RSA}.w \leftarrow \mathsf{RSA.MemWitUpOnAdd}(x, \mathsf{RSA}.w, \mathsf{RSA.upmsg}^h)$.
5. If $\mathsf{RSA.upmsg} \neq \perp$: $\mathsf{RSA}.w \leftarrow \mathsf{RSA.MemWitUpOnDel}(x, \mathsf{RSA}.w, \mathsf{RSA.upmsg})$
6. Return $u = (r, \mathsf{RSA}.w)$.

$\mathsf{NonMemWitUpOnDel}(a, x, u, \mathsf{upmsg})$ (executed by the witness holder upon receipt of upmsg):

1. Parse $(r, \mathsf{RSA}.w) = u$.
2. Parse $(\mathsf{RSA.upmsg}^l, \mathsf{RSA.upmsg}^h, r', \mathsf{RSA}.w', \mathsf{RSA.upmsg}) = \mathsf{upmsg}$.
3. If x is in the range denoted by r':
   (a) Return $(r', \mathsf{RSA}.w')$
4. Otherwise:
   (a) If $\mathsf{RSA.upmsg}^l \neq \perp$: $\mathsf{RSA}.w \leftarrow \mathsf{RSA.MemWitUpOnDel}(x, \mathsf{RSA}.w, \mathsf{RSA.upmsg}^l)$.
   (b) If $\mathsf{RSA.upmsg}^h \neq \perp$: $\mathsf{RSA}.w \leftarrow \mathsf{RSA.MemWitUpOnDel}(x, \mathsf{RSA}.w, \mathsf{RSA.upmsg}^h)$.
   (c) If $\mathsf{RSA.upmsg} \neq \perp$: $\mathsf{RSA}.w \leftarrow \mathsf{RSA.MemWitUpOnAdd}(x, \mathsf{RSA}.w, \mathsf{RSA.upmsg})$.
   (d) Return $u = (r, \mathsf{RSA}.w)$.

Fig. 19. Range-RSA witness holder algorithms (NonMemWitUpOnAdd and NonMemWitUpOnDel).

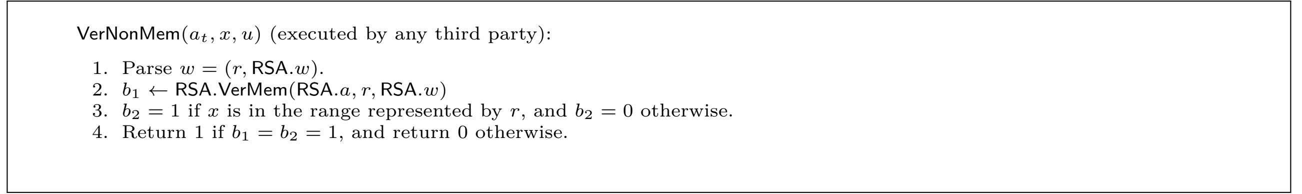

$\mathsf{VerNonMem}(a_t, x, u)$ (executed by any third party):

1. Parse $w = (r, \mathsf{RSA}.w)$.
2. $b_1 \leftarrow \mathsf{RSA.VerMem}(\mathsf{RSA}.a, r, \mathsf{RSA}.w)$
3. $b_2 = 1$ if x is in the range represented by r, and $b_2 = 0$ otherwise.
4. Return 1 if $b_1 = b_2 = 1$, and return 0 otherwise.

Fig. 20. Range-RSA third party algorithm (VerNonMem).

commitment to $\mathsf{low} * 2^{2l} + \mathsf{high} * 2^l + \mathsf{suffix}$, and use a commitment equality proof to show that the resulting commitment is to r.

Finally, to prove item 3, the witness holder will show that a committed value has been accumulated. Range-RSA uses the CL-RSA accumulator [11] as an underlying building block, so this proof will be done as described by Camenisch and Lysyanskaya [11].

Note that most of these proofs require interaction between the prover and the verifier. One could apply the Fiat-Shamir heuristic [21] to obtain a non-interactive zero-knowledge proof, but this would require assuming random oracles.

Zero-Knowledge Proofs of Member Knowledge in BraavosB*.*

To prove member knowledge for BraavosB in zero knowledge, a witness holder would have to compute the following zero-knowledge proof (described using Camenisch-Stadler [15] notation):

$$\begin{aligned} \mathsf{ZKP}[(x, r, \mathsf{ACC_A}.w, \mathsf{ACC_{NA}}.w) :& \\ &\wedge \mathsf{ACC_P}.\mathsf{VerMem}(\mathsf{ACC_P}.a, x, \mathsf{ACC_P}.w) \\ &\wedge \mathsf{ACC_N}.\mathsf{VerNonMem}(\mathsf{ACC_N}.a, x, \mathsf{ACC_N}.u) \\ ](\mathsf{ACC_{NA}}.a, \mathsf{ACC_A}.a)& \end{aligned}$$

Where $\mathsf{ACC_P}$ is the signature scheme $\mathsf{SIG}_{CL} = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ due to Camenisch and Lysyanskaya [12], and $\mathsf{ACC_N}$ is the range-RSA accumulator.

If integration into a larger system (where it is important to be able to link the witnesses used in the proof to other statements) is desired, the witness holder computes commitments to the membership witness $\mathsf{ACC_P}.w$

and non-membership witness $\mathsf{ACC_N}.u$: $(C_1, o_1) = \mathsf{Com.Commit}(\mathsf{ACC_P}.w)$ and $(C_2, o_2) = \mathsf{Com.Commit}(\mathsf{ACC_N}.u)$, where o_1 and o_2 are decommitment values. The proof is then enhanced, as follows:

$$\begin{aligned}&\mathsf{ZKP}[(x, \mathsf{ACC_P}.w, \mathsf{ACC_N}.u, o_1, o_2) :\\ &\quad \mathsf{Com.Verify}(C_1, \mathsf{ACC_P}.w, o_1)\\ &\quad \wedge \mathsf{Com.Verify}(C_2, \mathsf{ACC_N}.u, o_2)\\ &\quad \wedge \mathsf{ACC_P.VerMem}(\mathsf{ACC_P}.a, x, \mathsf{ACC_P}.w)\\ &\quad \wedge \mathsf{ACC_N.VerNonMem}(\mathsf{ACC_N}.a, x, \mathsf{ACC_N}.u)\\ &\quad](\mathsf{ACC_P}.a, \mathsf{ACC_N}.a, C_1, C_2)\end{aligned}$$

For concrete descriptions of the individual clauses of this proof, please refer to Fujisaki and Okamoto [22] and Camenisch and Lysyanskaya [12,11].

C Security of ARC

In this section we provide the security definitions for ARC. We first recall the revocation security definitions from [10] and then modify them to incorporate the Join algorithm.

C.1 Recalling Security Definitions by Camenisch et al.

Here, we recall the security definitions given by Camenisch et al. [10].

Correctness. Correctness requires that whenever an honestly computed revocation information RI is used, an honest user is able to successfully generate valid tokens.

Definition 3 (Revocation Correctness). *There exists a negligible function ν such that the following holds for all ordered sets $\mathcal{RS}$ of revocation handles and $rh' \notin \mathcal{RS}$:*

$$\Pr\Big[\mathsf{RevTokenVer}(rt, c', RI, rpk) = 0 \;:\; spar_g \leftarrow \mathsf{SPGen}_g(1^\lambda), spar_r \leftarrow \mathsf{SPGen}(spar_g), (C', o') \leftarrow \mathsf{Commit}(rh'), (rsk, rpk, RI) \leftarrow \mathsf{RKGen}(spar_r), RI \leftarrow \mathsf{Revoke}(rsk, RI, rh)\ \forall rh \in \mathcal{RS}, rt \leftarrow \mathsf{RevTokenGen}(rh', C', o', RI, rpk)\Big] \leq \nu(\lambda)\,.$$

Soundness. Revocation soundness captures the following: to make the verifier accept, the user must know the revocation handle contained in the commitment it computes a revocation token for. Further, nobody except for the revocation authority can come up with a new valid revocation information, i.e., the revocation information is always authentic. Finally, this revocation handle must not have been revoked in an earlier revocation step.

We, however, would like to slightly correct the original definition to eliminate a small flaw. Namely, for an empty revocation list, if $RI_{\mathcal{A}} = RI_0$, all revocation handles will potentially verify, as nothing is revoked, but RI_0 is not in $\mathcal{L}$.

Definition 4 (Revocation Soundness). *A revocation scheme is* sound *if there exists an efficient algorithm $\mathsf{E^r} = (\mathsf{E^r_1}, \mathsf{E^r_2})$, called the* extractor*, that satisfies the following properties:*

- $\mathsf{E^r_1}$ *outputs parameters and trapdoors, such that the parameters are indistinguishable from correctly computed system parameters:*

$$\{spar_r : spar_r \leftarrow \mathsf{SPGen}(\mathsf{SPGen_g}(1^\lambda, \lambda))\} \sim \{spar_r : (spar_r, \tau_\mathtt{r}) \leftarrow \mathsf{E^r_1}(\mathsf{SPGen_g}(1^\lambda))\}\,,$$

  and

– *for every efficient adversary $\mathcal{A}$ there exists a negligible function such that:*

$$\Pr\Big[\mathsf{RevTokenVer}(rt, C, RI_{\mathcal{A}}, rpk) = 1 \ \wedge\ \mathcal{L} \neq \emptyset \ \wedge \Big(\mathsf{Verify}(C, rh, o) = 0 \ \vee \nexists\, (RI_{\mathcal{A}}, *, *) \in \mathcal{L} \ \vee \exists\, (RI_{\mathcal{A}}, epoch, *), (RI', epoch', rh) \in \mathcal{L} \ :\ epoch' \leq epoch\Big) : (spar_r, \tau_{\mathtt{r}}) \leftarrow \mathsf{E}_1^{\mathtt{r}}(\mathsf{SPGen}_{\mathtt{g}}(1^{\lambda})),\ (rsk, rpk, RI_0) \leftarrow \mathsf{RKGen}(spar_r), (RI_{\mathcal{A}}, rt, C) \leftarrow \mathcal{A}^{\mathcal{O}^{\mathtt{ra}}}(rpk, RI_0, spar_r), (rh, o) \leftarrow \mathsf{E}_2^{\mathtt{r}}(\tau_{\mathtt{r}}, rpk, rt, C)\Big] \leq \nu(\lambda).$$

Here, $'*'$ denotes a wildcard, the oracle $\mathcal{O}^{\mathtt{ra}}$ simulates an honest revocation authority as follows, where initially $\mathcal{L} = \emptyset$ and $epoch = 0$:

– *On input $(\mathtt{revoke}, rh)$ for some revocation handle $rh \in \mathcal{RS}$, the oracle first revokes rh by computing $RI' \leftarrow \mathsf{Revoke}(rsk, RI, rh)$. It then updates its internal state as $RI \leftarrow RI'$ and $epoch \leftarrow epoch + 1$. It then adds $(RI, epoch, rh)$ to $\mathcal{L}$, and hands back the updated revocation information RI to the adversary.*

Privacy. Revocation privacy ensures that no adversary can tell which of two unrevoked revocation handles rh_0, rh_1 underlies a revocation token. It is formally defined through the following experiment.

Definition 5 (Revocation Privacy). *A revocation scheme is* private*, if for every efficient adversary $\mathcal{A}$ there exists a negligible function ν such that the following holds:*

$$\Pr\Big[b' = b \ \wedge\ rt_0 \neq \bot \ \wedge\ rt_1 \neq \bot \ : spar_g \leftarrow \mathsf{SPGen}_g(1^{\lambda}), \quad spar_r \leftarrow \mathsf{SPGen}(spar_g), (rpk, RI, rh_0, rh_1, \mathsf{state}) \leftarrow \mathcal{A}(spar_r), (C_i, o_i) \leftarrow \mathsf{Commit}(rh_i), rt_i \leftarrow \mathsf{RevTokenGen}(spar_r, rh_i, C_i, o_i, RI, rpk),\ i = 0, 1, b \leftarrow \{0,1\}, b' \leftarrow \mathcal{A}(C_b, rt_b, \mathsf{state})\Big] \leq \frac{1}{2} + \nu(\lambda)\,.$$

C.2 Security Definitions for ARC with Join

We now modify the security definitions of the ARC to incorporate the Join algorithm.

Correctness. It is required that whenever an honestly computed revocation information is used, an honest user is able to successfully generate valid tokens after Join is run on her rh and either her rh was not revoked before or rh was unrevoked afterwards (by running Join again on the same rh).

Definition 6 (Revocation Correctness). *There exists a negligible function ν such that the following holds for any $rh \in \mathcal{RS}$ and for all possible lists $\mathcal{RO}$ of operations (Join, t_j, rh) and $(\mathsf{Revoke}, t_r, rh)$ performed on the revocation handle rh:*

$$\Pr\Big[\mathsf{RevTokenVer}(rt, c', RI, rpk) = 0 \ \wedge \big(\exists(\mathsf{Join}, epoch, rh) \in \mathcal{RO} \ s.t.\ \nexists(\mathsf{Revoke}, epoch', rh) \in \mathcal{RO} \text{ for which } epoch' > epoch\big) : spar_g \leftarrow \mathsf{SPGen}_{\mathsf{g}}(1^{\lambda}), spar_r \leftarrow \mathsf{SPGen}(spar_g), (C', o') \leftarrow \mathsf{Commit}(rh), (rsk, rpk, RI) \leftarrow \mathsf{RKGen}(spar_r), rt \leftarrow \mathsf{RevTokenGen}(rh, C', o', RI, rpk)\Big] \leq \nu(\lambda)\ .$$

where epoch denotes sequence numbers of the join and revoke operations, respectively.

Soundness. First, we note that the soundness definition (Def. 4) can be simplified as follows: instead of providing the revocation information itself, the adversary can just reference the version of RI with respect to which it outputs the proof. The second winning condition, therefore, can be omitted as the "unforgeability" of the revocation information itself can be achieved by the RA signing the revocation information at the time of every update. This can be done outside of the scope of ARC and thus we remove signing revocation information from the algorithms.

Intuitively, the soundness requires the following: to make the verifier accept: (1) the user must know the revocation handle contained in the commitment it computes a revocation token for; (2) the revocation handle must not be in the revoked state with respect to the version of the revocation information it refers to.

Definition 7 (Revocation Soundness). *A revocation scheme is* sound *if there exists an efficient algorithm* $\mathsf{E}^{\mathtt{r}} = (\mathsf{E}^{\mathtt{r}}_1, \mathsf{E}^{\mathtt{r}}_2)$*, called the* extractor*, that satisfies the following properties:*

- $\mathsf{E}^{\mathtt{r}}_1$ *outputs parameters and trapdoors, such that the parameters are indistinguishable from correctly computed system parameters:*

$$\{spar_r : \ spar_r \leftarrow \mathsf{SPGen}(\mathsf{SPGen}_{\mathsf{g}}(1^\lambda, \lambda))\} \ \sim \\ \{spar_r : \ (spar_r, \tau_{\mathtt{r}}) \leftarrow \mathsf{E}^{\mathtt{r}}_1(\mathsf{SPGen}_{\mathsf{g}}(1^\lambda))\},$$

and
- *for every efficient adversary* $\mathcal{A}$ *there exists a negligible function such that:*

$$\begin{aligned}\Pr\Big[&\mathsf{RevTokenVer}(rt, C, RI, rpk) = 1 \ \wedge \\ &\Big(\mathsf{Verify}(C, rh, o) = 0 \ \vee (\forall\ (RI, \mathcal{V}) \in \mathcal{L} \ : \ rh \notin \mathcal{V})\Big) : \\ &(spar_r, \tau_{\mathtt{r}}) \leftarrow \mathsf{E}^{\mathtt{r}}_1(\mathsf{SPGen}_{\mathsf{g}}(1^\lambda)),\ (rsk, rpk, RI_0) \leftarrow \mathsf{RKGen}(spar_r), \\ &(rt, C) \leftarrow \mathcal{A}^{\mathcal{O}^{\mathtt{ra}}}(rpk, RI_0, spar_r), \\ &(rh, o) \leftarrow \mathsf{E}^{\mathtt{r}}_2(\tau_{\mathtt{r}}, rpk, rt, C)\Big] \le \nu(\lambda).\end{aligned}$$

Here, the oracle $\mathcal{O}^{\mathtt{ra}}$ *simulates an honest revocation authority as follows, where initially* $\mathcal{L} = \emptyset$ *and* $\mathcal{V} = \emptyset$*:*

- *On input* $(\mathtt{join}, rh)$*, the oracle first checks if* $rh \in \mathcal{V}$*, it does nothing. Else, it runs* $(w^{rh}, RI', rh') \leftarrow \mathsf{Join}(rsk, rpk, RI, rh)$ *and sets* $\mathcal{V} \leftarrow \mathcal{V} \cup \{rh'\}$*. It then updates its internal state as* $RI \leftarrow RI'$*. It then adds* $(RI, \mathcal{V})$ *to* $\mathcal{L}$*, and hands back the updated revocation information* RI' *to the adversary.*
- *On input* $(\mathtt{revoke}, rh)$ *for some revocation handle* $rh \in \mathcal{RS}$*, the oracle first checks if* $rh \in \mathcal{V}$*. If not it aborts. If yes, it revokes* rh *by computing* $RI' \leftarrow \mathsf{Revoke}(rsk, RI, rh)$*. It then sets* $\mathcal{V} \leftarrow \mathcal{V} \setminus \{rh\}$ *and adds* $(RI, \mathcal{V})$ *to* $\mathcal{L}$*, and hands back the updated revocation information* RI' *to the adversary.*

Privacy. Revocation privacy ensures that no adversary can tell which of two non-revoked revocation handles underlies a revocation token.

Definition 8 (Revocation Privacy). *A revocation scheme is* private*, if for every efficient adversary* $\mathcal{A}$ *there exists a negligible function* ν *such that the following holds:*

$$\begin{aligned}\Big|\Pr\Big[&b' = b \ \wedge \ rt_0 \neq \bot \ \wedge \ rt_1 \neq \bot \ : \\ &spar_g \leftarrow \mathsf{SPGen}_g(1^\lambda), \quad spar_r \leftarrow \mathsf{SPGen}(spar_g), \\ &(rpk, RI, rh_0, rh_1, \mathsf{state}) \leftarrow \mathcal{A}(spar_r), (C_i, o_i) \leftarrow \mathsf{Commit}(rh_i), \\ &rt_i \leftarrow \mathsf{RevTokenGen}(spar_r, rh_i, C_i, o_i, RI, rpk),\ i \in \{0,1\}, \\ &\quad b \leftarrow \{0,1\}, b' \leftarrow \mathcal{A}(C_b, rt_b, \mathsf{state})\Big] - \frac{1}{2}\Big| \le \nu(\lambda).\end{aligned}$$

Experiment $\mathsf{Join/Revoke\text{-}Unlinkability}_{\mathcal{A}}^{\mathsf{ARC}}(\lambda)$

$d \leftarrow \{0,1\}$;

$spar_g \leftarrow \mathsf{SPGen}_g(1^\lambda)$; $spar_r \leftarrow \mathsf{SPGen}_r(spar_g)$; $(rsk, rpk, RI) \leftarrow \mathsf{RKGen}(spar_r)$;

$(\mathsf{state}, rh_0, rh_1, b) \leftarrow \mathcal{A}^{\mathcal{O}^{\mathrm{ra}}}(spar_g, spar_r, rpk, RI)$,

where oracle $\mathcal{O}^{\mathrm{ra}}$ is an honest revocation authority as defined in the ARC-component.

If rh_0 and rh_1 are both in the same state (where the state could be one of never-joined, joined, or revoked) do nothing; else put them into the same state based on the adversary's choice:

Let RI be the current revocation information.

For $i = 0, 1$, let $(w^{rh}, RI, rh_i) \leftarrow \mathsf{Join}(rsk, rpk, RI, rh_i)$.

If $b = 1$, for $i = 0, 1$, let $RI \leftarrow \mathsf{Revoke}(rsk, rpk, RI, rh_i)$.

$d' \leftarrow \mathcal{A}^{\mathcal{O}^{\mathrm{ra}'}, \mathcal{O}^{\mathrm{user}'}}(\mathsf{state}, RI)$,

where oracle $\mathcal{O}^{\mathrm{ra}'}$ is a revocation authority with the following behavior:

On input $(\mathtt{joinrh}, rh)$:

Ignore, if $rh \in \{rh_0, rh_1\}$ or rh is currently joined.

Let $(w^{rh}, RI, rh) \leftarrow \mathsf{Join}(rsk, rpk, RI, rh)$.

Return RI, and the witness w^{rh}, to the adversary.

On input $(\mathtt{joinb}, b)$:

Ignore, if $b \notin \{0,1\}$ or $rh_{b\oplus d}$ is currently joined.

Let $(w^{rh_{b\oplus d}}, RI, rh_{b\oplus d}) \leftarrow \mathsf{Join}(rsk, rpk, RI, rh_{b\oplus d})$.

Return RI, and save the witness $w^{rh_{b\oplus d}}$ for further usage.

On input $(\mathtt{revoke}, rh)$:

If $rh \in \{rh_0, rh_1\}$, do:

If not both rh_0 and rh_1 are currently joined, ignore.

Let $RI_0 \leftarrow \mathsf{Revoke}(rsk, rpk, RI, rh_0)$; $RI_1 \leftarrow \mathsf{Revoke}(rsk, rpk, RI_0, rh_1)$.

Return (RI_0, RI_1) to the adversary;

else do:

Let $RI \leftarrow \mathsf{Revoke}(rsk, rpk, RI, rh)$.

Return RI to the adversary.

The oracle $\mathcal{O}^{\mathrm{user}'}$ behaves as follows:

On input $(\mathtt{generate}, rh)$, ignore, if $rh \notin \{rh_0, rh_1\}$ or not both rh_0 and rh_1 are in a joined state.

Let $(C, o) \leftarrow \mathsf{Commit}(rh_b)$.

Let $rt \leftarrow \mathsf{RevTokenGen}(spar_r, rh_b, C, o, RI, rpk)$.

Return (rt, C).

return 1, if $d' = d$

return 0.

Fig. 21. Join/Revoke Unlinkability

Join-Revoke Unlinkability. Join/Revoke Unlinkability ensures that no adversary can decide to which joining session a revocation corresponds to. This is related to the blindness of blind signature schemes, in the sense that the adversary has to guess which user out of two has joined, even if it can choose the revocation handles itself, and can arbitrarily join and revoke users, yet can also generate revocation tokens for all participants. Note, however, the revocation authority parameters are generated honestly, and thus this definition does not imply privacy.

Definition 9 (Join/Revoke Unlinkability). *A revocation component* ARC *is called join/revoke-unlinkable, if for any PPT adversary $\mathcal{A}$ there exists a negligible function ν such that:*

$$\left|\Pr[\mathsf{Join/Revoke\text{-}Unlinkability}_{\mathcal{A}}^{\mathsf{ARC}}(\lambda) = 1] - \frac{1}{2}\right| \leq \nu(\lambda)$$

The corresponding experiment is depicted in Figure 21.

Briefly, the experiment is defined as follows. Given the honestly generated parameters, the adversary outputs two revocation handles of its choice. If these handles do not share the same state (never-joined, joined, or revoked) then a bit b indicates which state both handles should be put in (either joined or revoked). Before outputting the

challenge revocation handles, the adversary has access to the oracle that acts as an honest revocation authority and can generate presentation tokens by itself.

After producing the challenge revocation handles, the adversary can query two oracles: the first one acts like a standard revocation authority with respect to all but the challenge revocation handles. The challenge ones it handles as follows. Upon a revocation request it revokes both challenge revocation handles if both are in the joined state, otherwise, ignores the request. Upon a join request it joins either the requested challenge revocation handle or the other one based on the bit d (if it is not in the joined state already). It does not return the corresponding witness to the adversary though, since it may leak the information about which handle was joined. Without the witness the adversary cannot generate presentation tokens for the challenge handles by itself. Thus, it can obtain those tokens by querying the second oracle that outputs presentation tokens for a challenge revocation handle only if both challenge handles are in the joined state. Note that the revocation information RI is altered after the different states.