Cryptanalysis of a Dynamic Universal Accumulator over Bilinear Groups*

Alex Biryukov[1], Aleksei Udovenko[2], Giuseppe Vitto[1](✉)

[1] DCS&SnT, University of Luxembourg, Esch-sur-Alzette, Luxembourg
name.surname@uni.lu
[2] CryptoExperts, Paris, France
aleksei@affine.group

Abstract. In this paper we cryptanalyse the two accumulator variants proposed by Au *et al.* [1], which we call the α-*based* construction and the common reference string-based ($\mathcal{CRS}$-*based*) construction. We show that if non-membership witnesses are issued according to the α-based construction, an attacker that has access to multiple witnesses is able to efficiently recover the secret accumulator parameter α and completely break its security. More precisely, if p is the order of the underlying bilinear group, the knowledge of $O(\log p \log \log p)$ non-membership witnesses permits to successfully recover α. Further optimizations and different attack scenarios allow to reduce the number of required witnesses to $O(\log p)$, together with practical attack complexity. Moreover, we show that accumulator's collision resistance can be broken if just one of these non-membership witnesses is known to the attacker. We then show how all these attacks for the α-based construction can be easily prevented by using instead a corrected expression for witnesses.

Although outside the original security model assumed by Au *et al.* but motivated by some possible concrete application of the scheme where the Manager must have exclusive rights for issuing witnesses (e.g. white/black list based authentication mechanisms), we show that if non-membership witnesses are issued using the $\mathcal{CRS}$-based construction and the $\mathcal{CRS}$ is kept secret by the Manager, an attacker accessing multiple witnesses can reconstruct the $\mathcal{CRS}$ and compute witnesses for arbitrary new elements. In particular, if the accumulator is initialized by adding m secret elements, the knowledge of m non-membership witnesses allows to succeed in such attack.

Keywords: accumulator · universal · dynamic · cryptanalysis · anonymous credentials

* This version of the contribution has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: `http://dx.doi.org/10.1007/978-3-030-75539-3_12`. Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use `https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms`.

This work was supported by the Luxembourg National Research Fund (FNR) project FinCrypt (C17/IS/11684537).

1 Introduction

A cryptographic accumulator scheme permits to aggregate values of a possibly very large set into a short digest, which is commonly referred to as the *accumulator value*. Unlike hash functions, where, similarly, (arbitrary) long data is mapped into a fixed length digest, accumulator schemes permit to additionally show whenever an element is accumulated or not, thanks to special values called *witnesses*. Depending on the accumulator design, we can have two kinds of witnesses: *membership witnesses*, which permit to show that an element is included into the accumulator, and *non-membership witnesses*, which, on the contrary, permit to show that an element is not included. Accumulator schemes which support both are called *universal* and the possibility to dynamically add and delete elements, give them the name of *dynamic accumulators*.

The first accumulator scheme was formalized by Benaloh and De Mare [3] in 1993 as a time-stamping protocol. Since then, many other accumulator schemes have been proposed and they play an important role in various protocols from set membership, authentication to (anonymous) credentials systems and cryptocurrency ledgers. However, there is only a small set of underlying cryptographic assumptions on which such accumulator primitives are based. Currently, three main families of accumulators can be distinguished in literature: schemes designed in groups of unknown order [3,2,12,17,22,18,7], schemes designed in groups of known order [21,15,1,13] and hash-based constructions [19,9,10,11,6]. Relevant to this paper are the schemes belonging to the second of these families, where the considered group is a prime order bilinear group. Moreover, when it comes to Dynamic Universal Accumulators (namely those that support dynamic addition and deletion of members and can maintain both membership and non-membership witnesses) we are down to just a few schemes.

In this paper we cryptanalyse one of these universal scheme proposed for bilinear groups, namely the Dynamic Universal Accumulator by Au *et al.* [1], which is zero-knowledge friendly and stood unscathed for 10 years of public scrutiny. This scheme comes in two variants which we called the *α-based construction* and the *$\mathcal{CRS}$-based construction*, respectively. For the first one, we show that the non-membership mechanism, designed to allow for more efficiency on the accumulator manager side, has a subtle cryptographic flaw which enables the adversary to efficiently recover the secret of the accumulator manager given just several hundred to few thousand non-membership witnesses (regardless of the number of accumulated elements).

As a consequence, the attacker can fully break the security of the scheme. Moreover, we show that given only *one* non-membership witness generated with this flawed mechanism, it is possible to efficiently invalidate the assumed collision resistance property of the accumulator by creating a membership witness for a non-accumulated element. Despite the presence of a valid security proof, this is possible because the provided security reduction covers the non-membership mechanism of the $\mathcal{CRS}$-based construction only and it doesn't take into account non-membership definition given for the α-based construction, which, in fact, resulted to be weak.

The second part of the paper investigates the $\mathcal{CRS}$-based variant: motivated by some concrete applications of the scheme where the Manager must have exclusive rights for issuing witnesses (e.g. white-/black-list based authentication mechanisms), we show that an adversary having access to a sufficient amount of witnesses is able to compute valid witnesses for unauthorized elements even when the Accumulator manager keeps secret all the information needed to compute such witnesses, i.e. the $\mathcal{CRS}$. In particular, if the accumulator is initialized by adding m secret elements, an attacker that has access to m non-membership witnesses would succeed in reconstructing the $\mathcal{CRS}$ and will then become able to issue membership and non-membership witnesses for any accumulated and non-accumulated elements, respectively.

In Section 2 we recall both variants of Au *et al.* accumulator scheme along with the security model and our attack scenarios. In Section 3 we detail how collision resistance does not hold when non-membership witnesses are issued accordingly to the α-based construction, while in Section 4 we present our first attack for the α-based construction which allows to fully recover the accumulator's secret α. In Section 4.3 we provide a complexity analysis in terms of time and non-membership witnesses needed and in Section 5 we discuss some further improvements to the α-recovery attack which lead, under different hypothesis, to two new attacks: a *random-y sieving attack* and a *chosen-y sieving attack*, described in Section 5.1 and 5.2, respectively. We implemented all these attacks and we compare, in Section 6, their success probability as a function of the total number of known witnesses needed. We further report another minor design vulnerability for the α-based construction in Section 7. Finally, in Section 8 we investigate the security of the $\mathcal{CRS}$-construction under some concrete attack scenarios and we present, in Section 8.2, the *Witness Forgery Attack* as well as possible countermeasures. A summary of our main contributions can be found in Table 1.

Construction	Ref.	Scenario	Witnesses	Time	Attack Result
α-based	Sec. 4	Random-y	$\mathcal{O}(\log p \log\log p)$	$O(\log^2 p)$	Recovery of α
	Sec. 5.1	Random-y	$\mathcal{O}(\log p \log\log p)$	$O((1+\ell/\log\log p)\log^2 p)$	Recovery of α
	Sec. 5.2	Chosen-y	$\mathcal{O}(\log p)$	$O(\ell \log^2 p/\log\log p)$	Recovery of α
	Sec. 3	Random-y	1	$\mathcal{O}(1)$	Break Collision Resistance
$\mathcal{CRS}$*-based*	Sec. 8.2	Random-y	m	$\mathcal{O}(m^2)$	Issue witnesses

Table 1. Time and non-membership witnesses required in our attacks on the Au *et al.* accumulator scheme for both α-based and $\mathcal{CRS}$-based construction. In this table, p denotes the order of the underlying bilinear group, m denotes the number of (secret) elements with which the accumulator is initialized, ℓ denotes the number of accumulations occurred in between the issues of non-membership witnesses. In the $\mathcal{CRS}$-based construction the $\mathcal{CRS}$ is unknown to the attacker.

2 Au *et al.* Dynamic Universal Accumulator

In their paper, Au and coauthors propose two different constructions for their Dynamic Universal Accumulator, depending on whether information is made available to the accumulator managers. The first requires the accumulator's secret parameter α and is suitable for a centralized entity which efficiently updates the accumulator value and issues witnesses to the users. The second instead, requires a common reference string $\mathcal{CRS}$ and allows to update the accumulator value and to issue witnesses without learning α, but less efficiently. We will refer to the first one as the *α-based* construction, while we will refer to the latter as the *$\mathcal{CRS}$-based* construction.

These two are interchangeable, in the sense that witnesses can be issued from time to time with one or the other construction. Moreover, we note that all operations done with the common reference string $\mathcal{CRS}$, can be done more efficiently by using α directly: hence, if the authority which generates α coincides with the Accumulator Manager, it is more convenient for the latter to always use the secret parameter α to perform operations and thus we will refer to the two constructions mainly to indicate the different defining equations for witnesses (in particular, non-membership witnesses).

We now detail a concrete instance of Au *et al.* accumulator scheme by using Type-I elliptic curves[3]. Where not explicitly stated, each operation refers to both the *α-based* and *$\mathcal{CRS}$-based* constructions.

Generation. Let E be an elliptic curve of embedding degree k over $\mathbb{F}_q$, which is provided with a symmetric bilinear group $\mathbb{G} = (p, G_1, G_T, P, e)$ such that $e : G_1 \times G_1 \to G_T$ is a non-degenerate bilinear map, G_1 is a subgroup of E generated by P, G_T is a subgroup of $(\mathbb{F}_{q^k})^*$ and $|G_1| = |G_T| = p$ is prime. The secret accumulator parameter α is randomly chosen from $\mathbb{Z}/p\mathbb{Z}^*$. The set of accumulatable elements is $\mathcal{ACC} = \mathbb{Z}/p\mathbb{Z} \setminus \{-\alpha\}$.

- ***$\mathcal{CRS}$-based construction.*** Let t be the maximum number of accumulatable elements. Then the common reference string $\mathcal{CRS}$ is computed as

$$\mathcal{CRS} = \{\, P,\ \alpha P,\ \alpha^2 P,\ \dots\ ,\ \alpha^t P \,\}$$

Accumulator updates.

- ***α-based construction.*** For any given set $\mathcal{Y}_V \subseteq \mathcal{ACC}$ let $f_V(x) \in \mathbb{Z}/p\mathbb{Z}[x]$ represent the polynomial

$$f_V(x) = \prod_{y \in \mathcal{Y}_V} (y + x)$$

  Given the secret accumulator parameter α, we say that an accumulator value $V \in G_1$ accumulates the elements in $\mathcal{Y}_V$ if $V = f_V(\alpha)P$.

[3] We note that Au *et al.* accumulator scheme and our attacks as well can be defined to work with any bilinear group.

An element $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$ is added to the accumulator value V, by computing $V' = (y+\alpha)V$ and letting $\mathcal{Y}_{V'} = \mathcal{Y}_V \cup \{y\}$. Similarly, an element $y \in \mathcal{Y}_V$ is removed from the accumulator value V, by computing $V' = \frac{1}{(y+\alpha)}V$ and letting $\mathcal{Y}_{V'} = \mathcal{Y}_V \setminus \{y\}$.

- ***$\mathcal{CRS}$-based construction.*** For any given set $\mathcal{Y}_V \subseteq \mathcal{ACC}$ such that $|\mathcal{Y}_V| \leq t$, let $f_V(x) \in \mathbb{Z}/p\mathbb{Z}[x]$ represent the polynomial

$$f_V(x) = \prod_{y \in \mathcal{Y}_V} (y+x) = \sum_{i=0}^{|\mathcal{Y}_V|} c_i x^i$$

Then, the accumulator value V which accumulates the elements in $\mathcal{Y}_V$ is computed using the $\mathcal{CRS}$ as $V = \sum_{i=0}^{|\mathcal{Y}_V|} c_i \cdot \alpha^i P$.

Witnesses Issuing.

- ***α-based construction.*** Given an element $y \in \mathcal{Y}_V$, the *membership witness* $w_{y,V} = C \in G_1$ with respect to the accumulator value V is issued as

$$C = \frac{1}{y+\alpha}V$$

Given an element $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$, the *non-membership witness* $\bar{w}_{y,V} = (C, d) \in G_1 \times \mathbb{Z}/p\mathbb{Z}$ with respect to the accumulator value V is issued[4] as

$$d = \big(f_V(\alpha) \bmod (y+\alpha) \big) \bmod p, \qquad C = \frac{f_V(\alpha) - d}{y+\alpha} P$$

- ***$\mathcal{CRS}$-based construction.*** Given an element $y \in \mathcal{Y}_V$, let $c(x) \in \mathbb{Z}/p\mathbb{Z}[x]$ be the polynomial such that $f_V(x) = c(x)(y+x)$. Then, the *membership witness* $w_{y,V}$ for y with respect to the accumulator value V is computed using the $\mathcal{CRS}$ as $w_{y,V} = c(\alpha)P$.
Given an element $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$, apply the Euclidean Algorithm to get the polynomial $c(x) \in \mathbb{Z}/p\mathbb{Z}[x]$ and the scalar $d \in \mathbb{Z}/p\mathbb{Z}$ such that $f_V(x) = c(x)(y+x) + d$. Then, the *non-membership witness* $\bar{w}_{y,V}$ for y with respect to the accumulator value V is computed from the $\mathcal{CRS}$ as $w_{y,V} = (c(\alpha)P, d)$.

Witness Update. When the accumulator value changes, users' witnesses are updated accordingly to the following operations:

[4] We assume that here $f_V(\alpha) = \prod_{y \in \mathcal{Y}_V}(y+\alpha)$ is computed over $\mathbb{Z}$. Alternatively, if this computation is done modulo p, then d would be equal to $f_V(\alpha) \mod p$ for a large fraction of elements $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$ and α can be easily recovered by factoring $f_V(x) - d$ over $\mathbb{Z}/p\mathbb{Z}[x]$.

- **On Addition**: suppose that a certain $y' \in \mathcal{ACC} \setminus \mathcal{Y}_V$ is added into V. Hence the new accumulator value is $V' = (y' + \alpha)V$ and $\mathcal{Y}_{V'} = \mathcal{Y}_V \cup \{y'\}$.
  Then, for any $y \in \mathcal{Y}_V$, $w_{y,V} = C$ is updated with respect to V' by computing
  $$C' = (y' - y)C + V$$
  and letting $w_{y,V'} = C'$.
  If, instead, $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$ with $y \neq y'$, its non-membership witness $\bar{w}_{y,V} = (C, d)$ is updated to $\bar{w}_{y,V'} = (C', d \cdot (y' - y)$, where C' is computed in the same way as in the case of membership witnesses.
- **On Deletion**: suppose that a certain $y' \in \mathcal{Y}_V$ is deleted from V. Hence the new accumulator value is $V' = \frac{1}{y'+\alpha}V$ and $\mathcal{Y}_{V'} = \mathcal{Y}_V \setminus \{y'\}$.
  Then, for any $y \in \mathcal{Y}_V$, $w_{y,V} = C$ is updated with respect to V' by computing
  $$C' = \frac{1}{y' - y}C - \frac{1}{y' - y}V'$$
  and letting $w_{y,V'} = C'$.
  If, instead, $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$, its witness $\bar{w}_{y,V} = (C, d)$ is updated to $\bar{w}_{y,V'} = (C', d \cdot \frac{1}{y'-y}$, where C' is computed in the same way as in the case of membership witnesses.

We note that in both cases the added or removed element y' has to be public in order to enable other users to update their witnesses.

Verification. A membership witness $w_{y,V} = C$ with respect to the accumulator value V is valid if it verifies the pairing equation $e(C, yP + \alpha P) = e(V, P)$. Similarly, a non-membership witness $\bar{w}_{y,V} = (C, d)$ is valid with respect to V if it verifies $e(C, yP + \alpha P)e(P, P)^d = e(V, P)$.

2.1 Security Model and Attack Scenarios

The security of the above accumulator scheme is intended in terms of *collision resistance*: in [1], this security property is shown under the t-SDH assumption [5]. Informally, collision resistance ensures that an adversary has negligible probability in forging a valid membership witness for a not-accumulated element and, respectively, a non-membership witness for an already accumulated element. In the following, we briefly recall its formal definition due to Derler et al. and we refer to [16] for more details:

Definition 1. (Collision Freeness [16]) *A cryptographic dynamic universal accumulator is collision-free if for any probabilistic polynomial time adversary $\mathcal{A}$ the following probability*

$$\mathbb{P}\left(\begin{array}{c} (sk_{acc}, pk_{acc}) \leftarrow Gen(1^\lambda) \;\;,\;\; (y, w_y, \bar{w}_y, \mathcal{Y}, V_{\mathcal{Y}}) \leftarrow \mathcal{A}^{\mathcal{O}}(pk_{acc}) \;: \\ (\; \mathrm{Verify}(pk_{acc}, V_{\mathcal{Y}}, w_y, y, \texttt{IsMembWit}) = \texttt{true} \quad \wedge\; y \notin \mathcal{Y} \;) \;\vee \\ (\; \mathrm{Verify}(pk_{acc}, V_{\mathcal{Y}}, \bar{w}_y, y, \texttt{IsNonMembWit}) = \texttt{true} \;\wedge\; y \in \mathcal{Y} \;) \end{array}\right)$$

is a negligible function in the security parameter λ and $\mathcal{O}$ is an oracle returning

- *the accumulator value $V_{\mathcal{Y}}$ resulting from the accumulation of elements of any given input set $\mathcal{Y}$,*
- *the membership witnesses w_{y^*} for any accumulated element y^*,*
- *the non-membership witnesses $\bar{w}_{y^*}$ for any freely chosen non accumulated element y^*.*

By using a secret and public accumulator key pair (sk_{acc}, pk_{acc}), this definition captures the trapdoor nature of Au et al. constructions: in fact, the secret accumulator parameter α corresponds to the formal accumulator secret key sk_{acc}, while pk_{acc} represents the public information, i.e. the bilinear group definition and the group elements needed for public witness verification. Furthermore, due to a result of Vitto and Biryukov [24, Lemma 1], the possibility to arbitrary query the above oracle $\mathcal{O}$ is equivalent to the knowledge of the common reference string $\mathcal{CRS}$, hence both variants can be restated in terms of the above definition and are substantially equivalent in terms of information the attacker has access to.

In next Sections we will show that the non-membership witness definition of the α-based construction is flawed and allows a probabilistic polynomial time attacker to recover the secret accumulator parameter α and thus break collision resistance. This flaw is not present in the non-membership witness definition of the $\mathcal{CRS}$-based construction −which, in fact, fully satisfy the security reduction under the $t-$SDH assumption− and hence the α-based construction can be easily fixed by using, instead, the non-membership witness defining equation of the other $\mathcal{CRS}-$based variant. In other words, a *"fixed"* α-based construction will correspond to a slightly more time-efficient (but asymptotically equivalent) version of the $\mathcal{CRS}-$based construction, where the $\mathcal{CRS}$ is not directly given to the attacker but can be computed in polynomial time [24, Lemma 1].

Motivated by this observation and by concrete applications of the scheme where the attacker cannot arbitrarily query an oracle returning witnesses for any freely chosen element, we show, in Section 8, that even when the Accumulator Manager keeps the $\mathcal{CRS}$ secret, the attacker is be able to efficiently recover it by accessing few non-membership witnesses, thus making him able to issue membership and non-membership witnesses accordingly to the $\mathcal{CRS}-$based defining equations, but not able to break collision resistance for this variant. We remark that this scenario is outside Au et al. security model −where such $\mathcal{CRS}$ is always available to the attacker which can further obtain witnesses from the oracle− but becomes relevant in all those concrete scenarios where the Manager wishes to have exclusive rights for issuing witnesses (and thus keeps the $\mathcal{CRS}$ secret), such us authentication mechanisms where witnesses are used as black-/white-list authentication tokens.

3 Breaking Collision Resistance in the α-based Construction

In the α-based construction, the knowledge of a single non-membership witness is enough to break the (assumed) collision resistance property of the accumulator scheme when the polynomial $f_V(x) \in \mathbb{Z}/p\mathbb{Z}[x]$ is fully known or, equivalently, the set of all accumulated elements is publicly known (which is typically the case).

In the security reduction provided in [1], it is required that given a non-accumulated element $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$ and its non-membership witness $\bar{w}_{y,V} = (C_y, d_y)$ with respect to the accumulator value V, the element $\tilde{d}_y \in \mathbb{Z}/p\mathbb{Z}$ verifies

$$\big(f_V(x) - \tilde{d}_y \bmod (y+x)\big) \equiv 0 \pmod p$$

which in turn corresponds to $\tilde{d}_y \equiv f_V(-y) \pmod p$, a condition enforced by the $\mathcal{CRS}-$based construction non-membership witness definition.

By using, instead, the defining equation for d_y provided in the α-based construction, the partial non-membership witness for y equals $d_y = \big(f_V(\alpha) \bmod (y+\alpha)\big) \bmod p$ and thus

$$d_y \equiv \tilde{d}_y \pmod p \quad \Rightarrow \quad \big(f_V(-y) \bmod (y+\alpha)\big) \equiv f_V(-y) \pmod p$$

holds only when $f_V(-y) < y+\alpha$, i.e. with negligible probability if V accumulates more than one element chosen uniformly at random from $\mathbb{Z}/p\mathbb{Z}$.

Now, if $d_y \not\equiv \tilde{d}_y \bmod p$, we have $f_V(x) - d_y \not\equiv 0 \bmod (y+x)$, and we can use Euclidean algorithm to find a polynomial $c(x) \in \mathbb{Z}/p\mathbb{Z}[x]$ and $r \in \mathbb{Z}/p\mathbb{Z}$ such that $f_V(x) - d_y = c(x)(y+x) + r$ in $\mathbb{Z}/p\mathbb{Z}[x]$. Then, by recalling that $C_y = \frac{f_V(\alpha)-d_y}{y+\alpha}P$, under the $t-$SDH assumption, the attacker uses the available $\mathcal{CRS} = \{P, \alpha P, \ldots, \alpha^t P\}$ to compute $c(\alpha)P$ and obtains a membership witness with respect to V for an arbitrary non accumulated element y as

$$C_y + \frac{d_y}{r}\big(\, C_y - c(\alpha)P\,\big) = C_y + \frac{d_y}{r}\left(\, C_y - C_y - \frac{r}{y+\alpha}P\right) = \frac{f_V(\alpha)}{y+\alpha}P = \frac{1}{y+\alpha}V$$

thus breaking the assumed collision resistance property. We note that this result doesn't invalidate the security proof provided by Au *et al.* in [1]: indeed, the reduction to the t-SDH assumption is shown for (non-membership) witnesses generated accordingly to the $\mathcal{CRS}$-based construction only, and thus, collision resistance can be guaranteed only for this latter construction.

We speculate that this flaw comes from the wrong assumption that

$$\big(f_V(x) \bmod (y+x)\big) \equiv \big(f_V(\alpha) \bmod (y+\alpha)\big) \pmod p$$

which, if true, would have implied security of non-membership witnesses issued accordingly to the α-based construction as well. The authors also declare [1, Section 2.2] that by using the secret accumulator value α, the Accumulator Manager can compute membership and non-membership witnesses in $O(1)$ time: this clearly cannot be true, since, regardless of the variant considered, the evaluation of the polynomial $f_V(x)$ and its reduction modulo a $\sim \log p$-bits integer requires (at least) $O(\deg f_V)$ time.

In the next Sections we will show that within the α-based construction, an attacker can efficiently recover the secret accumulator parameter α by accessing multiple non-membership witnesses, thus making him able to break collision resistance by computing membership witnesses for non-accumulated elements similarly as above, but also non-membership witnesses for accumulated elements.

4 The α-Recovery Attack for the α-Based Construction

From now on, we assume the secret parameter α and the accumulator value V along with the set of currently accumulated elements $\mathcal{Y}_V$ to be fixed.

The following attack on the α-based construction consists of two phases: the retrieval of the value $f_V(\alpha) \in \mathbb{Z}$ used to compute non-membership witnesses modulo many small primes and the full recovery of the accumulator secret parameter α.

4.1 Recovering $f_V(\alpha)$

Let $d_y = \big(f_V(\alpha) \bmod (y+\alpha)\big) \bmod p$ be a partial non-membership witness with respect to V for a certain element $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$, and let $\tilde{d}_y$ denote the integer $f_V(\alpha) \bmod (y+\alpha)$. We then have $d_y = \tilde{d}_y \bmod p$, and we are interested in how often d_y equals $\tilde{d}_y$ as integers. Attacker benefits from the cases when $y+\alpha < p$, since the reduction modulo p does nothing and $d_y = \tilde{d}_y$ for all y.

The worst case happens when α is maximal, i.e. $\alpha = p-1$. Indeed, in this case, if $y = 0$ then $y+\alpha < p$ and $d_y = \tilde{d}_y$ with probability 1; if instead $y > 0$ and $y \neq p-\alpha = 1$ the probability that $d_y = \tilde{d}_y$ is $\frac{p}{y+\alpha}$ and, hence, is minimal when compared to smaller values of α. Thus, with $\alpha = p-1$ the probability that d_y equals $\tilde{d}_y$ as integers ranges from 1 (when $y=0$) to almost $1/2$ (when $y = p-1$). Assuming that y is sampled uniformly at random, we can obtain the following lower bound on the probability (for arbitrary α):

$$\begin{aligned}
\underset{\substack{y\in\{0,\dots,p-1\}\\ y\neq p-\alpha\\ f_V(\alpha)\in\mathbb{Z}}}{\mathbb{P}}(d_y = \tilde{d}_y) &\geq \frac{1}{p-1}\left(1 + p\sum_{\tilde{y}=2}^{p-1}\frac{1}{\tilde{y}+p-1}\right)\\
&= \frac{p}{p-1}\left(\sum_{i=1}^{2p-2}\frac{1}{i} - \sum_{i=1}^{p-1}\frac{1}{i}\right) = \frac{p}{p-1}\left(H_{2p-2} - H_{p-1}\right)\\
&= \left(1+\frac{1}{p-1}\right)\cdot\left(\ln 2 - \frac{1}{4(p-1)} + o\left(p^{-1}\right)\right)\\
&= \ln 2 + \frac{4\ln 2 - 1}{4(p-1)} + o(p^{-1})\\
&> \ln 2. \qquad (1)
\end{aligned}$$

where H_n denotes the $n-$th Harmonic number, and the last inequality holds for all values of p used in practice.

Assume that $q|(y+\alpha)$ for a small prime $q \in \mathbb{Z}$ such that $q \ll y+\alpha$. If $d_y = \tilde{d}_y$ we have $f_V(\alpha) \equiv d_y \pmod{q}$ with probability 1, otherwise it happens with probability 0 since then $f_V(\alpha) \equiv d_y + p \pmod{q}$. If instead $q \nmid (y+\alpha)$, we assume $d_y \mod q$ to be random in $\mathbb{Z}/q\mathbb{Z}$ and thus $f_V(\alpha) \equiv d_y \pmod{q}$ happens with probability close to $\frac{1}{q}$.

More precisely,

$$\mathbb{P}\big(\, f_V(\alpha) \equiv d_y \pmod{q}\,\big) > \ln 2 \cdot \frac{1}{q} + \frac{q-1}{q^2} = \frac{(\ln 2 + 1)q - 1}{q^2} \tag{2}$$

while for any other $c \in \mathbb{Z}/q\mathbb{Z}$ such that $c \not\equiv d_y \pmod{q}$ we have

$$\mathbb{P}\big(\, f_V(\alpha) \equiv c \pmod{q}\,\big) < (1 - \ln 2) \cdot \frac{1}{q} + \frac{q-1}{q^2} = \frac{(2 - \ln 2)q - 1}{q^2} \tag{3}$$

In other words, the value $d_y \bmod q$ has a higher chance to be equal to $f_V(\alpha) \bmod q$ compared to any other value in $\mathbb{Z}/q\mathbb{Z}$.

We will use this fact to deduce $f_V(\alpha)$ modulo many different small primes. More precisely, suppose that an attacker has access to the elements $y_1, \dots, y_n$ together with the respective partial non-membership witnesses

$$d_{y_i} \equiv \big(f_V(\alpha) \mod (y_i + \alpha)\big) \bmod p$$

If q is a small prime and n is sufficiently large (see Section 4.3 for the analysis), $f_V(\alpha) \bmod q$ can be deduced by simply looking at the most frequent value among

$$d_{y_1} \bmod q,\ \dots\ ,\ d_{y_n} \bmod q$$

Once we compute $f_V(\alpha)$ modulo many different small primes $q_1, \dots, q_k$ such that $q_1 \cdot \ldots \cdot q_k > p$, we can proceed with the next phase of the attack: the full recovery of the secret parameter α.

4.2 Recovering $\boldsymbol{\alpha}$

If the discrete logarithm of any accumulator value is successfully retrieved modulo many different small primes whose product is greater than p, α can be recovered with (virtually) no additional partial non-membership witnesses. The main observation we will exploit is the following:

Observation 1. *Let q be an integer and let $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$ be a non-accumulated element such that its partial non-membership witness with respect to V satisfies $d_y = \tilde{d}_y$. Then $d_y \not\equiv f_V(\alpha) \pmod{q}$ implies that $q \nmid (y + \alpha)$, or, equivalently, $\alpha \not\equiv -y \pmod{q}$.*

From (1) it follows that for any given $q \in \mathbb{Z}$ and non-accumulated element y such that $(f_V(\alpha) - d_y) \not\equiv 0 \pmod{q}$, we have

$$\mathbb{P}\big(\alpha \not\equiv -y \pmod{q} \,|\, f_V(\alpha) \not\equiv d_y \pmod{q}\big) > 1 - \frac{(1 - \ln 2)q}{q^2 - (1 + \ln 2)q + 1} \approx 1 - \frac{1 - \ln 2}{q}$$

By considering all available non-membership witnesses, if q is small and n is sufficiently larger than q (see Section 4.3), we can deduce $\alpha \bmod q$ as the element in $\mathbb{Z}/q\mathbb{Z}$ which is the least frequent −or not occurring at all− among the residues

$$-y_{i_1} \bmod q\ ,\ \ \ldots\ \ ,\ -y_{i_j} \bmod q$$

such that $(f_V(\alpha) - d_{y_{i_k}}) \not\equiv 0 \bmod q$ for all $k = 1, \ldots, j$.

It follows that, if $q_1, \ldots, q_k$ are small primes such that $q_1 \cdot \ldots \cdot q_k > p$, from the values $f_V(\alpha) \bmod q_i$ −computed according to Section 4.1− and the values $\alpha \bmod q_i$, with $i \in [1, k]$, $\alpha \in \mathbb{Z}$ can be obtained by using the Chinese Remainder Theorem.

4.3 Estimating the minimum number of witnesses needed

We now give an asymptotic estimate of the minimum number of non-membership witnesses needed so that both phases of the above attack succeed with high probability. We will use the multiplicative Chernoff bound, which we briefly recall.

Theorem 2. (Chernoff Bound) *Let $X_1, \ldots, X_n$ be independent random variables taking values in $\{0,1\}$ and let $X = X_1 + \ldots + X_n$. Then, for any $\delta > 0$*

$$\mathbb{P}\big(\,X \le (1-\delta)\mathbb{E}[X]\,\big) \le e^{-\frac{\delta^2\mu}{2}} \qquad 0 \le \delta \le 1$$
$$\mathbb{P}\big(\,X \ge (1+\delta)\mathbb{E}[X]\,\big) \le e^{-\frac{\delta^2\mu}{2+\delta}} \qquad 0 \le \delta$$

Proof. See [20, Theorem 4.4, Theorem 4.5]. □

Our analysis will proceed as follows: first, we introduce two random variables to model, for a given small prime q, the behaviour of the values $f_V(\alpha) \bmod q$. Then, we will use Chernoff bound to first estimate the probability of wrongly guessing $f_V(\alpha) \bmod q$, and then deduce a value for n so that such probability is minimized for all primes q considered in the attack.

Let $q \in \mathbb{Z}$ be a fixed prime and let X_g be a random variable which counts the number of times $f_V(\alpha) \bmod q$ is among the values $d_1 \bmod q, \ldots, d_n \bmod q$. Similarly, let X_b be a random variable which counts the number of times a certain residue $t \in \mathbb{Z}/q\mathbb{Z}$ not equal to $f_V(\alpha) \bmod q$ is among the values $d_1 \bmod q$, …, $d_n \bmod q$. Then

$$\mathbb{E}[X_g] = n \cdot \frac{(\ln 2 + 1)q - 1}{q^2} \approx (\ln 2 + 1)\frac{n}{q}$$

$$\mathbb{E}[X_b] = n \cdot \frac{(2 - \ln 2)q - 1}{q^2} \approx (2 - \ln 2)\frac{n}{q}$$

By applying Theorem 2, we can estimate the probability that X_g and X_b crosses $\frac{\mathbb{E}[X_g]+\mathbb{E}[X_b]}{2} = \frac{3n}{2q}$ as

$$\mathbb{P}\left(X_g \le \frac{3n}{2q}\right) = \mathbb{P}\left(X_g \le \left(1 - \frac{2\ln 2 - 1}{2\ln 2 + 2}\right)\mathbb{E}[X_g]\right) < e^{-\frac{n}{91q}} \doteq e_{q,g}$$

$$\mathbb{P}\left(X_b \geq \frac{3n}{2q}\right) = \mathbb{P}\left(X_b \geq \left(1 + \frac{2\ln 2 - 1}{4 - 2\ln 2}\right)\mathbb{E}[X_b]\right) < e^{-\frac{n}{76q}} \doteq e_{q,b}$$

and we minimize these inequalities by requiring that

$$1 - (1 - e_{q,g})(1 - e_{q,b})^{q-1} \approx e_{q,g} + (q-1)e_{q,b} \doteq s_q$$

is small for each prime q considered in this attack phase. Thus, if $q = max(q_1, \dots, q_k)$, we can bound the sum

$$\sum_{i=1}^{k} s_{q_i} \leq q s_q = q(e^{-\frac{n}{91q}} + (q-1)e^{-\frac{n}{76q}}) \approx e^{-\frac{n}{91q} + \log q} + e^{-\frac{n}{76q} + 2\log q}$$

and we make it small by taking $n = O(q \log q)$.

In order to apply the Chinese Remainder Theorem for the full recovery of α we need that $q_1 \cdot \ldots \cdot q_k > p$. If $q_1, \dots, q_k$ are chosen to be the first k primes, we can use an estimation for the first Chebyshev function growth rate to obtain $\ln(q_1 \cdot ... \cdot q_k) = (1 + o(1)) \cdot k \ln k \sim q_k$ by Prime Number Theorem and thus $q_k > \ln p$. We then conclude that

$$n = O(\log p \log \log p)$$

non-membership witnesses are enough to recover $f_V(\alpha) \bmod q_1 \cdot \ldots \cdot q_k$ with high probability.

We note that by using Chernoff bound in order to estimate the minimum number of witnesses needed to recover α, it can be shown, similarly as done above for $f_V(\alpha)$, that $O(\log p \log \log p)$ non-membership witnesses are enough to identify with high probability $\alpha \bmod q_1 \cdot \ldots \cdot q_k = \alpha$.

The time complexity is dominated by

$$(\#\text{ primes } q) \times (\#\text{ witnesses}) = O\left(\frac{\log p}{\log \log p}\right) \times O(\log p \log \log p)$$

which is equal to $O(\log^2 p)$.

5 Improving the α-Recovery Attack

We will now improve the α-Recovery Attack outlined in Section 4 by giving some variants under two different attack scenarios, depending on whether the attacker has access to non-membership witnesses for *random-y* or *chosen-y*.[5] These improvements will further reduce the number of non-membership witnesses needed to fully recover the secret accumulator parameter α to a small multiple of $\log p$.

The main idea behind the improved attack is to keep removing wrong candidates for $\alpha \bmod q$ for small primes q (*sieving*), until only the correct one is left. As in the previous attack, full value of α is then reconstructed using the Chinese Remainder Theorem.

[5] We observe that according to Definition 1, the attacker has access to an oracle which returns witnesses for any *chosen-y*. However, in concrete instances of the accumulator scheme, an attacker might have access only to witnesses for *random* values y.

Collecting Witnesses Issued at Different States In the α-Recovery Attack described in Section 4, $O(\log p \log\log p)$ non-membership witnesses issued with respect to the same accumulator value V are needed in order to fully recover α. In the following attacks we drop this condition and allow non-membership witnesses to be issued with respect to different accumulator values $f_1(\alpha)P = V_1, \ldots, f_\ell(\alpha)P = V_\ell$, but we require that no deletions occur between the accumulator states V_1 and V_ℓ. In this case, since the sequence of elements added must be public to permit witness updates, we have that the polynomial functions $g_{i,j}(x) \in \mathbb{Z}/p\mathbb{Z}$ such that $f_j(\alpha) = g_{i,j}(\alpha)f_i(\alpha)$ for any $\alpha \in \mathbb{Z}/p\mathbb{Z}$, can be publicly computed for any $i, j \in [1, \ell]$. It follows that, given a small prime q, once $\alpha \bmod q$ and $f_i(\alpha) \bmod q$ for some $i \in [1, \ell]$ are correctly computed, $f_j(\alpha) \bmod q$ can be computed as $g_{i,j}(\alpha)f_i(\alpha) \bmod q$ for any $j \in [1, \ell]$ such that $j > i$.

The requirement that no deletion operation should occur if the collected witnesses were issued at different states, comes from the fact that the accumulator can be initialized by accumulating some values which are kept secret by the Accumulator Manager.

It follows that, whenever the polynomial $f_1(x) \in \mathbb{Z}/p\mathbb{Z}$ is publicly known (or, equivalently, the set of all accumulated elements $\mathcal{Y}_{V_1}$) for a certain accumulator value V_1, we can remove the condition that no later deletion operations occur during attack execution, since the knowledge of $\alpha \bmod q$ is enough to compute $f_i(\alpha) \bmod q$ for any $i \in [1, \ell]$. Thus any non-membership witnesses issued from V_1 on can be used to recover α.

Removing reduction modulo *p*. We show that, under some practical assumptions, it is possible to eliminate with high probability the noise given by the reduction modulo p performed by the Accumulator Manager when he issues a non-membership witness. That is, we recover $\tilde{d}_{y_i} = f_{V_j}(\alpha) \mod (y_i + \alpha)$ for a large fraction of pairs (y_i, V_j), given the partial non-membership witnesses $d_{y_i} = \big(f_{V_j}(\alpha) \mod (y_i + \alpha)\big) \bmod p$ collected with respect to different accumulator values V_j with $j > 1$.

Aiming at this, we first observe that from the fact that $0 \le y, \alpha < p$ for any given $y \in \mathcal{ACC} \setminus \mathcal{Y}_V$, the partial non-membership witness d_y for y with respect to V can be expressed in terms of $\tilde{d}_y$ in one of the following way:

(1) $d_y = f_V(\alpha) \mod (y + \alpha) = \tilde{d}_y$,
(2) $d_y = \big(f_V(\alpha) \mod (y + \alpha)\big) - p = \tilde{d}_y - p$.

Since p is odd, whenever $y + \alpha$ is even, these two cases can be easily distinguished modulo 2: indeed, in the first case $d_y \equiv f_V(\alpha) \pmod 2$, while in the second case $d_y \not\equiv f_V(\alpha) \pmod 2$.

This observation effectively allows to correctly compute $\tilde{d}_y$ half of the times given a correct guess for $\alpha \bmod 2$ and $f_V(\alpha) \bmod 2$. Indeed, given a set of partial non-membership witnesses $d_{y_1}, \ldots, d_{y_n}$ with respect to V, each guess of $\alpha \bmod 2$ and $f_V(\alpha) \bmod 2$ will split the witnesses in two subsets, namely one where the corresponding elements y_i satisfy $y_i + \alpha \equiv 0 \pmod 2$ (and thus $\tilde{d}_{y_i}$ can be correctly recovered), and the other where this doesn't happen.

Checking if $\alpha \bmod 2$ and $f_V(\alpha) \bmod 2$ were actually correct guesses can be done observing how the attacks described in Section 5.1 and 5.2 (or in Section 4 if witnesses are issued with respect to the same accumulator value) behaves with respect to the subset of witnesses that permitted to recover the values $\tilde{d}_{y_i}$. In case of a wrong guess, indeed, it will not possible to distinguish α and $f_{V_i}(\alpha)$ modulo some different small primes q: in this case the attack can be stopped and a new guess should be considered. On the other hand, a correct guess will permit to correctly recover α and $f_{V_i}(\alpha)$ modulo few more primes q greater than 2. Since, whenever $\alpha \bmod q$ and $f_V(\alpha) \bmod q$ are known, $\tilde{d}_y$ can be correctly recovered, analogously to the modulo 2 case, for all those y such that $y + \alpha$ is divisible by q, this implies that it is possible to iteratively recover more and more correct values $\tilde{d}_{y_i}$ given the initial set of considered witnesses.

Repeating this procedure for small primes q up to r, it allows to recover $\tilde{d}_{y_i}$ for those y_i that are divisible by at least one prime not exceeding r. This fraction tends to $1 - \varphi(r\#)/(r\#)$ as y_i tend to infinity, where φ is the Euler's totient function and $r\#$ denotes the product of all primes not exceeding r. For example, setting $r = 101$ allows to recover $\tilde{d}_{y_i}$ for about 88% of all available witnesses. We conclude that $\tilde{d}_{y_i}$ can be recovered for practically all $i \in [1, n]$.

In the case where witnesses are issued with respect to different accumulator values $V_1, \ldots, V_\ell$, as remarked above, the knowledge of $\alpha \bmod q$ and $f_{V_1}(\alpha) \bmod q$ allows to compute $f_{V_j}(\alpha) \mod q$ for all V_j with $j > 1$, so the modulo p noise reduction can be easily performed independently on when the witnesses are issued.

5.1 The Random-y Sieving Attack

In this scenario we assume that all elements y_i for which the partial non-membership witnesses d_{y_i} are available to the adversary, are sampled uniformly at random from $\mathbb{Z}/p\mathbb{Z}$. Furthermore these witnesses are pre-processed accordingly to the method described above, in order to eliminate the noise given by reduction modulo p.

Recovering $\boldsymbol{\alpha}$ mod $\boldsymbol{q}$. Let q be a small prime, i.e. $q = O(\log p)$, and let $\mathcal{Y}_\alpha$ be the set containing all pairs $(y_i, \tilde{d}_{y_i})$ such that $y_i + \alpha \equiv 0 \pmod q$ for a certain guess $\alpha \mod q$. If the latter is guessed wrongly, then the values $\tilde{d}_{y_i}$ modulo q are distributed uniformly and independently from the values $f_{V_i}(\alpha) \mod q$. On the other hand, if the guess is correct, then $\tilde{d}_{y_i} \equiv f_{V_i}(\alpha) \pmod q$.

Even in the case when $f_{V_1}(\alpha) \mod q$ is unknown, $f_{V_i}(\alpha) \mod q$ can be recovered from the first occurrence of y_i in the set $\mathcal{Y}_\alpha$ and verified at all further occurrences, since all $f_{V_j}(\alpha) \mod q$ can be computed for any $j \geq i$. It follows that we can easily distinguish if a guess for $\alpha \mod q$ is either correct or not.

The attack succeeds if for every wrong guess $\alpha^\times$ of $\alpha \mod q$ we observe a contradiction within the pairs in $\mathcal{Y}_{\alpha^\times}$. It's easy to see that if $|\mathcal{Y}_{\alpha^\times}| = t$, the probability to observe at least one contradiction is $1 - 1/q^{t-1}$. Thus, by ensuring a constant number t of elements in $\mathcal{Y}_{\alpha^\times}$ given each $\alpha^\times \neq \alpha \mod q$ is sufficient to make the probability of false positives negligible. This requires availability of $O(q \log q)$ witnesses in total.

Recovering $\boldsymbol{\alpha}$. The final step is the same as in the previous attacks: the secret value α is recovered by repeating the process for different small primes q and then by applying the Chinese Remainder Theorem. Furthermore, if for some primes q there are multiple candidates of $\alpha \mod q$, such primes can be simply omitted from the application of the Chinese Remainder Theorem. In this case, in order to fully recover $\alpha \in \mathbb{Z}$, the maximum prime q that has to be considered must be larger than $\ln p$ by a constant factor. We conclude that $O(q \log q) = O(\log p \log \log p)$ witnesses are sufficient for full recovery of α with overwhelming probability.

The time complexity of the attack is dominated by guessing $\alpha \mod q$ for each q considered. Note that for a wrong guess of $\alpha \mod q$, we can expect on average a constant amount of witnesses to check before an inconsistency is observed; this amount is thus enough to identify the correct value. For each such guess, nearly all accumulator states in the history have to be considered in order to take into account all additions to the accumulator. However, the non-membership witnesses issued in each state can be classified by guesses of $\alpha \mod q$ in a single scan for each prime q.

We conclude that the time complexity is dominated by

$$(\#\text{ primes } q) \times (q \text{ guesses of } \alpha \bmod q) \times (\#\text{ of accumulator states})$$

and by classifying all non-membership witnesses for each prime q

$$(\#\text{ primes } q) \times (\#\text{ witnesses})$$

The final complexity is $O((1 + \ell/\log \log p) \log^2 p)$.

5.2 The Chosen-y Sieving Attack

If the adversary is allowed to choose the elements y_i for which the partial non-membership witnesses are issued, no matter with respect to which accumulator state, the amount of required witnesses can be further reduced by a $\log \log p$ factor.

First, we assume that the adversary chooses the elements y_i non-adaptively, i.e. before the accumulator is initialized. The idea is simply to use consecutive values, that is $y_0 = r$, $y_1 = r + 1, \ldots, y_i = r + i, \ldots$, for some $r \in \mathbb{Z}/p\mathbb{Z}$. This choice fills equally all sets $\mathcal{Y}_{\tilde{\alpha}}$ for all $\tilde{\alpha} \in \mathbb{Z}/q\mathbb{Z}$ and small q, where $\tilde{\alpha}$ represents either a correct guess for $\alpha \bmod q$ or a wrong guess $\alpha^{\times}$. As a result, $t = O(q)$ elements are enough to make the size of each set $\mathcal{Y}_{\tilde{\alpha}}$ at least equal to t. The full total number of required non-membership witnesses is then reduced to $O(q) = O(\log p)$. The time complexity then is improved by a factor $\log \log p$ in the case when ℓ is small: $O(\ell \log^2 p/ \log \log p)$.

We now consider the case when the adversary can adaptively chose the elements y_i. Note that, on average, we need only $2 + 1/(q-1)$ elements in each set $Y_{\alpha^\times}$ to discard the wrong guess of $\alpha \mod q$, for all q. The adaptive choice allows to choose y_i such that $(y_i + \alpha^\times) \equiv 0 \pmod q$ specifically for those $\alpha^\times$ which are not discarded yet. Furthermore, the Chinese Remainder Theorem allows us to combine such adaptive queries for all chosen primes q simultaneously. As a result, approximately $2 \ln p$ witnesses for adaptively chosen elements are sufficient for the full recovery of α. This improves the constant factor of the non-adaptive attack in term of number of non-membership witnesses required.

Remark 1. As described at the beginning of this Section, non-membership witnesses can be issued with respect to different successive accumulator values $V_1, \ldots, V_\ell$, within which no deletion operation occurs. If the value $f_{V_1}(x) \in \mathbb{Z}[x]$ is known to the adversary (or equivalently the set of all accumulated elements in V_1), only $\ln p$ non-membership witnesses issued for adaptively chosen elements are sufficient to recover α. In this case, indeed, instead of verifying uniqueness of elements in the set $\mathcal{Y}_{\alpha^\times}$, we can directly compare our guess to the value $f_{V_j}(\alpha) \bmod q$ given from $f_{V_1}(\alpha)$, thus requiring $1 + 1/(q-1)$ elements on average.

6 Experimental Results

We implemented the *α-Recovery Attack* from Section 4 and both the *random-y* and the *non-adaptive chosen-y sieving* attacks from Sections 5.1 and 5.2.

For the verification purpose we used a random 512-bit prime p. We measured the success rate of the attacks with respect to the number of available non-membership witnesses. The α-Recovery Attack applies to a single accumulator state, and for the sieving attacks, the number of state changes of the accumulator was 10 times less than the number of issued witnesses. The initial state of the accumulator in all attacks was assumed to be secret. Each attack was executed 100 times per each analyzed number of available non-membership witnesses. The sieving attacks were considered successful if at most 2^{10} candidates for α were obtained and the correct α was among them. The results are illustrated in Figure 1.

The α-Recovery Attack, while being simple, requires a significant amount of witnesses to achieve a high success rate, more than $20000 \approx 10 \ln p \ln \ln p$ witnesses and finishes in less than 5 seconds. The random-y sieving attack achieves almost full success rate with about $6000 \approx 3 \ln p \ln \ln p$ available witnesses and completes in less than 10 seconds. The chosen-y sieving attack requires less than $2000 \approx 4 \ln p$ witnesses to achieve almost perfect success rate and completes in less than 4 seconds. All timings include the generation of witnesses. The experiments were performed on a laptop with Linux Mint 19.3 OS and an Intel Core i5-10210U CPU clocked at 1.60GHz.

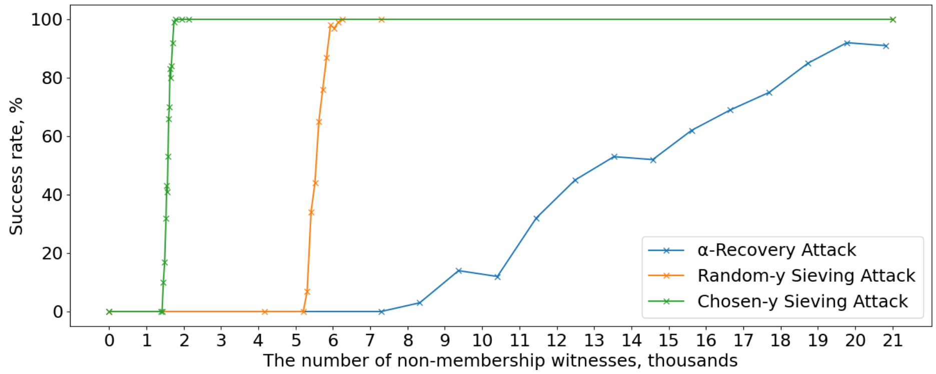

Fig. 1. Attacks experimental success rate as a function of the total number of available witnesses.

7 Weak Non-membership Witnesses

In the α-based construction, non-membership witness definition is affected by another minor design vulnerability: given a non-membership witness $\bar{w}_{y,V} = (C_y, d_y)$ with respect to an accumulator value V, if $d_y \equiv f_V(\alpha) \bmod p$, then $C_y = O$.

Those *"weak non-membership witnesses"* are issued with non-negligible probability in the security parameter λ when only one element is accumulated. Assume, indeed, that $V = (y' + \alpha)P$ for a certain element $y' \in \mathcal{ACC}$. Then, for any element $y \in \mathcal{ACC}$ such that $y' < y$, the corresponding non-membership witness $\bar{w}_{y,V}$ with respect to V is issued as

$$d_y = \big(y' + \alpha \bmod (y + \alpha)\big) \bmod p = (y' + \alpha) \bmod p$$

and thus $C_y = O$. In this case, as soon as the element y' becomes public (e.g. is removed), the accumulator secret parameter can be easily obtained as $\alpha = (d_y - y') \bmod p$.

8 Preventing Witness Forgery in the $\mathcal{CRS}$-based Construction

All the attacks we have presented so far are ineffective when witnesses (more precisely, non-membership witnesses) are issued according to the defining equations given for the $\mathcal{CRS}$-based construction.

We note that the knowledge of the $\mathcal{CRS}$ is functionally equivalent to the knowledge of α when the set of currently accumulated elements is fully known: indeed, besides accumulator updates, the $\mathcal{CRS}$ permits to issue both membership and non-membership witnesses for arbitrary elements, with the difference that the knowledge of α permits to break collision-resistance, while the knowledge of the $\mathcal{CRS}$ does not. Furthermore, despite what we saw in Section 3, witnesses definition in the $\mathcal{CRS}$-based construction satisfy the hypothesis for the $t-$SDH security reduction provided by Au *et al.*, i.e. collision-resistance is enforced when the $\mathcal{CRS}$ is used to issue witnesses.

Depending on the use-case application of the accumulator scheme, the possibility to publicly issue witnesses for arbitrary elements could be undesirable: for example, this is relevant when the accumulator scheme is used as a privacy-preserving authorization mechanism, i.e. an Anonymous Credential System. Suppose, indeed, that in this scenario the accumulator value V accumulates revoked users' identities and the non-revoked ones authenticate themselves showing the possession of a valid non-membership witness $\bar{w}_{y,V}$ for an identity y, both issued by a trusted Authentication Authority. If an attacker has access to the $\mathcal{CRS}$, he will be able to forge a random pair of credentials $(y', w_{y',V})$ and then he could authenticate himself, even if the Authentication Authority never issued the identity y' nor the corresponding witness. This is especially the case when a zero knowledge protocol is instantiated during users' credentials verification since it is impossible to distinguish between a zero knowledge proof for an authorized identity y and a proof for the never issued, but valid, identity y'.

In the following we will investigate the $\mathcal{CRS}$-based construction under this scenario, i.e. assuming the Accumulator Manager to be the only authority allowed to issue witnesses. We stress that resistance to witness forgeries is outside the security model provided by Au *et al.*, where the attacker can generate as many witnesses as he wishes, and the attacks described in the following do not break any security properties assumed for the $\mathcal{CRS}-$based construction by the respective authors.

In the next two Sections, we will discuss how witness forgery for never-authorized elements can be prevented, namely: a) the manager constructs the set $\mathcal{Y}_V$ of currently accumulated elements in such a way that it is infeasible to fully reconstruct it; b) the common reference string $\mathcal{CRS}$ is not published and an attacker cannot reconstruct it.

8.1 How to ensure some accumulated elements remain unknown

Given an accumulator value V, assume $\mathcal{Y}_V$ is the union of the disjoint sets $\mathcal{Y}_{V_0}$, whose elements are used exclusively to initialize the accumulator value from P to V_0, and $\mathcal{Y}_{id} = \mathcal{Y}_V \setminus \mathcal{Y}_{V_0}$, the set of currently accumulated elements for which a membership witness have been issued.

Since the elements in $\mathcal{Y}_{id}$ must be public to enable users to update their witnesses[6], the reconstruction of $\mathcal{Y}_V = \mathcal{Y}_{V_0} \cup \mathcal{Y}_{id}$ can be prevented only if $\mathcal{Y}_{V_0}$ remains, at least partially, unknown.

[6] The very first element for which a membership witness is issued can remain unknown if there are no other users which need to update their witnesses. In this case, we assume that this elements belongs to $\mathcal{Y}_0$.

From $\mathcal{Y}_V = \mathcal{Y}_{V_0} \cup \mathcal{Y}_{id}$ and $\mathcal{Y}_{V_0} \cap \mathcal{Y}_{id} = \emptyset$, it follows that the polynomial $f_V(x)$ can be written as

$$f_V(x) = f_0(x) \cdot f_{id}(x) = \prod_{y_i \in \mathcal{Y}_{V_0}} (y_i + x) \prod_{y_j \in \mathcal{Y}_{id}} (y_j + x)$$

When non-membership witnesses are generated according to the $\mathcal{CRS}$-construction, as soon as an attacker has access to $deg(f_{id}) \geq deg(f_0)$, $|\mathcal{Y}_{V_0}|$ partial non-membership witnesses for the elements $y_1, \ldots, y_{|\mathcal{Y}_{V_0}|}$, i.e.

$$d_{y_i} \equiv f_V(-y_i) \equiv f_0(-y_i) \cdot f_{id}(-y_i) \pmod p$$

he will be able to reconstruct the unknown set $\mathcal{Y}_{V_0}$. Indeed, with the knowledge of $\mathcal{Y}_{id}$, the polynomial $f_{id}(x)$ can be easily obtained and it is then possible to compute the $|\mathcal{Y}_{V_0}|$ pairs

$$\Big(-y_i, f_0(-y_i) \Big) = \Big(-y_i, \frac{d_{y_i}}{f_{id}(-y_i)} \Big)$$

With these pairs, the attacker is able to uniquely interpolate, using for example Lagrange interpolation, the monic polynomial $f_0(x) \mod p$ whose roots are the elements in $\mathcal{Y}_{V_0}$.[7]

The reconstruction of the set $\mathcal{Y}_V$ can be prevented by initializing the accumulator with a number of random elements which is greater than the total number of issuable non-membership witnesses: this clearly avoids the possibility to interpolate $f_0(x)$, even in the case when the attacker has access to all issued non-membership witnesses.

We note, however, that this approach has some disadvantages. First of all, the maximum number of issuable non-membership witnesses has to be set at generation time and cannot be increased once the first witness is issued, since all further accumulated elements will be public to allow witness updates. When this number is reasonable big, let's say 1 billion, the Accumulator Manager needs to evaluate at least a 1-billion degree polynomial when issuing any new non-membership witnesses, an operation that becomes more and more expensive as the number of accumulated elements increases. On the other hand, by decreasing it, the Accumulator Manager can issue the non-membership witnesses in a less expensive way, but only to a smaller set of users.

8.2 Recovering the $\mathcal{CRS}$

Alternatively to the countermeasure proposed in Section 8.1, it's natural to wonder if unauthorized witness forgery can be prevented by just keeping the $\mathcal{CRS}$ secret from the attacker.

We will now show that by executing what we will refer to as *The Witness Forgery Attack*, an attacker that has access to multiple witnesses can successfully recover the $\mathcal{CRS}$, even if the Accumulator Manager keeps it secret.

[7] Since $f_0(x)$ is monic, only $deg(f_0)$ evaluations are needed to uniquely interpolate it.

The main observation on which this attack is based on is that given any partial witness C_y (no matter if it is a membership or a non-membership one) for an element y with respect to the accumulator value V, it can be expressed as $C_y = g_y(\alpha)P$ for a polynomial $g_y(x) \in \mathbb{Z}/p\mathbb{Z}[x]$ which depends on y and $f_V(x)$ (i.e. $f_V(x) = g_y(x)(y+x) + d_y$ for some $d_y \in \mathbb{Z}/p\mathbb{Z}$).

Assume the attacker has access to $n \geq |\mathcal{Y}_V| = m$ partial non-membership witnesses

$$C_{y_1} = g_1(\alpha)P, \ldots, C_{y_n} = g_n(\alpha)P$$

with respect to V. From Section 8.1, we know that he is able to fully recover the polynomial $f_V(x)$ and so he can explicitly compute from the elements $y_1, \ldots, y_n$ the n polynomials $g_1(x), \ldots, g_n(x)$ in $\mathbb{Z}/p\mathbb{Z}[x]$, each of degree $m-1$. We note that by randomly choosing m out of these n polynomials, they will be linearly independent with probability

$$\frac{1}{p^{m^2}} \cdot \prod_{k=0}^{m-1}(p^m - p^k) = \prod_{k=1}^{m}\left(1 - \frac{1}{p^k}\right) \approx 1$$

and so we assume, without loss of generality, that $g_1(x), \ldots, g_m(x)$ are independent. It follows that for any fixed $i \in [0, \ldots, m-1]$, there exist computable not-all-zero coefficients $a_1, \ldots, a_m \in \mathbb{Z}/p\mathbb{Z}$ such that

$$x^i = a_1 g_1(x) + \ldots + a_m g_m(x)$$

and so

$$\alpha^i P = a_1 C_{y_1} + \ldots + a_m C_{y_m}$$

In other words, the partial common reference string

$$\mathcal{CRS}_m \doteq \{P, \alpha P, \ldots, \alpha^{m-1}P\}$$

can be obtained from these witnesses and this will enable the attacker to compute membership and non-membership witnesses with respect to V for any accumulated and non-accumulated element, respectively.

We note that it is more convenient to execute the above attack with respect to the accumulator value V_0 and the polynomial $f_{V_0}(x)$: in fact, any non-membership witness for a never added element which is issued with respect to a later accumulator value than V_0, can be iteratively transformed back to a non-membership witness with respect to V_0 by just inverting the non-membership witness update formula outlined in Section 2. Once both $f_{V_0}(x)$ and $\mathcal{CRS}_{|\mathcal{Y}_{V_0}|}$ are computed, the attacker can issue witnesses with respect to V_0 for elements in and not in $\mathcal{Y}_{V_0}$ and update them with respect to the latest accumulator value as usual. Clearly, since it is possible to issue many different non-membership witnesses with respect to V_0, this implies that by updating them, these non-membership witnesses can be used to iteratively expand the previously computed partial common reference string $\mathcal{CRS}_{|\mathcal{Y}_{V_0}|}$.

Attack 1: The Witness Forgery Attack

Input : $n \geq |\mathcal{Y}_{V_0}|$ non-membership witnesses for never accumulated elements, the accumulator history (accumulator values and added/removed elements)

Output: a non-membership witness for a non-accumulated element or a membership witness for an accumulated one with respect to V

1. *Un-update* all non-membership witnesses with respect to V_0 inverting witness update formula and using accumulator history.
2. Interpolate the polynomial $f_{V_0}(x) = \prod_{y_i \in \mathcal{Y}_{V_0}} (y_i + x)$ from witnesses.
3. Use Euclidean Algorithm to find $g_i(x)$ and d_{y_i} such that $f_{V_0}(x) = g_i(x)(y_i + x) + d_{y_i}$ for every element y_i, $i = 1, \ldots, n$
4. Use linear algebra to write x^j as a linear combinations of $g_1(x), \ldots, g_n(x)$ for any $j = 0, \ldots, |\mathcal{Y}_{V_0}| - 1$
5. Obtain $\mathcal{CRS}_{|\mathcal{Y}_{V_0}|}$ from witnesses.
6. Use $\mathcal{CRS}_{|\mathcal{Y}_{V_0}|}$ and $f_{V_0}(x)$ to issue many different non-membership witnesses with respect to V_0.
7. Use the additional non-membership witnesses issued to expand the common reference string to $\mathcal{CRS}_{|\mathcal{Y}_V|}$.
8. Issue membership and non-membership witnesses with respect to the accumulator value V.

More precisely, given an accumulator value V we know that

$$V = \left(\prod_{y_i \in \mathcal{Y}_V \setminus \mathcal{Y}_{V_0}} (y + \alpha) \right) V_0 = f_V(\alpha) P$$

where $f_V(x)$ can be publicly computed from the published witness update information if the monic polynomial $f_{V_0}(x)$ is recovered by the attacker through interpolation, as outlined in Section 8.1.

Once the attacker successfully computes $\mathcal{CRS}_{|\mathcal{Y}_{V_0}|}$, they use it to issue (a multiple of) $|\mathcal{Y}_V| - |\mathcal{Y}_{V_0}|$ additional non-membership witnesses for random elements with respect to V_0, he updates them with respect to V and expands its starting set of elements and witnesses. Then, for each element y_i in this bigger set, he computes the corresponding polynomial $g_i(x)$ of degree $deg(f_V) - 1$ such that $f_V(x) = g_i(x)(y_i + x) + d_{y_i}$. At this point and similarly as before, the attacker can explicitly write a linear combinations of computable polynomials which equals x^i for any i such that $deg(f_{V_0}) - 1 < i \leq deg(f_V) - 1$, and thus can expand the previously computed $\mathcal{CRS}_{deg(f_{V_0})}$ to $\mathcal{CRS}_{deg(f_V)}$. In conclusion, an attacker would be able to forge witnesses with respect to the latest accumulator value by accessing only $|\mathcal{Y}_{V_0}|$ non-membership witnesses. The whole attack is summarized in Attack 1.

Similarly as discussed in Section 8.1, this attack can be prevented if the total number of issued non-membership witnesses is less than $|\mathcal{Y}_{V_0}|$.

9 Conclusions

In this paper, we cryptanalysed the Dynamic Universal Accumulator scheme proposed by Au *et al.* [1], investigating the security of the two constructions proposed, to which we refer as the α*-based* and the $\mathcal{CRS}$*-based* construction.

For the first construction we have shown several attacks which allow to recover the accumulator secret parameter α and thus break its collision resistance. More precisely, if p is the order of the underlying bilinear group, an attacker that has access to $O(\log p \log\log p)$ non-membership witnesses for random elements will be able to fully recover α, no matter how many elements are accumulated. If instead the elements can be chosen by the attacker, the number of required witnesses reduces down to just $O(\log p)$, thus making the attack linear in the size of the accumulator secret α. Furthermore, we showed how accumulator collision resistance can be broken in the α-based construction given *one* non-membership witness and we described also another minor design flaw.

For the second, i.e. the $\mathcal{CRS}$-based construction, we investigated resistance to witness forgeries under the hypothesis that the Accumulator Manager has the exclusive right to issue witnesses (as in authentication mechanisms) and thus keeps the $\mathcal{CRS}$ private. We have shown that an attacker that has access to multiple witnesses is able to reconstruct the Accumulator Manager $\mathcal{CRS}$, which would then enable him to compute witnesses for arbitrary elements. In particular, if the accumulator is initialized by accumulating m secret elements, m witnesses suffices to recover the secret $\mathcal{CRS}$.

Countermeasures We have shown that the α-based construction of Au *et al.* Dynamic Universal Accumulator is insecure, however one can still use the witness defining equations provided in the alternative $\mathcal{CRS}$-based construction, which is collision-resistant under the t-SDH assumption. There is one caveat: knowledge of $\mathcal{CRS}$ will enable an attacker to issue witnesses for arbitrary elements. If this needs to be avoided (ex. in authentication mechanisms), then $\mathcal{CRS}$ should be kept secret and the accumulator properly initialized. Namely, the accumulator manager needs to define an upper limit m to the total number of issuable non-membership witnesses and has to initialize the accumulator by adding $m + 1$ secret elements in order to prevent Attack 1.

Acknowledgements

We thank the anonymous reviewers for their helpful comments and suggestions. This work was supported by the Luxembourg National Research Fund (FNR) project FinCrypt (C17/IS/11684537).

References

1. Au, M. H., Tsang, P. P., Susilo, W. and Mu, Y. Dynamic Universal Accumulators for DDH Groups and Their Application to Attribute-Based Anonymous Credential Systems. *In CT-RSA*, Springer LNCS, **5473**, 295–308 (2009)

2. Baric, N. and Pfitzmann, B. Collision-free Accumulators and Fail-stop Signature Schemes Without Trees. *In EUROCRYPT*, 480–494, (1997)
3. Benaloh, J. and de Mare, M. One-way Accumulators: A Decentralized Alternative to Digital Signatures. *In EUROCRYPT*, 274–285 (1993)
4. Boneh, D., Boyen, X., Short signatures without random oracles. *In Advances in Cryptology – EUROCRYPT 2004*, Springer LNCS, 3027, 56–73 (2004)
5. Boneh, D. and Boyen, X. Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology*, **21(2)**, 149–177, (2008).
6. Boneh, D. and Corrigan-Gibbs, H. Bivariate polynomials modulo composites and their applications. *In International Conference on the Theory and Application of Cryptology and Information Security*, 42–62 (2014)
7. Boneh, D., Bünz, B. and Fisch, B. Batching techniques for accumulators with applications to IOPs and stateless blockchains. *Cryptology ePrint Archive*, Report 2018/1188, (2018)
8. Bosma, W., Cannon, J. J., Fieker, C., Steel, A. (eds.), *Handbook of Magma functions*, Edition 2.16 (2010)
9. Buldas, A., Laud, P. and Lipmaa, H. Accountable Certificate Management Using Undeniable Attestations. *In ACM CCS*, 9–17 (2000)
10. Buldas, A., Laud, P. and Lipmaa, H. Eliminating Counterevidence with Applications to Accountable Certificate Management. *Journal of Computer Security*, 10:2002, (2002)
11. Camacho, P., Hevia, A., Kiwi, M. A. and Opazo, R. Strong Accumulators from Collision-Resistant Hashing. *In ISC*, Springer LNCS, **4222**, 471–486 (2008)
12. Camenisch, J. and Lysyanskaya, A. Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials. *In CRYPTO*, 61–76 (2002)
13. Camenisch, J. and Soriente, C. An Accumulator Based on Bilinear Maps and Efficient Revocation for Anonymous Credentials. *In PKC 2009*, Springer LNCS, **5443**, 481–500 (2009)
14. Chang, Y.-C. and Lu, C.-J. Oblivious Polynomial Evaluation and Oblivious Neural Learning. *In Advances in Cryptology — ASIACRYPT 2001*, 369–384 (2001)
15. Damgård, I. and Triandopoulos, N. Supporting Non-membership Proofs with Bilinear-map Accumulators. *IACR Cryptology ePrint Archive*, 538 (2008)
16. Derler, D., Hanser, C. and Slamanig, D. Revisiting cryptographic accumulators, additional properties and relations to other primitives. *In CT-RSA 2015*, 127–144 (2015)
17. Li, J., Li, N. and Xue, R. Universal Accumulators with Efficient Nonmembership Proofs. *In ACNS*, Springer LNCS, **4521**, 253–269 (2007)
18. Lipmaa, H. Secure Accumulators from Euclidean Rings without Trusted Setup. *In ACNS*, Springer LNCS, **7341**, 224–240 (2012)
19. Merkle, R. C. A certified digital signature. *In Advances in Cryptology – CRYPTO 1989*, 218-–238 (1989)
20. Mitzenmacher M. and Upfal E. Probability and Computing: Randomized Algorithms and Probabilistic Analysis. Cambridge University Press (2005)
21. Nguyen, L. Accumulators from Bilinear Pairings and Applications. *In CT-RSA*, Springer LNCS, **3376**, 275–292 (2005)
22. Sander, T. Efficient Accumulators Without Trapdoor. *In ICICS*, 252–262 (1999)
23. The Sage Developers, *SageMath, the Sage Mathematics Software System (Version 9.0)*, `https://www.sagemath.org`, (2020)
24. Vitto G. and Biryukov A. Dynamic Universal Accumulator with Batch Update over Bilinear Groups. *IACR Cryptology ePrint Archive*, 777 (2020)