

Revisiting BBS Signatures

Stefano Tessaro  and Chenzhi Zhu 

Paul G. Allen School of Computer Science & Engineering
University of Washington, Seattle, US
{tessaro,zhucz20}@cs.washington.edu

Abstract. BBS signatures were implicitly proposed by Boneh, Boyen, and Shacham (CRYPTO '04) as part of their group signature scheme, and explicitly cast as stand-alone signatures by Camenisch and Lysyanskaya (CRYPTO '04). A provably secure version, called BBS+, was then devised by Au, Susilo, and Mu (SCN '06), and is currently the object of a standardization effort which has led to a recent RFC draft. BBS+ signatures are suitable for use within anonymous credential and DAA systems, as their algebraic structure enables efficient proofs of knowledge of message-signature pairs that support partial disclosure.

BBS+ signatures consist of one group element and two scalars. As our first contribution, we prove that a variant of BBS+ producing shorter signatures, consisting only of one group element and one scalar, is also secure. The resulting scheme is essentially the original BBS proposal, which was lacking a proof of security. Here we show it satisfies, under the q -SDH assumption, the same provable security guarantees as BBS+. We also provide a complementary tight analysis in the algebraic group model, which heuristically justifies instantiations with potentially shorter signatures.

Furthermore, we devise simplified and shorter zero-knowledge proofs of knowledge of a BBS message-signature pair that support partial disclosure of the message. Over the BLS12-381 curve, our proofs are 896 bits shorter than the prior proposal by Camenisch, Drijvers, and Lehmann (TRUST '16), which is also adopted by the RFC draft.

Finally, we show that BBS satisfies one-more unforgeability in the algebraic group model in a scenario, arising in the context of credentials, where the signer can be asked to sign arbitrary group elements, meant to be commitments, without seeing their openings.

1 Introduction

The seminal works of Camenisch and Lysyanskaya [16,17] highlighted how certain digital signature schemes with suitable algebraic structures are amenable to applications such as anonymous credentials, direct anonymous attestation (DAA), and group signatures. These schemes easily enable the signing of a *commitment*, typically by being algebraically compatible with a Pedersen commitment [31], and support very efficient zero-knowledge proofs of knowledge of a valid message-signature pair.

This paper revisits and improves BBS signatures [10,6,15], one of the most efficient pairing-based schemes with these properties, which has recently been in the midst of renewed interest in the context of decentralized identity. This has led to reference implementations [1,2], to a standardization effort by the W3C Verifiable Credentials Working group, and to an RFC draft [28]. BBS is also a building block for DAA [18,12,15], and is used by Intel SGX’s EPID protocol [13]. Furthermore, BBS signatures are theoretically interesting, due to their simplicity and efficiency. Most applications, and the RFC draft, consider the provably-secure version of BBS referred to as BBS+ [15,6], whose signatures consist of *one* group element in \mathbb{G}_1 and *two* scalars in \mathbb{Z}_p , where p is the group order.

OUR CONTRIBUTIONS, IN A NUTSHELL. As our first main contribution, we prove the strong unforgeability of a variant of BBS+ which produces shorter signatures only consisting of one group element and one scalar. The resulting scheme is in fact the original BBS signature scheme by Boneh, Boyen, and Shacham [10] as stated by Camenisch and Lysyanskaya [17], for which however no proof of security was known. Our new proof gives us a more efficient version of the scheme that can replace BBS+ in applications and standards with no loss, and re-affirms the security of prior works (e.g., [18,12]) which already used this optimized version but relied on an incorrect proof.

Furthermore, we provide a tighter security proof in the Algebraic Group Model [23], which also supports potentially shorter signatures. We also optimize the associated proofs of knowledge of BBS signatures, achieving substantial savings over the current state-of-the-art [28]. Finally, we study the security of BBS in contexts where group elements are signed, and show that the scheme satisfies, in the AGM, a security property which is a natural weakening of what is achieved by structure-preserving signatures [5].

BBS+. The BBS+ scheme was proposed by Au, Susilo, and Mu [6], based on ideas by [10,17], and proved secure under the q -SDH assumption. The proof was then adapted to type-3 pairings by Camenisch, Drijvers, and Lehmann [15]. It signs vectors $\mathbf{m} \in \mathbb{Z}_p^\ell$. To do so, the public parameters consist of $\ell + 2$ generators $g_1, h_0, h_1, \dots, h_\ell \in \mathbb{G}_1$, and a signature has the format (A, e, s) , where $s, e \in \mathbb{Z}_p$ are randomly chosen, and

$$A = \left(g_1 h_0^s \prod_{i=1}^{\ell} h_i^{\mathbf{m}[i]} \right)^{\frac{1}{x+e}}.$$

Here, $x \in \mathbb{Z}_p$ is the secret key, and given the public key $X_2 = g_2^x \in \mathbb{G}_2$, and a pairing, it is easy to verify a valid BBS+ signature.

SECURITY FOR BBS SIGNATURES. The only difference between BBS and BBS+ is the additional term h_0^s in the latter, which mandates the inclusion of s in the signature. A natural question is whether this term is necessary, or instead an artifact of the proofs [6,15]. Indeed, no attack seems to affect plain BBS, without the term h_0^s , but prior proof attempts (e.g., [12]) contained fundamental errors.

We prove that (plain) BBS signatures are indeed secure under the q -SDH assumption. The concrete security guarantees are essentially identical to those established for BBS+, and this suggests a more efficient drop-in replacement for BBS+ in existing applications. Our techniques close in particular gaps left by incorrect proofs, and can be used to prove exculpability of the original BBS group signature scheme [10].

TIGHT AGM BOUNDS. Our new proof, not unlike the prior proofs for BBS+, is not tight, i.e., it incurs a multiplicative loss equal to the number of signing queries q . As a strong hint that this loss may be artificial, we give a tight proof for BBS signatures in the Algebraic Group Model [23].

Our AGM analysis also addresses a different artificial aspect of the standard-model analysis, namely the random choice of e values from \mathbb{Z}_p . Instead, our AGM analysis merely asks that these values are unlikely to collide, and their collision probability becomes a term (meant to be negligible) in the security bound. This allows for more flexibility, in that the e values could be generated from a counter or (assuming random oracles) as a hash of the message. It also suggests a BBS variant, which we call *truncated BBS*, where e is chosen from $\mathbb{Z}_{2^{2\lambda}}$, where λ is the desired security level (typically, $\lambda = 128$). On BLS12-381, this does not have any benefit. However, as in all schemes based on q -SDH, one may want to assess the potential impact of attacks such as those by Brown and Gallant [14] and Cheon [20] and choose an even bigger curve—in that case, truncation of the scalar may become effective.

SIGNING COMMITMENTS. An important question is to which extent BBS can be thought as a signature scheme signing a user-supplied group element, i.e., an element $C \in \mathbb{G}_1$ is signed as $(C^{\frac{1}{x+e}}, e)$. Indeed, in the context of blind issuance of credentials, one can think of $C = g_1 h_1^{m_1} \cdots h_\ell^{m_\ell}$ as a commitment sent from the user to the signer, and the signer’s response $(C^{1/(x+e)}, e)$ is a valid BBS signature on $\mathbf{m} = (m_1, \dots, m_\ell)$. It is not hard to see that this form of BBS does not yield a secure signature scheme over group elements as e.g., a signature on C can easily be transformed, by squaring it, into a signature on C^2 . However, we show that, in the AGM, BBS satisfies a form of one-more unforgeability, where obtaining signatures of q group elements does not enable the attacker to produce valid BBS signatures (e.g., by “opening” these group elements as commitments) on more than q messages. This is sufficient in the context of blind issuance.

ZERO-KNOWLEDGE POKS. We also revisit the problem of proving knowledge of a BBS message-signature pair with new zero-knowledge proofs of knowledge which are shorter than state-of-the-art solutions adopted in the RFC draft [28] and initially proposed in [15]. To prove knowledge of a BBS message-signature pair (\mathbf{m}, σ) , without revealing k out of the ℓ components of \mathbf{m} , our proof (when compiled as a NIZK via the Fiat-Shamir transform) consists of 2 elements in \mathbb{G}_1 , as well as $k + 3$ scalars in \mathbb{Z}_p . The proof adopted in the RFC draft, in contrast, consists of 3 elements in \mathbb{G}_1 , and $k + 5$ scalars. While a reduction by one scalar is possible due to our removal of the random value s from a signature, the remaining optimizations are the result of a different approach.

RELATED SCHEMES. We note that when signing individual elements of \mathbb{Z}_p , the simpler Boneh-Boyen signatures [9] would typically outperform BBS. The closest scheme to BBS is the one by Pointcheval-Sanders (PS) [32]. PS signatures consist of two group elements, and are comparably efficient to the short version of BBS from this paper. However, both the public and the secret keys grow linearly with ℓ , the length of the message vector to be signed, whereas in BBS they consist of a single element. (The group generators in BBS can be generated as the output of a hash function, and they should not be part of the key materials.) PS signatures feature properties which BBS does not possess, including re-randomizability and aggregability. The latter property is essential for their use in the recent Coconut system [33], for which BBS does not appear suitable.

OUTLINE. Our new proof for BBS is given in Section 3, followed by our AGM analysis in Section 4. Our new zero-knowledge proofs are given in Section 5, and our analysis of BBS as a signature scheme on group elements is in Section 6. We give a technical overview next.

1.1 Technical Overview

NEW BBS PROOF. It is instructive to first review existing proofs [6,15] for BBS+. To this end, we consider the special case where we sign a *single* scalar $m \in \mathbb{Z}_p$, i.e., a signature under secret key $x \in \mathbb{Z}_p$ takes form, for random $s, e \in \mathbb{Z}_p$,

$$\sigma = (A, s, e), \text{ where } A = (g_1 h_0^s h_1^m)^{\frac{1}{x+e}}.$$

If an attacker obtains q adaptively chosen signatures (A_i, s_i, e_i) for message $m_i \in \mathbb{Z}_p$ and finally produces a valid forgery (A^*, e^*, s^*) for a message $m^* \in \mathbb{Z}_p$, we can identify three cases, which are to be addressed differently:

- (1) There exists an $i \in [q]$ such that $A_i = A^*$ and $e_i = e^*$
- (2) There exists an $i \in [q]$ such that $A_i \neq A^*$ and $e_i = e^*$
- (3) $e^* \notin \{e_1, \dots, e_q\}$.

The most challenging case is (2). Indeed, (1) implies that $g_1 h_0^{s_i} h_1^{m_i} = g_1 h_0^{s^*} h_1^{m^*}$, while $(s_i, m_i) \neq (s^*, m^*)$, which in turn implies a break of the discrete logarithm assumption in \mathbb{G}_1 . For (3), instead, one resorts to a by-now classical technique by Boneh and Boyen [9]. The key point here is that to break q -SDH, given $g_1, g_1^x, \dots, g_1^{x^q} \in \mathbb{G}_1$, along with g_2, g_2^x in \mathbb{G}_2 , it is enough to compute $g_1^{p(x)/(x+e)}$, for a polynomial $p(X)$ which is not divisible by $X + e$. This indeed allows us to recover $g_1^{1/(x+e)}$, which gives us a valid q -SDH solution.

To do so, the reduction picks e_1, \dots, e_q ahead of time. It uses g_2^x as the public key, but uses new generators $\bar{g}_1 = g_1^{\theta p(x)}$, $h_0 = g_1^{\alpha p(x)}$ and $h_1 = g_1^{\beta p(x)}$ for \mathbb{G}_1 , where $p(X) = \prod_{i=1}^q (X + e_i)$ and $\alpha, \beta, \theta \leftarrow \mathbb{Z}_p$. Note that \bar{g}_1, h_0, h_1 can easily be computed from $g_1^{x^i}$ for $i \in [q]$, since $p(X)$ has degree q , and that for any m, s , and $i \in [q]$, the reduction can always simulate a signature $(\bar{g}_1 h_0^s h_1^m)^{\frac{1}{x+e_i}}$, since $p(X)$ is divisible by $X + e_i$. Moreover, a forgery for $e^* \notin \{e_1, \dots, e_q\}$, allows us to compute $(e^*, g_1^{\frac{1}{x+e^*}})$, and break q -SDH.

HANDLING CASE (2). The value s was crucial in [6,15] to deal with (2). To see how it was used, let us assume that we can actually guess the index i for which (2) occurs. Then, with $p'(\mathbf{X}) = \prod_{j \neq i} (\mathbf{X} + e_j)$, the reduction can set

$$\bar{g}_1 = g_1^{\theta p'(x)}, \quad h_0 = \bar{g}_1^{\frac{(x+e_i)\delta-1}{\alpha}}, \quad h_1 = h_0^\beta.$$

Queries for $j \neq i$ can be answered as above for any (s, m) , since $p'(\mathbf{X})$ is divisible by $\mathbf{X} + e_i$. In contrast, for the i -th query, on message m_i , the reduction can only answer for the specific choice of $s_i = \alpha - \beta \cdot m_i$, since

$$A_i = (\bar{g}_1 h_0^{s_i} h_1^{m_i})^{\frac{1}{x+e_i}} = \left(\bar{g}_1^{(x+e_i)\delta} \right)^{\frac{1}{x+e_i}} = \bar{g}_1^\delta.$$

Despite the fact that s_i depends on m_i , one can show that its distribution is uniform. If the attacker now produces a forgery with $e^* = e_i$, it means that we have

$$A^* = \left(\bar{g}_1 h_0^{s^*} h_1^{m^*} \right)^{\frac{1}{x+e_i}} = \bar{g}_1^{\frac{1}{x+e_i} \left(1 + \frac{(x+e_i)\delta-1}{\alpha} (s^* + \beta m^*) \right)}$$

and the reduction can solve q -SDH because $1 + \frac{(x+e_i)\delta-1}{\alpha} (s^* + \beta m^*)$ is not divisible by $\mathbf{X} + e_i$.

OUR IMPROVEMENT. The reduction for BBS+ programs s in a message-dependent way to handle (2). Our main idea here is to let e play this role instead, thus dispensing with the use of s , and in fact obtaining a slightly simpler reduction. Concretely, for BBS, we drop h_0 , as it is not needed any more, and we set up

$$\bar{g}_1 = g_1^{\alpha p'(x)(x+\varepsilon_i)}, \quad h_1 = g_1^{\beta p'(x)}.$$

Here, $p'(\mathbf{X})$ is as above, and $\alpha, \beta, \varepsilon_i$ are random, and, most importantly, ε_i will not necessarily equal e_i . Now, every query $j \neq i$ can be answered as before. To answer the i -th query, however, we first observe that

$$C_i = \bar{g}_1 h_1^{m_i} = g_1^{\alpha p'(x)(x+\varepsilon_i)} \cdot g_1^{\beta p'(x)m_i} = g_1^{\alpha p'(x)(x+\varepsilon_i + \frac{\beta}{\alpha} m_i)}.$$

We are going to show that if we set $e_i = \varepsilon_i + \frac{\beta}{\alpha} m_i$, not only we can compute $A_i = C_i^{\frac{1}{x+e_i}}$, but *also*, this e_i has the right distribution. This last argument is somewhat involved. For example, it turns out that if message m_i is such that $x + \varepsilon_i + \frac{\beta}{\alpha} m_i = 0$, the distribution of e_i is not correct. Luckily, however, this is the only case, and moreover, if this indeed happened, this would mean $x = -\varepsilon_i - \frac{\beta}{\alpha} m_i$, and we could break q -SDH directly. If we then obtain a forgery (A^*, e^*) with $e^* = e_i$ for a message m^* , we note that the discrete logarithm of A^* is

$$\text{DL}_{g_1}(A^*) = \frac{\alpha p'(x)(x + \varepsilon_i + \frac{\beta}{\alpha} m^*)}{x + \varepsilon_i + \frac{\beta}{\alpha} m_i}.$$

However, one can show that $\mathbf{X} + \varepsilon_i + \frac{\beta}{\alpha} m_i$ does not divide $\alpha p'(\mathbf{X})(\mathbf{X} + \varepsilon_i + \frac{\beta}{\alpha} m^*)$ if $m_i \neq m^*$.

AGM SECURITY. In the AGM, we restrict our focus to algebraic adversaries, and in the case of BBS, this means that the adversary outputs a forgery A^* and its representation in terms of g_1, h_1, \dots, h_ℓ , as well as the \mathbb{G}_1 -part the prior signatures A_1, \dots, A_q . We note that the discrete logarithm of each A_i equals $\varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(x)$, where x is the secret key, \mathbf{y} is the vectors of discrete logarithms of \mathbf{h}_1 relative to g_1 , and

$$\varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(X) = \frac{1 + \langle \mathbf{y}, \mathbf{m}_i \rangle}{X + e_i}.$$

Analogously, the discrete logarithm of A^* for a forgery A^*, e^* for a message \mathbf{m}^* equals $\varphi_{\mathbf{m}^*, e^*}^{\mathbf{y}}(x)$. Further, the algebraic adversary gives us a representation of $\varphi_{\mathbf{m}^*, e^*}^{\mathbf{y}}(x)$ as an affine combination of the $\varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(x)$'s. Our key observation is that unless some very specific properties are satisfied by \mathbf{y} , the *function* $\varphi_{\mathbf{m}^*, e^*}^{\mathbf{y}}(X)$, as opposed to its evaluation at x , *cannot* be expressed as an affine combination of the functions $\varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(X)$. Therefore, x must be a zero to a (known) polynomial of degree at most q , and it can be recovered by factoring this polynomial. It also turns out that whenever the choice of \mathbf{y} does not allow this argument to go through, we are going to be able to recover a non-trivial discrete-logarithm relation, and break the discrete logarithm problem directly.

BBS SIGNATURES OF COMMITTED VALUES. Our AGM proof will enable us to also study the scenario where the adversary can query an oracle on an arbitrary $C \in \mathbb{G}_1$ and obtain $C^{\frac{1}{x+e}}$ for a random e . We show that in the AGM it is impossible, except with negligible probability, to come up with $q + 1$ valid BBS signatures upon querying this oracle q times. The main difficulty is that an AGM adversary here can query this oracle with group elements which are combinations of the outputs of previous queries. However, we are going to show how an algebraic adversary making q oracle queries can be simulated by one making queries to the actual BBS signing oracle. To do this, we rely on a property of our AGM proof above, namely that the statement holds even if the values e_1, e_2, \dots are known to the adversary beforehand, and we use this to give an inductive argument which shows how to build these signing queries in order to emulate the oracle signing a group element instead.

NEW PROOFS OF KNOWLEDGE. We give new Σ -protocols to prove knowledge of a message-signature pair for BBS, given, possibly, partial knowledge of the message. Our basic observation is that valid signature (A, e) for \mathbf{m} satisfies $e(A, X_2) = e(B, g_2)$, where $B = C(\mathbf{m})A^{-e}$, and $C(\mathbf{m}) = g_1 \prod_{i=1}^{\ell} h_i^{\mathbf{m}[i]}$. For the case where \mathbf{m} is fully known to the verifier, for example, our prover commits to a randomized version of A, B , namely $\bar{A} = A^r$ and $\bar{B} = B^r = C(\mathbf{m})^r \bar{A}^{-e}$. Then, the prover engages in a homomorphism proof [29] to show knowledge of a representation (r, e) of \bar{B} to the base $C(\mathbf{m})$ and \bar{A} .

2 Preliminaries

NOTATION. We will use the shorthand $[n] = \{1, \dots, n\}$. We will denote formal variables in polynomials with sans-serif letters X, Y, \dots , and for any modulus p ,

we let $\mathbb{Z}_p[X]$ be the ring of formal polynomials $a(X) = \sum_{i=0}^d a_i X^i$ with coefficients in \mathbb{Z}_p . As usual, d is the *degree* of $a(X)$.

Throughout the paper, we adopt as far as possible the concrete security and efficiency approach, and avoid qualitative statements. We refer to “efficient” informally to stress that an algorithm is meant to be as efficient as possible, but make theorems precise by giving explicit reductions in their proofs.

GROUPS AND PAIRINGS. We work with prime-order groups. For such a group \mathbb{G} , we denote by $1_{\mathbb{G}}$ the identity element, and let $\mathbb{G}^* = \mathbb{G} \setminus \{1_{\mathbb{G}}\}$ be the set of $p-1$ generators. We use multiplicative notation, and generally denote group elements with upper case letters, scalars with lower case ones, with the exception of generators. For a generator $g \in \mathbb{G}^*$ and a group element $X \in \mathbb{G}$, we also let $\text{DL}_g(X)$ be the discrete logarithm $x \in \mathbb{Z}_p$ of X to the base g , i.e., $g^x = X$.

For prime-order groups $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$, a *bilinear map* is an efficiently computable function $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ which satisfies both (1) *bi-linearity*, i.e., $e(A^x, B^y) = e(A, B)^{xy}$ for all $A \in \mathbb{G}_1$, $B \in \mathbb{G}_2$, and $x, y \in \mathbb{Z}_p$, and (2) *non-triviality*, i.e., $e(g_1, g_2) \in \mathbb{G}_T^*$ for all generators $g_1 \in \mathbb{G}_1^*$ and $g_2 \in \mathbb{G}_2^*$. We normally consider a *group parameters generation algorithm* GGen such that $\text{GGen}(1^\lambda)$ outputs a description $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$ such that $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are groups of order p , $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a bilinear map, and p is a λ -bit prime.

Our treatment is compatible with type-3 pairings (cf. e.g. [24]) like BLS curves [7] which allow for some of the most efficient implementations, e.g., using BLS12-381 [11]. The most relevant property is that $\mathbb{G}_1 \neq \mathbb{G}_2$ and no efficiently computable homomorphism $\mathbb{G}_2 \rightarrow \mathbb{G}_1$ exists.

SIGNATURE SCHEMES. A *signature scheme* SS consists of a setup algorithm SS.Setup , a *key generation algorithm* SS.KG , a (possibly randomized) *signing algorithm* SS.Sign , and a deterministic *verification algorithm* SS.Ver , satisfying the usual syntax and correctness requirement. In particular, SS.Setup outputs parameters par , upon which all algorithms depend. We also let the message space $\text{SS.M} = \text{SS.M}(\text{par})$ depend on par . (We implicitly assume that the signing algorithm returns an error symbol \perp , which is not a valid signature, if the message is not in the message space.) We target *strong unforgeability*, which is defined by Game $\text{SUF}_{\text{SS}}^A(\lambda)$ in Figure 1. The corresponding advantage metric is

$$\text{Adv}_{\text{SS}}^{\text{suf}}(\mathcal{A}, \lambda) = \Pr \left[\text{SUF}_{\text{SS}}^A(\lambda) \right] .$$

THE SECURITY ASSUMPTIONS. We will use the following variant of the *q-Strong Diffie-Hellman* (q -SDH) assumption, as defined by Boneh and Boyen [9] in a format meant to support type-3 pairings. It is formalized by Game $q\text{-SDH}_{\text{GGen}}^A(\lambda)$ on the left of Figure 2. We also consider the related *q-Discrete Logarithm* (q -DL) assumption, as formulated on the right of Figure 2 by Game $q\text{-DL}_{\text{GGen}}^A(\lambda)$, which only differs in the winning condition. We associate with these games the corresponding advantage metrics

$$\text{Adv}_{\text{GGen}}^{q\text{-sdh}}(\mathcal{A}, \lambda) = \Pr \left[q\text{-SDH}_{\text{GGen}}^A(\lambda) \right] , \quad \text{Adv}_{\text{GGen}}^{q\text{-dl}}(\mathcal{A}, \lambda) = \Pr \left[q\text{-DL}_{\text{GGen}}^A(\lambda) \right] . \quad (1)$$

Game $\text{SUF}_{\text{SS}}^{\mathcal{A}}(\lambda)$: $S \leftarrow \emptyset$, $\text{par} \leftarrow \text{SS.Setup}(1^\lambda)$ $(\text{vk}, \text{sk}) \leftarrow \text{SS.KG}(\text{par})$ $(M^*, \sigma^*) \leftarrow \mathcal{A}^{\text{SIGN}}(\text{par}, \text{vk})$ If $(M^*, \sigma^*) \notin S \wedge \text{SS.Ver}(\text{par}, \text{vk}, M^*, \sigma^*)$ then Return true Return false	Oracle $\text{SIGN}(M)$: $\sigma \leftarrow \text{SS.Sign}(\text{par}, \text{sk}, M)$ If $\sigma \neq \perp$ then $S \leftarrow S \cup \{(M, \sigma)\}$ Return σ
--	---

Fig. 1. Strong unforgeability.

Game $q\text{-SDH}_{\text{GGen}}^{\mathcal{A}}(\lambda)$: $\text{par} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow \text{GGen}(1^\lambda)$ $g_1 \leftarrow \mathbb{G}_1^*$, $g_2 \leftarrow \mathbb{G}_2^*$ $x \leftarrow \mathbb{Z}_p$ $(c, Z) \leftarrow \mathcal{A}(\text{par}, g_1, (g_1^{x^i})_{i \in [q]}, g_2, g_2^x)$ Return $Z = g_1^{1/(x+c)}$	Game $q\text{-DL}_{\text{GGen}}^{\mathcal{A}}(\lambda)$: $\text{par} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow \text{GGen}(1^\lambda)$ $g_1 \leftarrow \mathbb{G}_1^*$, $g_2 \leftarrow \mathbb{G}_2^*$ $x \leftarrow \mathbb{Z}_p$ $x' \leftarrow \mathcal{A}(\text{par}, g_1, (g_1^{x^i})_{i \in [q]}, g_2, g_2^x)$ Return $x' = x$
--	--

Fig. 2. Assumptions. The assumptions could also be defined with respect to fixed generators, but this would invalidate some of our security proofs.

We note that the q -SDH assumption implies the q -DL assumption for any q , as finding x implies finding $g_1^{1/(x+c)}$ for any c . The converse is not known to be true *in general*, but it is true for algebraic adversaries [8]. Notation-wise, we drop q whenever it equals one, and refer to the resulting assumption as the *Discrete Logarithm* (DL) assumption.

Remark 1. Our security proofs will repeatedly rely on the observation (due to Boneh and Boyen [9]) that, given $g_1, g_1^x, g_1^{x^2}, \dots, g_1^{x^q}$, computing $A = g_1^{p(x)/(x-e)}$, for any known non-zero polynomial $p(X) \in \mathbb{Z}_p[X]$ with degree at most q such that $p(e) \neq 0$, leads to a break q -SDH. This is because, by the polynomial remainder theorem, we can write $p(X) = d(X)(X - e) + r$, where the remainder $r = p(e) \in \mathbb{Z}_p$ is a non-zero integer mod p , whereas $d(X)$ has degree at most $q - 1$. Therefore, $A = g_1^{d(x)+r/(x-e)}$, and also,

$$(Ag_1^{-d(x)})^{1/r} = g_1^{\frac{1}{x-e}}$$

can be efficiently computed, and $(-e, g_1^{1/(x-e)})$ is a q -SDH solution.

3 New Proof for (Short) BBS Signatures

3.1 Description and implementation details

Figure 3 describes a version of BBS with shorter signatures than BBS+ [6]. We refer formally to this scheme as $\text{BBS} = \text{BBS}[\text{GGen}, \mathcal{D}_e, \ell]$, where GGen is a group parameter generator, \mathcal{D}_e is a distribution over \mathbb{Z}_p , and $\ell \geq 1$ a parameter. We

Algorithm BBS.Setup(1^λ) : $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow \text{GGen}(1^\lambda)$ $g_1 \leftarrow \$ \mathbb{G}_1^*, \mathbf{h}_1 \leftarrow \$ \mathbb{G}_1^\ell, g_2 \leftarrow \$ \mathbb{G}_2^*$ $\text{par} \leftarrow (p, g_1, \mathbf{h}_1, g_2, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$ Return par	Algorithm BBS.Sign($\text{sk} = x, \mathbf{m}$) : $C \leftarrow g_1 \prod_i \mathbf{h}_1[i]^{\mathbf{m}[i]}$ $e \leftarrow \$ \mathcal{D}_e$ $A \leftarrow C^{\frac{1}{x+e}}$ Return $\sigma = (A, e)$
Algorithm BBS.KG(par) : $(p, g_1, \mathbf{h}_1, g_2, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow \text{par}$ $x \leftarrow \$ \mathbb{Z}_p; X_2 \leftarrow g_2^x$ $\text{sk} \leftarrow x; \text{vk} \leftarrow X_2$ Return (sk, vk)	Algorithm BBS.Ver($\text{vk}, \mathbf{m}, \sigma = (A, e)$) : $C \leftarrow g_1 \prod_i \mathbf{h}_1[i]^{\mathbf{m}[i]}$ Return $\mathbf{e}(A, g_2^e \text{vk}) = \mathbf{e}(C, g_2)$

Fig. 3. BBS signature. The scheme is parameterized by GGen , \mathcal{D}_e , and the message length $\ell = \ell(\lambda) \geq 1$. Here, group operations are in the groups \mathbb{G}_1 and \mathbb{G}_2 determined by the parameters.

omit \mathcal{D}_e whenever it is understood to be the uniform distribution over \mathbb{Z}_p , and ℓ whenever it is clear from the context. Here, the message space is $\text{BBS.M} = \mathbb{Z}_p^\ell$, and it depends on the parameters in that the modulus p is determined by GGen via BBS.Setup . There is an unlikely event that the inversion to compute $1/(x+e)$ during signature issuance fails because $x+e=0$ —for ease of syntax, we use the convention that $1/0=0$. The BBS^+ scheme is a special case where each signed message \mathbf{m} is such that its first component, $\mathbf{m}[1]$, is randomly chosen. (And, therefore, needs to be made part of the signature.)

MODELING CHOICES. Our modeling is similar to that of [10,6,15], in that in particular we fix the message length via ℓ . One can easily accommodate unbounded-length messages, as in practice, the generators in \mathbf{h}_1 do not need to be fixed beforehand, and $\mathbf{h}_1[i]$ can be the output of a hash function (modeled as a random oracle) on some input that depends on i . This allows us to also sign messages in \mathbb{Z}_p^+ , given a suitable encoding. (The RFC draft [28] suggests hashing a length-dependent set of parameters into the first message block, although more efficient encodings certainly exist.)

We also model BBS as randomized, as this feature may be useful in some contexts, but we can de-randomize the scheme by applying a PRF to \mathbf{m} , or (more efficiently) to C , to derive e .

3.2 Security Analysis

We show security of BBS in the standard model, under the q -SDH assumption, for the setting where \mathcal{D}_e is the uniform distribution over \mathbb{Z}_p . Here, q is the number of signing queries issued by the signer. Note that this theorem *also* implies security of BBS^+ , as it corresponds to a special case of BBS where the first block of each signed message is randomly chosen, and included in the signature.

Theorem 1 (Security of BBS). *Let GGen be a group parameter generator, producing groups of order $p = p(\lambda)$. For every SUF adversary \mathcal{A} issuing at most*

$q = q(\lambda)$ signing queries, there exist adversaries \mathcal{B}_1 , \mathcal{B}_2 , and \mathcal{B}_3 such that

$$\begin{aligned} \text{Adv}_{\text{BBS}[\text{GGen}]}^{\text{sf}}(\mathcal{A}, \lambda) &\leq q \cdot \text{Adv}_{\text{GGen}}^{q\text{-sdh}}(\mathcal{B}_1, \lambda) + \text{Adv}_{\text{GGen}}^{\text{dl}}(\mathcal{B}_2, \lambda) \\ &\quad + \text{Adv}_{\text{GGen}}^{q\text{-sdh}}(\mathcal{B}_3, \lambda) + \frac{q^2}{2p} + \frac{q+2}{p}. \end{aligned}$$

The adversaries \mathcal{B}_1 , \mathcal{B}_2 and \mathcal{B}_3 are given explicitly in the proof, and run in time roughly comparable to that of \mathcal{A} .

The proof of the theorem is given in Section 3.3 below. The concrete bound is essentially the same as prior analyses of BBS+ [6,15], and we incur a factor q loss in the reduction. Below, in Theorem 2, we give a tight reduction to q -DL in the algebraic group model, which suggests this loss may be artificial.

DISCUSSION OF CONCRETE PARAMETERS. Even assuming the tight bound as the correct one, the reliance on q -SDH raises the question of the extent to which parameters should accommodate for Cheon’s attack [20] on q -SDH/ q -DL, which achieves complexity (roughly) $O(\sqrt{p/q})$ for certain choices of q . The RFC [28] suggests the use of BLS12-381 [11], which gives (roughly) a 256-bit group order. We could accommodate for roughly $q = 2^{36}$, for example, while still having 110-bit security. (This type of reasoning was for example adopted to justify BLS12-381 in zkSNARKs [4].) But even then, we observe that the only way we know to *actually* break BBS via Cheon’s attack is the reduction by Jao and Yoshida [27], which requires all signatures to be on the same message, with different e -values.¹ Not only this situation is unlikely to arise, but it *does not* occur if we de-randomize BBS, which is the choice the RFC also adopted for BBS+. It is an excellent question to see whether a similar attack exists even for de-randomized BBS.

3.3 Proof of Theorem 1

Let us consider an interaction of the adversary \mathcal{A} in the SUF game, where the adversary finally outputs a forgery (\mathbf{m}^*, σ^*) , where $\sigma^* = (A^*, e^*)$. We define three events, depending on the specific format of the forgery:

- **Forge₁**: This is the event where σ^* is a valid forgery, and a prior SIGN query has output a signature $\sigma_i = (A_i, e_i)$ for $A_i \neq A^*$, $e_i = e_i^*$, and some message $\mathbf{m}_i \neq \mathbf{m}^*$.
- **Forge₂**: This is the event where σ^* is a valid forgery, and a prior SIGN query output the *same* signature $\sigma_i = \sigma^*$ for a message $\mathbf{m}_i \neq \mathbf{m}^*$, or the forgery A^* equals $1_{\mathbb{G}_1}$.
- **Forge₃**: This is the event where σ^* is a valid forgery and completely fresh, i.e., neither of **Forge₁** or **Forge₂** occurs.

¹ Roughly, their attack considers the setting where $g^{\frac{1}{x+e_i}}$ is obtained for multiple e_i ’s.

As the union of these three events equal the event that (\mathbf{m}^*, σ^*) is a valid forgery, the union bound yields

$$\text{Adv}_{\text{BBS}}^{\text{uf}}(\mathcal{A}, \lambda) \leq \Pr[\text{Forge}_1] + \Pr[\text{Forge}_2] + \Pr[\text{Forge}_3] .$$

We will proceed in upper bounding these three probabilities via separate reductions. The hardest case is the analysis of Forge_1 , and this is where our proof differs from prior work. The analyses of Forge_2 and Forge_3 are essentially the same as in the original analysis of BBS+. The theorem statement then follows by combining Lemmas 1, 2, and 3, which we state next. The proof of Lemma 1 is given below in Section 3.4, whereas the proofs of Lemmas 2 and 3, which are more standard, are deferred to the full version.

Lemma 1 (Analysis of Forge_1). *There exists a q -sdh adversary \mathcal{B}_1 such that*

$$\Pr[\text{Forge}_1] \leq q \cdot \text{Adv}_{\text{GGen}}^{q\text{-sdh}}(\mathcal{B}_1, \lambda) + \frac{q^2}{2p} + \frac{1}{p} .$$

Lemma 2 (Analysis of Forge_2). *There exists a dl adversary \mathcal{B}_2 such that*

$$\Pr[\text{Forge}_2] \leq \text{Adv}_{\text{GGen}}^{\text{dl}}(\mathcal{B}_2, \lambda) + \frac{1}{p} .$$

Lemma 3 (Analysis of Forge_3). *There exists a q -sdh adversary \mathcal{B}_3 such that*

$$\Pr[\text{Forge}_3] \leq \text{Adv}_{\text{GGen}}^{q\text{-sdh}}(\mathcal{B}_3, \lambda) + \frac{q}{p} .$$

3.4 Proof of Lemma 1

We give an overview of the adversary \mathcal{B}_1 that underlies the reduction to q -SDH for Forge_1 . The formal pseudocode description is in Figure 4, although we omit there some lengthier and tedious descriptions of how to compute certain elements, and give them here in the text instead. Recall that q is a bound on the number of generated signatures, i.e., the number of queries to SIGN issued by the adversary \mathcal{A} . We assume here that exactly q queries are made, without loss of generality.

Given the q -SDH setup $g_1, X_{1,1} = g_1^x, \dots, X_{1,q} = g_1^{x^q}, g_2, X_2 = g_2^x$, the adversary \mathcal{B}_1 first generates a suitable setup to run \mathcal{A} . In particular, it picks random values $\varepsilon_1, \dots, \varepsilon_q \leftarrow_{\$} \mathbb{Z}_p$, as well as randomizers $\alpha \leftarrow_{\$} \mathbb{Z}_q^*$ and $\beta_1, \dots, \beta_q \leftarrow_{\$} \mathbb{Z}_q$. Then, the generators $\bar{g}_1 \in \mathbb{G}_1^*$ and $\mathbf{h}_1 \in \mathbb{G}_1^\ell$ are set to

$$\bar{g}_1 = g_1^{\alpha \cdot p(x) \cdot (x + \varepsilon_{i^*})} , \quad \mathbf{h}_1[i] \leftarrow g_1^{\beta_i \cdot p(x)} \text{ for all } i \in [\ell],$$

where $i^* \leftarrow_{\$} [q]$ and

$$p(X) = \prod_{i \in [q] \setminus \{i^*\}} (X + \varepsilon_i) .$$

<p>Adversary $\mathcal{B}_1(\text{par}, g_1, (X_{1,i})_{i \in [q]}, g_2, X_{2,1}) :$</p> <p>$(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathbf{e}) \leftarrow \text{par}$</p> <p>$i^* \leftarrow [q]; \text{cnt} \leftarrow 0; \text{Sigs} \leftarrow \emptyset; x^* \leftarrow \perp$</p> <p>$\varepsilon_1, \dots, \varepsilon_q \leftarrow \mathbb{Z}_p$</p> <p>$\alpha \leftarrow \mathbb{Z}_p^*; \beta_1, \beta_2, \dots, \beta_\ell \leftarrow \mathbb{Z}_p$</p> <p>$\bar{g}_1 \leftarrow g_1^{\alpha \cdot p(x) \cdot (x + \varepsilon_{i^*})}$</p> <p>For $i = 1$ to ℓ do $\mathbf{h}_1[i] \leftarrow g_1^{\beta_i \cdot p(x)}$</p> <p>$X_2 \leftarrow X_{2,1}; \bar{\text{par}} \leftarrow (p, \bar{g}_1, \mathbf{h}_1, g_2, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathbf{e})$</p> <p>$(\mathbf{m}^*, A^*, e^*) \leftarrow \mathcal{A}^{\text{SIGN}}(\bar{\text{par}}, X_2)$</p> <p>$C^* \leftarrow \bar{g}_1 \prod_{i=1}^q \mathbf{h}_1[i]^{\mathbf{m}^*[i]}$</p> <p>If $e_{i^*} = e^* \wedge \mathbf{e}(A^*, X_2 g_2^{e^*}) = \mathbf{e}(C^*, g_2) \wedge (A^*, e^*) \notin \text{Sigs}$ then</p> <p> If $x^* \neq \perp$ then</p> <p> Return $(0, g_1^{1/x^*})$ // direct break</p> <p> If $e_{i^*} \notin \{e_i\}_{i \in [q] \setminus \{i^*\}}$ then</p> <p> $\gamma \leftarrow \sum_{i=1}^\ell \beta_i (\mathbf{m}^*[i] - \mathbf{m}_{i^*}[i])$</p> <p> Return $(e_{i^*}, [A^* \cdot (A_{i^*}^{-1})]^\frac{1}{\gamma})$</p>	<p>Oracle $\text{SIGN}(\mathbf{m}) :$</p> <p>$\text{cnt} \leftarrow \text{cnt} + 1, \mathbf{m}_{\text{cnt}} \leftarrow \mathbf{m}$</p> <p>$C_{\text{cnt}} \leftarrow \bar{g}_1 \prod_{i=1}^q \mathbf{h}_1[i]^{\mathbf{m}[i]}$</p> <p>If $\text{cnt} \neq i^*$ then</p> <p> $e_{\text{cnt}} \leftarrow \varepsilon_{\text{cnt}}$</p> <p>Else</p> <p> $e_{\text{cnt}} \leftarrow \varepsilon_{i^*} + \sum_{i=1}^\ell \frac{\beta_i}{\alpha} \mathbf{m}[i]$</p> <p> If $C_{\text{cnt}} = 1_{\mathbb{G}_1}$ then</p> <p> $x^* \leftarrow \{x' \in \{-e_{\text{cnt}}\} \cup \{\varepsilon_i\}_{i \neq i^*} \mid g_1^{x'} = X_{1,1}\}$</p> <p> $e_{\text{cnt}} \leftarrow \frac{\varepsilon_{\text{cnt}}}{1 - e_{\text{cnt}}}$</p> <p>$A_{\text{cnt}} \leftarrow C_{\text{cnt}}^{\frac{1}{x + e_{\text{cnt}}}}$</p> <p>$\sigma_{\text{cnt}} \leftarrow (A_{\text{cnt}}, e_{\text{cnt}})$</p> <p>$\text{Sigs} \leftarrow \text{Sigs} \cup \{\sigma_{\text{cnt}}\}$</p> <p>Return σ_{cnt}</p>
---	---

Fig. 4. Adversary \mathcal{B}_1 in the proof of Lemma 1. Recall that once $\varepsilon_1, \dots, \varepsilon_q$ are fixed and understood from the context, we use the shorthand $p(X) = \prod_{i \in [q] \setminus \{i^*\}} (X + \varepsilon_i)$ for convenience. In the pseudo-code, we omit the explicit computations of \bar{g}_1, \mathbf{h}_1 , and A_{cnt} from C_{cnt} , which are detailed in the text.

It is not hard to see that \bar{g}_1 and \mathbf{h}_1 can be computed efficiently from part of the q -SDH setup $g_1, X_{1,1}, \dots, X_{1,q}$. Moreover, at least informally, it should be clear that as long as $x \notin \{-\varepsilon_1, \dots, -\varepsilon_q\}$, the distributions of \bar{g}_1 and \mathbf{h}_1 are correct, i.e., they are uniform in \mathbb{G}_1^* and \mathbb{G}_1^ℓ , respectively, since $g_1 \in \mathbb{G}_1^*$. (The formal argument about the correctness of distributions is given below, and this is only meant to serve as some intuition.) We stress that our simulation will not be correct if $x \in \{-\varepsilon_1, \dots, -\varepsilon_q\}$, so it is easiest to assume that this is not the case to understand the rest of the reduction.

HANDLING SIGNING QUERIES. The oracle SIGN then simulates the correct signing oracle, keeping a query counter cnt . Whenever $\text{cnt} \neq i^*$, it is not hard to see that SIGN can easily answer the query using $e_{\text{cnt}} = \varepsilon_{\text{cnt}}$. Indeed, if $x + \varepsilon_{\text{cnt}} \neq 0$, on input \mathbf{m} , the simulate SIGN can compute

$$A_{\text{cnt}} = \left(\bar{g}_1 \prod_{i=1}^\ell \mathbf{h}_1[i]^{\mathbf{m}[i]} \right)^{\frac{1}{x + \varepsilon_{\text{cnt}}}} = g_1^{p_i(x) [\alpha(x + \varepsilon_{i^*}) + \sum_{i=1}^\ell \beta_i \mathbf{m}[i]]},$$

where $p_i(x) = \prod_{i \notin \{\text{cnt}, i^*\}} (x + \varepsilon_i)$. It is easy to detect $x + \varepsilon_{\text{cnt}} = 0$, and in that case, $A_{\text{cnt}} = 1_{\mathbb{G}_1}$ by definition.

Crucially, when $\text{cnt} = i^*$ the adversary \mathcal{B}_1 answers the signing query differently. We observe first that, with $C_{\text{cnt}} \leftarrow \bar{g}_1 \prod_{i=1}^q \mathbf{h}_1[i]^{\mathbf{m}[i]}$,

$$\begin{aligned} \text{DL}_{g_1}(C_{\text{cnt}}) &= \alpha \cdot p(x) \cdot (x + \varepsilon_{i^*}) + \sum_{i=1}^{\ell} \beta_i \mathbf{m}[i] p(x) \\ &= \alpha \cdot p(x) \left(x + \varepsilon_{i^*} + \sum_{i=1}^{\ell} \frac{\beta_i}{\alpha} \mathbf{m}[i] \right). \end{aligned} \quad (2)$$

Here, there are two cases. Either $C_{\text{cnt}} = 1_{\mathbb{G}_1}$, and then we return $A_{\text{cnt}} = 1_{\mathbb{G}_1}$, along with a $e_{\text{cnt}} = \varepsilon_{\text{cnt}}$. Alternatively, and more interestingly, if $C_{\text{cnt}} \neq 1_{\mathbb{G}_1}$, we set $e_{i^*} = \varepsilon_{i^*} + \sum_{i=1}^{\ell} \frac{\beta_i}{\alpha} \mathbf{m}[i]$, and

$$A_{i^*} = C_{i^*}^{\frac{1}{x+e_{i^*}}} = g_1^{\alpha \cdot p(x)},$$

which can be efficiently computed. The bulk of our analysis below will show that if $C_{i^*} \neq 1_{\mathbb{G}_1}$, then we indeed generate a random e_{i^*} in this way.

Note that by equation (2) if $C_{i^*} = 1_{\mathbb{G}_1}$, then $x = -\varepsilon_{i^*} - \sum_{i=1}^{\ell} \frac{\beta_i}{\alpha} \mathbf{m}[i]$ or $x \in \{\varepsilon_i\}_{i \neq i^*}$, and hence we can directly break of q -SDH. (The variable x^* here stores the recovered discrete logarithm.) To simplify the analysis below, in this case, it is convenient for the reduction \mathcal{B}_1 to defer breaking q -SDH to end, and return the signature $(1_{\mathbb{G}_1}, \varepsilon_{i^*})$ instead.

EXTRACTING A SOLUTION. Assume now that \mathcal{A} outputs a valid forgery \mathbf{m}^*, σ^* , where $\sigma^* = (A^*, e^*)$, $e^* = e_{i^*}$, and $A^* \neq A_{i^*}$. Further, assume that $C_{i^*} \neq 1_{\mathbb{G}_1}$, which implies that $x + e_{i^*} \neq 0$ and $p(x) \neq 0$. (If this was not true, as highlighted above, we would have extracted x already.) Then,

$$\begin{aligned} \text{DL}_{g_1}(A^*) &= \alpha p(x) \frac{x + \varepsilon_{i^*} + \sum_{i=1}^{\ell} \frac{\beta_i}{\alpha} \mathbf{m}^*[i]}{x + \varepsilon_{i^*} + \sum_{i=1}^{\ell} \frac{\beta_i}{\alpha} \mathbf{m}[i]} \\ &= \alpha \cdot p(x) + p(x) \frac{\sum_{i=1}^{\ell} \beta_i (\mathbf{m}^*[i] - \mathbf{m}[i])}{x + e_{i^*}}. \end{aligned}$$

Further, because $A_{i^*} \neq A^*$ but $e^* = e_{i^*}$, we also have $\gamma = \sum_{i=1}^{\ell} \beta_i (\mathbf{m}^*[i] - \mathbf{m}[i]) \neq 0$, and then

$$[A^* \cdot (A_{i^*}^{-1})]^{\frac{1}{\gamma}} = g_1^{\frac{p(x)}{x+e_{i^*}}}.$$

If $e_{i^*} \notin \{e_i\}_{i \in [q] \setminus \{i^*\}}$, $X + e_{i^*}$ does not divide $p(X)$. We can then compute $g_1^{\frac{1}{x+e_{i^*}}}$ using Remark 1 and break q -SDH.

FORMAL ANALYSIS. We now proceed with a formal analysis to show that the probability guarantees for \mathcal{B}_1 as stated in the lemma indeed hold. To this end, we use $\Pr_0[\cdot]$ to denote probabilities in the experiment $\text{SUF}_{\text{BBS}}^{\mathcal{A}}(\lambda)$, where \mathcal{A} plays the SUF game against BBS. Similarly, we use $\Pr_1[\cdot]$ to denote probabilities in the simulated experiment where \mathcal{A} is run within \mathcal{B}_1 in Game $q\text{-SDH}_{\text{Gen}}^{\mathcal{B}_1}(\lambda)$.

In both experiments, we can define the event Forge_1 , as it only depends on the output of the adversary and the property of this output relative to its earlier signing query. Moreover, let $\text{Forge}_1^{(i)}$ for $i \in [q]$ be the event that Forge_1 happens and the i -th query is the first signing query that satisfies the condition for Forge_1 to happen. Let GoodE be the event that all e_i 's are distinct. Note that \mathcal{B}_1 is guaranteed to output a q -SDH solution if $\text{Forge}_1^{(i)}$ happens and $i = i^*$ and, moreover, GoodE also occurs. Also, note that the probability that GoodE and $\text{Forge}_1^{(i)}$ occurs is independent of whether $i = i^*$ or not, and therefore

$$\begin{aligned} \text{Adv}_{\text{GGen}}^{q\text{-sdh}}(\mathcal{B}_1, \lambda) &\geq \sum_{i=1}^q \Pr_1 [\text{GoodE} \wedge \text{Forge}_1^{(i)}] \cdot \Pr_1 [i^* = i] \\ &= \frac{1}{q} \cdot \sum_{i=1}^q \Pr_1 [\text{GoodE} \wedge \text{Forge}_1^{(i)}] = \frac{1}{q} \cdot \Pr_1 [\text{GoodE} \wedge \text{Forge}_1] . \end{aligned}$$

We rely on the following central lemma, which in particular shows that the simulation of \mathcal{A} 's execution within \mathcal{B} is nearly correct. While the intuition has been given above, the formal proof is rather subtle and we rely on the H-coefficient method [30,19] to prove the following.

Lemma 4. $\Pr_0 [\text{GoodE} \wedge \text{Forge}_1] - \Pr_1 [\text{GoodE} \wedge \text{Forge}_1] \leq \frac{q}{p}$.

Before turning to a proof of the lemma in Section 3.5 below, we observe that plugging the inequality of the lemma into the above yields

$$\begin{aligned} \text{Adv}_{\text{GGen}}^{q\text{-sdh}}(\mathcal{B}_1, \lambda) &\geq \frac{1}{q} \cdot \Pr_0 [\text{GoodE} \wedge \text{Forge}_1] - \frac{1}{p} \\ &\geq \frac{1}{q} (\Pr_0 [\text{Forge}_1] - \Pr_0 [\overline{\text{GoodE}}]) - \frac{1}{p} . \end{aligned}$$

On the other hand, $\Pr_0 [\overline{\text{GoodE}}] \leq \binom{q}{2} \frac{1}{p} \leq \frac{q^2}{2p}$, and thus we obtain the bound in Lemma 1 by re-arranging terms.

3.5 Proof of Lemma 4

We assume \mathcal{A} to be deterministic without loss of generality. We describe the transcripts of the interaction of \mathcal{A} in the SUF and within \mathcal{B}_1 as part of the q -SDH experiment, respectively, via the following two random variables

$$\begin{aligned} T_0 &= (g_1, g_2, \mathbf{h}_1, x, i^*, (\mathbf{m}_1, e_1), \dots, (\mathbf{m}_q, e_q)) , \\ T_1 &= (\bar{g}_1, g_2, \mathbf{h}_1, x, i^*, (\mathbf{m}_1, e_1), \dots, (\mathbf{m}_q, e_q)) \end{aligned}$$

where in T_0 , i^* is sampled uniformly from $[q]$, independently of everything else. We do not include X_2 , as $X_2 = g_2^x$ in both experiments. Moreover, in both experiments, the first component A_1, A_2, \dots of the the answer to each signature query is removed, as it is also a deterministic function of the rest of the transcript. Similarly, the final forgery (A^*, e^*) is also a function of T_0/T_1 . For this reason,

we note that the event $\text{GoodE} \wedge \text{Forge}_1$ is deterministically determined from T_0 and T_1 , in their respective experiments, i.e., there exists a Boolean function ϕ such that $\phi(T_b) = 1$ if and only if the event happens in the corresponding experiment. Therefore,

$$\Pr_0[\text{GoodE} \wedge \text{Forge}_1] - \Pr_1[\text{GoodE} \wedge \text{Forge}_1] \leq \text{SD}(T_0, T_1),$$

where $\text{SD}(T_0, T_1) = \frac{1}{2} \sum_{\tau} |\Pr[T_0 = \tau] - \Pr[T_1 = \tau]|$ is the total variation distance, which we upper bound by a special case of Patarin's H-coefficient method [30], which we introduce on the way. (We use the formalism from [26] here.)

INTERPOLATION PROBABILITIES. Concretely, for any potential value τ of T_b for $b \in \{0, 1\}$, which we denote as

$$\tau = (\underline{g}_1, \underline{g}_2, \underline{h}_1, \underline{x}, \underline{i}^*, (\underline{m}_1, \underline{e}_1), \dots, (\underline{m}_q, \underline{e}_q)),$$

we let $\mathbf{p}_0(\tau)$ and $\mathbf{p}_1(\tau)$ be its interpolation probabilities, i.e., the probabilities that we pick randomness in the respective experiment that would lead to transcript τ if queries $\underline{m}_1, \dots, \underline{m}_q$ are fixed ahead of time. (These probabilities are independent of \mathcal{A} , and only depend on τ and the randomness of the experiment.) We want to isolate the following type of good transcript.

Definition 1 (Good transcript). *We call τ good if $\underline{x} \notin \{-\underline{e}_1, \dots, -\underline{e}_q\}$. Otherwise, τ is bad.*

We are then going to prove that for all good transcripts τ , $\mathbf{p}_0(\tau) = \mathbf{p}_1(\tau)$. This is enough to conclude the proof, as it implies that

$$\text{SD}(T_0, T_1) \leq \Pr[T_0 \text{ is bad}] = \Pr_0[x \in \{-e_1, \dots, -e_q\}] \leq \frac{q}{p}.$$

To compute $\mathbf{p}_1(\tau)$ for a good transcript τ , we assume that the generator g_1 given to \mathcal{B}_1 is fixed. (Of course, it is actually sampled randomly from \mathbb{G}_1^* as part of the q -SDH instance, but the interpolation probability is the same conditioned on any particular choice, and thus we fix it.) The randomness then consists of $i^* \leftarrow [q]$, $x \leftarrow \mathbb{Z}_p$, $\alpha \leftarrow \mathbb{Z}_p^*$, $\varepsilon_1, \dots, \varepsilon_q \leftarrow \mathbb{Z}_p$, and $g_2 \leftarrow \mathbb{G}_2^*$. To generate the transcript τ , we need in particular

$$i^* = \underline{i}^*, g_2 = \underline{g}_2, x = \underline{x}, (\varepsilon_i)_{i \in [q] \setminus \{\underline{i}^*\}} = (\underline{e}_i)_{i \in [q] \setminus \{\underline{i}^*\}}$$

and as these values are chosen independently, this is true with probability

$$\frac{1}{q} \cdot \frac{1}{p-1} \cdot \frac{1}{p} \cdot \frac{1}{p^{q-1}} = \frac{1}{q(p-1)p^q}$$

over the choice of $i^*, g_2, x, \{\varepsilon_i\}_{i \in [q] \setminus \{\underline{i}^*\}}$. Let us assume this initial part of the transcript is consistent. We also need $\beta_1, \dots, \beta_\ell$ to ensure $\underline{h}_1 = \underline{h}_1$, which, conditioned on $x = \underline{x}$, is equivalent to the fact that

$$p(\underline{x}) \cdot \beta_i = \text{DL}_{g_1}(\underline{h}_1[i]) \quad \text{for all } i \in [\ell].$$

Because τ is good, $p(\underline{x}) \neq 0$, and therefore the ℓ equalities hold with probability $1/p^\ell$ over the choice of $\beta_1, \dots, \beta_\ell$.

Finally, conditioned on $i^*, g_2, x, \{\varepsilon_i\}_{i \in [g] \setminus \{i^*\}}, \{\beta\}_{i \in [\ell]}$ satisfying all above constraints, we need ε_{i^*} and α to ensure that

$$\bar{g}_1 = g_1^{\alpha p(\underline{x})(\underline{x} + \varepsilon_{i^*})} = \underline{g}_1, \quad e_{i^*} = \underline{e}_{i^*}.$$

There are two cases here, depending on $\underline{m} = \underline{m}_{i^*}$ and the associated value

$$\underline{C} = \underline{g}_1 \prod_{i=1}^{\ell} \underline{h}_1[i]^{\underline{m}[i]}.$$

Case 1: $\underline{C} = 1_{\mathbb{G}_1}$. Then, in this case, \mathcal{B}_1 ensures $e_{i^*} = \varepsilon_{i^*}$, and this value, which is uniform, equals \underline{e}_{i^*} with probability $1/p$. Conditioned on this,

$$p(\underline{x})(\underline{x} + \varepsilon_{i^*}) = p(\underline{x})(\underline{x} + \underline{e}_{i^*}) \neq 0$$

because τ is good. Thus $\underline{g}_1 = g_1^{\alpha p(\underline{x})(\underline{x} + \varepsilon_{i^*})}$ holds with probability $1/(p-1)$ over the choice of α from \mathbb{Z}_p^* .

Case 2: $\underline{C} \neq 1_{\mathbb{G}_1}$. For convenience, we write $\underline{e} = \underline{e}_{i^*}$, $\varepsilon = \varepsilon_{i^*}$, $\underline{a} = \text{DL}_{g_1}(\underline{g}_1)$, and $\underline{y} = \sum_i \beta_i \underline{m}[i]$. Here, $\alpha \leftarrow \mathbb{Z}_p^*$ and $\varepsilon \leftarrow \mathbb{Z}_p$ need to satisfy

$$\begin{aligned} \alpha p(\underline{x})(\underline{x} + \varepsilon) &= \underline{a} \\ \varepsilon + \frac{1}{\alpha} \underline{y} &= \underline{e} \end{aligned}$$

The second equation directly implies that

$$\varepsilon = \underline{e} - \frac{1}{\alpha} \cdot \underline{y}. \quad (3)$$

Substituting this into the first equation yields

$$\alpha \cdot p(\underline{x}) \left(\underline{x} + \underline{e} - \frac{1}{\alpha} \cdot \underline{y} \right) = p(\underline{x})(\alpha(\underline{x} + \underline{e}) - \underline{y}) = \underline{a}.$$

Re-arranging terms we get

$$\alpha = \frac{\underline{a}/p(\underline{x}) + \underline{y}}{\underline{x} + \underline{e}}. \quad (4)$$

This is indeed a value in \mathbb{Z}_p^* for two reasons. First off, $\underline{x} + \underline{e} \neq 0$ as τ is good. Second, $\underline{a}/p(\underline{x}) + \underline{y} \neq 0$. Indeed, if instead $\underline{a}/p(\underline{x}) + \underline{y} = 0$ were true, we would have

$$\underline{C} = \underline{g}_1 \prod_{i=1}^{\ell} \underline{h}_1[i]^{\underline{m}[i]} = \underline{g}_1^{\underline{a}} \prod_{i=1}^{\ell} g_1^{p(\underline{x}) \cdot \beta_i \cdot \underline{m}[i]} = g_1^{\underline{a} + p(\underline{x}) \underline{y}} = g_1^0 = 1_{\mathbb{G}_1},$$

a contradiction with the fact that we are in Case 2. Therefore, the probability over the choice of α, ε that (3) and (4) are both satisfied is $\frac{1}{p(p-1)}$.

Game $\text{SUF} + \text{GGen}, \text{eG}, \text{eS}(\lambda)$: $\text{Sigs} \leftarrow \emptyset; \text{cnt} \leftarrow 0$ $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \text{e}) \leftarrow \text{GGen}(1^\lambda)$ $g_1 \leftarrow \mathbb{G}_1^*, \mathbf{h}_1 \leftarrow \mathbb{G}_1, g_2 \leftarrow \mathbb{G}_2^*$ $\text{par} \leftarrow (p, g_1, \mathbf{h}_1, g_2, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \text{e})$ $\text{st}_e \leftarrow \text{eG}(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \text{e})$ $x \leftarrow \mathbb{Z}_p; X_2 \leftarrow g_2^x; \text{sk} \leftarrow x; \text{vk} \leftarrow X_2$ $(\mathbf{m}^*, (A^*, e^*)) \leftarrow \mathcal{A}^{\text{SIGN}}(\text{par}, \text{st}_e, \text{vk})$ If $(\mathbf{m}^*, (A^*, e^*)) \notin \text{Sigs}$ then $C^* \leftarrow g_1 \prod_{i=1}^{\ell} \mathbf{h}_1[i]^{m^*[i]}$ If $\text{e}(A^*, X_2 g_2^{e^*}) = \text{e}(C^*, g_2)$ then return true Return false	Oracle $\text{SIGN}(\mathbf{m})$: $\text{cnt} \leftarrow \text{cnt} + 1$ $e_{\text{cnt}} \leftarrow \text{eS}(\text{st}_e, \text{cnt})$ $C_{\text{cnt}} \leftarrow g_1 \prod_{i=1}^{\ell} \mathbf{h}_1[i]^{m[i]}$ $A_{\text{cnt}} \leftarrow C_{\text{cnt}}^{1/(x+e_{\text{cnt}})}$ $\text{Sigs} \leftarrow \text{Sigs} \cup \{(\mathbf{m}, (A_{\text{cnt}}, e_{\text{cnt}}))\}$ Return A_{cnt}
--	---

Fig. 5. Stronger security for BBS. Stronger ad-hoc unforgeability achieved by BBS in the AGM, where the e_i 's are sampled deterministically from an algorithm that uses an initially generated state st_e , known to the adversary.

Therefore, in summary, we have

$$\mathbf{p}_1(\tau) = \frac{1}{q(p-1)p^q} \cdot \frac{1}{p^\ell} \cdot \frac{1}{p(p-1)} = \frac{1}{q(p-1)^2 p^{q+\ell+1}}.$$

It is not hard to observe that we also have

$$\mathbf{p}_0(\tau) = \frac{1}{p^{q+\ell+1}(p-1)^2 q},$$

because g_1, g_2 are uniform over \mathbb{G}_1^* , \bar{h}_1 is uniform over \mathbb{G}_1^ℓ , and x, e_1, \dots, e_q are uniform in \mathbb{Z}_p , and i^* is uniform in $[q]$.

4 Tighter Proofs for BBS in the AGM

This section complements the above standard-model analysis with a tight analysis of BBS in the *algebraic group model* (AGM) [23]. In addition, we prove here that security holds even if the attacker is given the values e_1, e_2, \dots *ahead of time*, and we allow these values to be sampled from a more general distribution. The former fact will be helpful later in Section 6. The latter fact will allow for instantiations of BBS with shorter signatures in some contexts, as we explain further below.

STRONGER SECURITY. We formalize our security goal in terms of Game $\text{SUF} +$ in Figure 5. This is *not* a generic security game, as it is specific to BBS, but clearly, it does imply its strong unforgeability in a number of settings when the scheme instantiation corresponds to a particular pick to eG and eS . The ad-hoc feature is that we allow part of the signature (namely, the e value in a pair (A, e)) to be generated initially. To model this, in addition to the group parameter

generator GGen , the game is parameterized by a pair of algorithms, eG and eS , where eG , on input the group parameters, outputs a state st_e , and then $\text{eS}(\text{st}_e, i)$ (deterministically) outputs the value e_i used for the i -th signature. The initial state st_e is given to the adversary, who can run eS to pre-compute the e_i 's. It will be convenient to define the collision probability

$$\delta_{\text{eG}, \text{eS}}(q, \lambda) = \Pr \left[\begin{array}{l} \text{par} \leftarrow_s \text{GGen}(1^\lambda) \\ \text{st}_e \leftarrow_s \text{eG}(\text{par}) \end{array} : \exists 1 \leq i < j \leq q : \text{eS}(\text{st}_e, i) = \text{eS}(\text{st}_e, j) \right] .$$

We also define the advantage metric

$$\text{Adv}_{\text{GGen}, \text{eG}, \text{eS}}^{\text{suf}+}(\mathcal{A}, \lambda) = \Pr [\text{SUF} +_{\text{GGen}, \text{eG}, \text{eS}}^{\mathcal{A}}(\lambda)] .$$

ALGEBRAIC SECURITY. We are now ready to state our main theorem, which is proved below in Section 4.1. We dispense with a full formalization of the AGM [23], as its use is relatively straightforward here. Namely, we consider *algebraic adversaries* that provide an explanation of the element $A^* \in \mathbb{G}_1$ contained in the forgery in terms of all previously seen group elements in \mathbb{G}_1 , which include the generators g_1, \mathbf{h}_1 , as well as the issued signatures. (Because we consider type-3 pairings, we do not include \mathbb{G}_2 elements in these representations.) We also give our reduction here to q -DL, as opposed to q -SDH as in the case of Theorem 1, but note that the assumptions are equivalent in the AGM.

Theorem 2 (AGM Security of BBS). *Let GGen be a group parameter generator, producing groups of order $p(\lambda)$, and let eG , eS as above. For every algebraic $\text{SUF}+$ adversary \mathcal{A} issuing at most q signing queries, there exist adversaries \mathcal{B}_1 and \mathcal{B}_2 such that*

$$\text{Adv}_{\text{GGen}, \text{eG}, \text{eS}}^{\text{suf}+}(\mathcal{A}, \lambda) \leq \text{Adv}_{\text{GGen}}^{q\text{-dl}}(\mathcal{B}_1, \lambda) + \text{Adv}_{\text{GGen}}^{\text{dl}}(\mathcal{B}_2, \lambda) + \delta_{\text{eG}, \text{eS}}(q, \lambda) + \frac{1}{p(\lambda)} .$$

The adversaries \mathcal{B}_1 and \mathcal{B}_2 are given explicitly in the proof, and have running times comparable to that of \mathcal{A} . The adversary \mathcal{B}_1 need to additionally factor a polynomial of degree (at most) q .

The only property required from eG and eS is that $\delta_{\text{eG}, \text{eS}}(q, \lambda)$ is small. We note that the lack of collisions is a necessary condition. Indeed, if we generate two signatures $(A, e), (A', e)$ for messages \mathbf{m} and \mathbf{m}' , respectively, it is easy to verify that $((A \cdot A')^{\frac{1}{2}}, e)$ is a signature for $\frac{1}{2}(\mathbf{m} + \mathbf{m}')$, where $\frac{1}{2}$ is the inverse of 2 mod p .

The above theorem supports the security (in the AGM) of some interesting and natural instantiations of BBS with shorter signatures, which we discuss next. **COUNTER BBS.** One natural instantiation, which we refer to as *Counter BBS*, generates the e_i 's from a counter, i.e., $\text{eS}(\text{st}_e, i) = i$. This can be advantageous if the signer can reliably maintain such a counter. Signatures then would consist of a group element in \mathbb{G}_1 and then $\log q$ additional bits, where q is an upper bound on the number of issued signatures. In particular, one could safely set $q = 2^{50}$ in many applications, leading to very short signatures.

TRUNCATED BBS. A different application scenario considers a conservative instantiation that uses a 384-bit group \mathbb{G}_1 to prevent Cheon's attack [20]. Then, the above bound allows us to choose the e_i 's from $\mathbb{Z}_{2^{256}}$, as opposed to \mathbb{Z}_p for a 384-bit prime p , hence saving 128-bit of signature length. We refer to the resulting scheme as *Truncated BBS*. While we need to rely on the AGM to trust this optimization, we do note that the uniformity of the e_i 's needed by Theorem 1 appears to be an artifact of the proof, and does not appear to prevent actual attacks.

4.1 Proof of Theorem 2

Before we turn to the construction of the adversaries \mathcal{B}_1 and \mathcal{B}_2 , and their formal analysis, we introduce the algebraic framework that will guide their construction.

ALGEBRAIC FRAMEWORK. To start with, in an execution of $\text{SUF} + \mathcal{A}_{\text{GGen}, \text{eG}, \text{eS}}^A(\lambda)$, it is convenient to associate the discrete logarithms of group elements in \mathbb{G}_1 with formal rational functions (which are then evaluated in the actual execution to obtain the discrete logarithm). In particular, let us denote the discrete logarithms of \mathbf{h}_1 to the base g_1 by the vector $\mathbf{y} \in \mathbb{Z}_p^\ell$. Then, the i -th SIGN query for $\mathbf{m} \in \mathbb{Z}_p^\ell$, where $e_i = e$, returns a value with discrete logarithm $\varphi_{\mathbf{m}, e}^{\mathbf{y}}(x)$, where

$$\varphi_{\mathbf{m}, e}^{\mathbf{y}}(X) = \frac{1 + \langle \mathbf{y}, \mathbf{m} \rangle}{X + e_i}.$$

Here, $\langle \mathbf{x}, \mathbf{y} \rangle$ denotes inner product in \mathbb{Z}_p . It turns out that these functions are essentially linearly independent, except for some unfortunate configurations for \mathbf{y} . This is captured by the following central lemma.

Lemma 5. *Let $e_1, \dots, e_q \in \mathbb{Z}_p$ be distinct, let $\mathbf{y} \in \mathbb{Z}_p^\ell$, and let $\mathbf{m}_1, \dots, \mathbf{m}_q \in \mathbb{Z}_p^\ell$. Further, let $(\mathbf{m}^*, e^*) \notin \{(\mathbf{m}_i, e_i)\}_{i \in [q]}$. Then, assume that there exist $\lambda_1, \dots, \lambda_q, \gamma \in \mathbb{Z}_p$ such that*

$$\varphi_{\mathbf{m}^*, e^*}^{\mathbf{y}}(X) = \sum_{i=1}^q \lambda_i \cdot \varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(X) + \gamma. \quad (5)$$

Then, one of the following two conditions must be true:

- (i) *There exists $i \in [q]$ such that $e^* = e_i$ and $1 - \lambda_i + \langle \mathbf{y}, \mathbf{m}^* - \lambda_i \cdot \mathbf{m}_i \rangle = 0$.*
- (ii) *We have $e^* \notin \{e_1, \dots, e_q\}$, but $1 + \langle \mathbf{y}, \mathbf{m}^* \rangle = 0$.*

Proof. To verify (i), assume indeed that $e^* \in \{e_1, \dots, e_q\}$, and wlog, let $e^* = e_1$. We multiply both sides of (5) by $p(X) = \prod_{i=1}^\ell (X + e_i)$, and after re-arranging terms, we get

$$(1 - \lambda_1 + \langle \mathbf{y}, \mathbf{m}^* - \lambda_1 \cdot \mathbf{m}_1 \rangle) \cdot p_1(X) = \gamma \cdot p(X) + \sum_{i=2}^q \lambda_i \cdot \varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(X) \cdot p_i(X), \quad (6)$$

where we have used the shorthand $p_i(X) = \prod_{j \in [q] \setminus \{i\}} (X + e_j) = p(X)/(X + e_i)$. We claim that the LHS and RHS of (6) cannot be identical functions, unless

$1 - \lambda_1 + \langle \mathbf{y}, \mathbf{m}^* - \lambda_1 \cdot \mathbf{m}_1 \rangle = 0$. Indeed, the RHS is always divisible by $\mathbf{X} + e_1$, because either $\lambda_2, \dots, \lambda_q$ are all 0, in which case this is vacuously true, or $p(\mathbf{X})$ and $p_i(\mathbf{X})$ for $i \geq 2$ are divisible by $(\mathbf{X} + e_1)$. In contrast, if $1 - \lambda_1 + \langle \mathbf{y}, \mathbf{m}^* - \lambda_1 \cdot \mathbf{m}_1 \rangle \neq 0$, then the RHS is not divisible by $\mathbf{X} + e_1$ because $p_1(e_1) \neq 0$.

Let us consider instead the case $e^* \notin \{e_1, \dots, e_q\}$. For notational convenience, we let $e_{q+1} = e^*$, $p'(\mathbf{X}) = \prod_{i \in [q+1]} (\mathbf{X} + e_i)$, and $p'_k(\mathbf{X}) = \prod_{i \in [q+1] \setminus \{k\}} (\mathbf{X} + e_i)$. Then, multiplying both sides of (5) by $p'(\mathbf{X})$ yields

$$(1 + \langle \mathbf{y}, \mathbf{m}^* \rangle) \cdot p'_{q+1}(\mathbf{X}) = \gamma \cdot p'(\mathbf{X}) + \sum_{i=2}^q \lambda_i \cdot \varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(\mathbf{X}) \cdot p'_i(\mathbf{X}).$$

We notice that if $1 + \langle \mathbf{y}, \mathbf{m}^* \rangle \neq 0$ the LHS is non-zero, and not divisible by $\mathbf{X} + e_{q+1}$, as $p'_{q+1}(e_{q+1}) \neq 0$. In contrast, the RHS is always divisible by $\mathbf{X} + e_{q+1}$. A contradiction. \square

OVERVIEW OF THE REDUCTION. Let \mathcal{A} be an algebraic adversary in Game $\text{SUF} + \mathcal{A}_{\text{Gen}, \text{eG}, \text{eS}}(\lambda)$. It initially receives group elements $g_1 \in \mathbb{G}_1$, $\mathbf{h}_1 \in \mathbb{G}_1^\ell$, along with \mathbb{G}_2 elements $g_2, X_2 = g_2^x$. For each signing query, she also gets $A_i \in \mathbb{G}_1$. Finally, when producing a forgery $(\mathbf{m}^*, (A^*, e^*))$, by virtue of being algebraic, the adversary \mathcal{A} also provides a representation $(\gamma_0, \gamma_1, \dots, \gamma_\ell, \lambda_1, \dots, \lambda_q) \in \mathbb{Z}_p^{q+\ell+1}$ of A^* such that

$$A^* = g_1^{\gamma_0} \prod_{i=1}^{\ell} \mathbf{h}_1[i]^{\gamma_i} \prod_{i=1}^q A_i^{\lambda_i} = (C^*)^{\frac{1}{x+e^*}} = g_1^{\varphi_{\mathbf{m}^*, e^*}^{\mathbf{y}}(x)},$$

where $\mathbf{y}[i] = \text{DL}_{g_1}(\mathbf{h}_1[i])$ for all $i \in [\ell]$. Further, we have $A_i = g_1^{\varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(x)}$. Therefore, setting $\gamma = \gamma_0 + \sum_{i \in [\ell]} \gamma_i \cdot \mathbf{y}[i]$, this implies in particular that

$$\gamma + \sum_{i=1}^q \lambda_i \cdot \varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(x) - \varphi_{\mathbf{m}^*, e^*}^{\mathbf{y}}(x) = 0.$$

Let us now assume that the two conditions in Lemma 5 do not hold, and that e_1, \dots, e_q are distinct. Then, Lemma 5 implies that

$$\rho(\mathbf{X}) = \gamma + \sum_{i=1}^q \lambda_i \cdot \varphi_{\mathbf{m}_i, e_i}^{\mathbf{y}}(\mathbf{X}) - \varphi_{\mathbf{m}^*, e^*}^{\mathbf{y}}(\mathbf{X}) \neq 0,$$

and therefore x is one of its zeros. Assuming that $x \notin \{-e_1, \dots, -e_q, -e^*\}$, such zeros can be obtained by factoring the non-zero polynomial

$$q(\mathbf{X}) = \rho(\mathbf{X}) \cdot \prod_{e \in \{e_1, \dots, e_q, e^*\}} (\mathbf{X} + e),$$

which has degree at most $q + 1$. One of the zeros has to equal x .

We still need to handle the case where either $1 + \langle \mathbf{y}, \mathbf{m}^* \rangle = 0$ or $1 - \lambda_i + \langle \mathbf{y}, \mathbf{m}^* - \lambda_i \cdot \mathbf{m}_i \rangle = 0$. It is however not hard to see that this gives us non-trivial discrete logarithm relation, and we can use this to compute the discrete logarithm directly.

FORMAL REDUCTION. To formalize the above analysis, we consider three events during the execution of Game $\text{SUF} + \mathcal{A}_{\text{Gen}, \text{eG}, \text{eS}}(\lambda)$:

- **Forge**: This is the event that \mathcal{A} outputs a successful forgery and wins the game.
- **Rel**: This is the event that the forgery is for a message \mathbf{m}^* such that either Condition (i) or (ii) of Lemma 5 holds.
- **Col**: Is the event that there exist distinct $i, j \in [q]$ with $e_i = e_j$.

Then, by the law of total probability,

$$\begin{aligned} \text{Adv}_{\text{Gen}, \text{eG}, \text{eS}}^{\text{suf}+}(\mathcal{A}, \lambda) &= \Pr[\text{Forge} \wedge \overline{\text{Rel}}] + \Pr[\text{Forge} \wedge \text{Rel}] \\ &\leq \Pr[\text{Forge} \wedge \overline{\text{Rel}} \wedge \text{Col}] + \Pr[\text{Forge} \wedge \overline{\text{Rel}} \wedge \overline{\text{Col}}] + \Pr[\text{Rel}] \\ &\leq \Pr[\text{Col}] + \Pr[\text{Forge} \wedge \overline{\text{Rel}} \wedge \overline{\text{Col}}] + \Pr[\text{Rel}] . \end{aligned}$$

By definition, we know that $\Pr[\text{Col}] = \delta_{\text{eG}, \text{eS}}(q, \lambda)$. We now give adversaries \mathcal{B}_1 and \mathcal{B}_2 such that

$$\Pr[\text{Forge} \wedge \overline{\text{Rel}} \wedge \overline{\text{Col}}] \leq \text{Adv}_{\text{Gen}}^{q\text{-dl}}(\mathcal{B}_1) , \quad \Pr[\text{Rel}] \leq \text{Adv}_{\text{Gen}}^{\text{dl}}(\mathcal{B}_2) + \frac{1}{p} .$$

THE ADVERSARY \mathcal{B}_1 . If $x \notin \{-e_1, \dots, -e_q\}$ (which can be checked right away, and gives x), the reduction simulates the original generator g_1 as $\bar{g}_1 = g_1^{\alpha p(x)}$, where $p(x) = \prod_{i \in [q]} (x + e_i)$. (This can be computed given the inputs $X_{1,i} = g^{x_i}$ for $i \in [q]$.) Since $\alpha \in \mathbb{Z}_p^*$ and $p(x) \neq 0$, the simulation is perfect. Also, we can easily compute the answers to **SIGN** queries due to our choice of \bar{g}_1 . The adversary then checks that $q(X)$ is non-zero (which is implied by $\overline{\text{Rel}}$), and if so, proceeds to compute its zeros. It is easy to verify that the adversary succeeds with probability at least $\Pr[\text{Forge} \wedge \overline{\text{Rel}} \wedge \overline{\text{Col}}]$.

THE ADVERSARY \mathcal{B}_2 . The construction is somewhat standard. Let us assume one of the two conditions leading to **Rel** occurs in Game $\text{SUF} + \mathcal{A}_{\text{Gen}, \text{eG}, \text{eS}}(\lambda)$. First, if (i) occurs for some $i \in [q]$, then

$$g_1^{1-\lambda_i} \prod_{j=1}^{\ell} h_1^{\mathbf{m}^*[j] - \lambda_i \mathbf{m}_i[j]} = 1_{\mathbb{G}_1} . \quad (7)$$

In contrast, if (ii) occurs, then

$$g_1 \prod_{j=1}^{\ell} h_1^{\mathbf{m}_i[j]} = 1_{\mathbb{G}_1} . \quad (8)$$

Therefore, in both cases, we obtain a non-zero vector $(a, \mathbf{b}) \in \mathbb{Z}_p^{\ell+1}$ such that $g_1^a \prod_{i \in [\ell]} h_1[i]^{b[i]} = 1_{\mathbb{G}_1}$. Given the DL instance $X_{1,1} = g_1^x \in \mathbb{G}_1$, the adversary \mathcal{B}_2 simulates the generator by picking $\alpha_i, \beta_i \xleftarrow{\$} \mathbb{Z}_p$ for $i \in [\ell]$, and lets

$$\bar{g}_1 = X_{1,1} = g_1^x , \quad h_1[i] = X_{1,1}^{\alpha_i} g_1^{-\beta_i} = g_1^{\alpha_i x - \beta_i} \quad \text{for } i \in [\ell] .$$

This simulates the right distribution if $X_{1,1} \neq 1_{\mathbb{G}_1}$, which is ensured beforehand. Then, as outlined above, the adversary simulates a correct execution with \mathcal{A} , and checks if we obtain a non-trivial relation as above. If so, we are given a non-zero $(a, \mathbf{b}) \in \mathbb{Z}_p^{\ell+1}$ such that

$$ax + \sum_{i=1}^{\ell} \mathbf{b}[i] \cdot (\alpha_i x - \beta_i) = 0 ,$$

and, in turn, we get

$$x = \frac{\sum_{i=1}^{\ell} \beta_i \mathbf{b}[i]}{a + \sum_{i=1}^{\ell} \alpha_i \mathbf{b}[i]} .$$

Note that this is well defined, unless $a + \sum_{i=1}^{\ell} \alpha_i \mathbf{b}[i] = 0$. However, because $(a, \mathbf{b}) \neq \mathbf{0}$, and the fact that the α_i 's are uniform and independent given the adversary's view, this happens with probability at most $1/p$. This concludes the proof. \square

5 Efficient Proofs of Knowledge for BBS Signatures

We discuss zero-knowledge proofs of knowledge (zkPoK) of a BBS message-signature pair (\mathbf{m}, σ) that are shorter than those from [15] adopted by the RFC draft [28]. If we do not want to reveal k of the components of \mathbf{m} , our proof (when compiled as a NIZK via the Fiat-Shamir transform) consists of 2 elements in \mathbb{G}_1 , as well as $k+3$ scalars in \mathbb{Z}_p . The prior proof, in contrast, consists of 3 elements in \mathbb{G}_1 , and $k+5$ scalars. The benefit of our new proofs is also computational: We save 4 group exponentiations in \mathbb{G}_1 and 3 scalar multiplications for the prover, and save 3 group exponentiations in \mathbb{G}_1 for the verifier. We note that the CDL proofs are tailored at BBS+, but even for the latter scheme, we can achieve savings, as we can think of a BBS+ signature for \mathbf{m} as a BBS signature for (s, \mathbf{m}) , for a secret s , and merely increase k by one.

5.1 Proofs of Knowledge for Signatures

We consider zkPoKs associated with a signature scheme \mathcal{SS} with message space $\mathcal{SS}.\mathcal{M}(\text{par}) = \mathcal{M}(\text{par})^\ell$ for some set \mathcal{M} that can depend on the public parameters par , and some understood vector length $\ell = \ell(\lambda)$. We give proofs of knowledge of a signature consistent with a partial message vector $\mathbf{m} \in (\mathcal{M} \cup \{\star\})^\ell$. For any two such partial messages \mathbf{m}, \mathbf{m}' , we denote $\mathbf{m} \subseteq \mathbf{m}'$ if for all $i \in [\ell]$, $\mathbf{m}[i] \neq \star$ implies $\mathbf{m}[i] = \mathbf{m}'[i]$.

We only consider three-move public-coin protocols between a *prover* and a *verifier*, described by a tuple $\text{PoK} = (\text{PoK.P}_1, \text{PoK.P}_2, \text{PoK.C}, \text{PoK.V})$. Informally, we think of running the protocol in settings where the parameters for \mathcal{SS} are available, i.e., $\text{par} \leftarrow \mathcal{SS}.\text{Setup}(1^\lambda)$, $(\text{sk}, \text{vk}) \leftarrow \mathcal{SS}.\text{KG}(\text{par})$, and the protocol is run as follows, on private input (\mathbf{m}, σ) , where $\mathbf{m} \in \mathcal{M}^\ell$, and public input $\mathbf{m}' \in (\mathcal{M} \cup \{\star\})^\ell$:

Distribution $\text{Real}_{\text{PoK}, \text{SS}}^{\mathcal{A}}(\lambda)$:	Distribution $\text{Ideal}_{\text{PoK}, \text{SS}}^{\mathcal{A}, \mathcal{S}, \mathcal{L}}(\lambda)$:
$\text{par} \leftarrow \text{SS.Setup}(1^\lambda)$	$\text{par} \leftarrow \text{SS.Setup}(1^\lambda)$
$(\text{sk}, \text{vk}) \leftarrow \text{SS.KG}(\text{par})$	$(\text{sk}, \text{vk}) \leftarrow \text{SS.KG}(\text{par})$
$(\mathbf{m}', \mathbf{m}, \sigma) \leftarrow \mathcal{A}(\text{par}, \text{sk}, \text{vk})$	$(\mathbf{m}', \mathbf{m}, \sigma) \leftarrow \mathcal{A}(\text{par}, \text{sk}, \text{vk})$
If $\mathbf{m}' \subseteq \mathbf{m} \wedge \text{SS.Ver}(\text{par}, \text{vk}, (\mathbf{m}, \sigma))$	If $\mathbf{m}' \subseteq \mathbf{m} \wedge \text{SS.Ver}(\text{par}, \text{vk}, (\mathbf{m}, \sigma))$
then	then
$(a, \text{stp}) \leftarrow \text{PoK.P}_1(\text{par}, \text{vk}, \mathbf{m}', (\mathbf{m}, \sigma))$	$z \leftarrow \mathcal{L}(\text{par}, \text{sk})$
$c \leftarrow \text{PoK.C}(\text{par}, \text{vk})$	$(a, c, s) \leftarrow \mathcal{S}(\text{par}, \text{vk}, \mathbf{m}', z)$
$s \leftarrow \text{PoK.P}_2(\text{stp}, c)$	Return $(\text{sk}, \text{vk}, \mathbf{m}', (a, c, s))$
Return $(\text{sk}, \text{vk}, \mathbf{m}', (a, c, s))$	Return \perp
Return \perp	

Fig. 6. Distributions for the definition of HVZK

- (1) The prover initially takes inputs par , vk , and a candidate signature-message pair (\mathbf{m}, σ) , and outputs $(a, \text{stp}) \leftarrow \text{PoK.P}_1(\text{par}, \text{vk}, \mathbf{m}', (\mathbf{m}, \sigma))$. The message a is sent to the verifier.
- (2) The verifier outputs $c \leftarrow \text{PoK.C}(\text{par}, \text{vk})$, and the *challenge* c is sent to the prover.
- (3) The prover outputs $s \leftarrow \text{PoK.P}_2(\text{stp}, c)$, and sends s to the verifier.
- (4) Finally, the verifier outputs a Boolean value $\text{PoK.V}(\text{par}, \text{vk}, \mathbf{m}', a, c, s) \in \{\text{true}, \text{false}\}$.

We say that PoK is correct if, whenever $\text{SS.Ver}(\text{par}, \text{vk}, (\mathbf{m}, \sigma)) = \text{true}$ and $\mathbf{m}' \subseteq \mathbf{m}$, then the verifier also outputs **true**.

SPECIAL SOUNDNESS. We target *special soundness*. To this end, we say that $(\text{par}, \text{vk}, \mathbf{m}', a, c, s)$ is an *accepting* transcript if par is a valid output of SS.Setup , vk is a valid output of $\text{SS.KG}(\text{par})$, c is a valid output of $\text{PoK.C}(\text{par}, \text{vk})$, and $\text{PoK.V}(\text{par}, \text{vk}, \mathbf{m}', a, c, s) = \text{true}$.

Definition 2. We say that PoK as above is special-sound if there exists an efficient algorithm Extract which, given any two valid transcripts $(\text{par}, \text{vk}, \mathbf{m}', a, c, s)$, $(\text{par}, \text{vk}, \mathbf{m}', a, c', s')$ such that $c \neq c'$, then $(\mathbf{m}, \sigma) \leftarrow \text{Extract}(\text{par}, \text{vk}, a, (c, s), (c', s'))$ is such that $\text{SS.Ver}(\text{par}, \text{vk}, (\mathbf{m}, \sigma)) = \text{true}$ and $\mathbf{m}' \subseteq \mathbf{m}$.

We do not specify more general soundness goals further, as the use of special soundness will largely depend on the concrete security game modeling the security of the system using the PoK .

HONEST-VERIFIER ZERO-KNOWLEDGE. The protocols we give will be shown to be honest-verifier zero-knowledge, which suffices for their use as NIZKs via the Fiat-Shamir transform. We will in fact weaken the notion to allow for some *leakage* of the parameters given to the simulator. In particular, we model such leakage as a (possibly randomized) function $\mathcal{L}(\text{par}, \text{sk})$ taking as input the parameters and the signing key.

Definition 3 (HVZK). The protocol PoK for SS as above is perfectly \mathcal{L} -honest-verifier zero-knowledge (\mathcal{L} -HVZK) if there exists an efficient simulator \mathcal{S} such

that for all \mathcal{A} and $\lambda \in \mathbb{N}$, the distributions $\text{Real}_{\text{PoK,SS}}^{\mathcal{A}}(\lambda)$ and $\text{Ideal}_{\text{PoK,SS}}^{\mathcal{A},\mathcal{S},\mathcal{L}}(\lambda)$ given in Figure 6 are identical.

5.2 Protocols

FULL DISCLOSURE. We start with the protocol for the case $\mathbf{m} = \mathbf{m}'$, i.e., the full-disclosure case. Recall that a BBS signature for a message $\mathbf{m} \in \mathbb{Z}_p^\ell$ takes form $\sigma = (A = C(\mathbf{m})^{1/(x+e)}, e)$, where $C(\mathbf{m}) = g_1 \prod_{i=1}^\ell h_1[i]^{\mathbf{m}[i]}$. We assume from now on that $A \neq 1_{\mathbb{G}_1}$, and this assumption is almost without loss of generality, as a valid signature with $A = 1_{\mathbb{G}_1}$ implies finding a non-trivial DLOG relation, as $C(\mathbf{m}) = 1_{\mathbb{G}_1}$ would be true as well. Recall that the signature is valid if

$$\mathbf{e}(A, g_2^e X_2) = \mathbf{e}(C(\mathbf{m}), g_2) ,$$

where $X_2 = g_2^x$ is the verification key. However, by bilinearity,

$$\mathbf{e}(A, g_2^e X_2) = \mathbf{e}(A, g_2^e) \cdot \mathbf{e}(A, X_2) = \mathbf{e}(A^e, g_2) \cdot \mathbf{e}(A, X_2) .$$

And therefore, one can equivalently check that

$$\mathbf{e}(A, X_2) = \mathbf{e}(C(\mathbf{m}) \cdot A^{-e}, g_2) . \quad (9)$$

The main idea is to provide suitably randomized versions of A and $B = C(\mathbf{m}) \cdot A^{-e}$, which can be computed from a signature (A, e) , and extend this with a proof of correctness attesting to the format of these values. In particular, we use $\bar{A} = A^r$ and $\bar{B} = (C(\mathbf{m}) \cdot A^{-e})^r$, for which we still have $e(\bar{A}, X_2) = e(\bar{B}, g_2)$.

PROTOCOL DESCRIPTION. Concretely, we consider the following Σ -protocol:

- Given a signature (A, e) for the message \mathbf{m} with $A \neq 1_{\mathbb{G}_1}$, the prover picks $r \leftarrow \mathbb{Z}_p^*$, and computes

$$\bar{A} \leftarrow A^r , \quad \bar{B} \leftarrow (C(\mathbf{m}) A^{-e})^r = C(\mathbf{m})^r \bar{A}^{-e} .$$

It also picks $\alpha, \beta \leftarrow \mathbb{Z}_p$, and computes $U \leftarrow C(\mathbf{m})^\alpha \bar{A}^\beta$. It sends (\bar{A}, \bar{B}, U) to the verifier.

- The verifier picks a random challenge $c \leftarrow \mathbb{Z}_p$, and sends it to the prover
- The prover responds with (s, t) , where

$$s \leftarrow \alpha + r \cdot c , \quad t \leftarrow \beta - e \cdot c .$$

- The verifier accepts if and only if

$$\mathbf{e}(\bar{A}, X_2) = \mathbf{e}(\bar{B}, g_2) , \quad U \cdot \bar{B}^c = C(\mathbf{m})^s \bar{A}^t .$$

SPECIAL-SOUNDNESS. It is not hard to see that the protocol is special sound. Indeed, given \bar{A}, \bar{B}, U , as well as $c_1 \neq c_2$, and (s_1, s_1, t_1, t_2) such that

$$\mathbf{e}(\bar{A}, X_2) = \mathbf{e}(\bar{B}, g_2), \quad U \cdot \bar{B}^{c_i} = C(\mathbf{m})^{t_i} \bar{A}^{s_i} \quad \text{for } i = 1, 2,$$

we can first extract r and e such that $\bar{B} = C(\mathbf{m})^r \bar{A}^{-e}$, because

$$\bar{B}^{c_1 - c_2} = C(\mathbf{m})^{t_1 - t_2} \bar{A}^{s_1 - s_2},$$

and thus we can set $r = (t_1 - t_2)/(c_1 - c_2)$ and $e = (s_2 - s_1)/(c_1 - c_2)$. If $r \neq 0$, then $(A = \bar{A}^{r^{-1}}, e)$ is a valid signature on \mathbf{m} , because $\mathbf{e}(\bar{A}, X_2) = \mathbf{e}(\bar{B}, g_2)$ implies that $\mathbf{e}(\bar{A}^{r^{-1}}, X_2) = \mathbf{e}(\bar{B}^{r^{-1}}, g_2)$, and $\bar{B}^{r^{-1}} = C(\mathbf{m})A^{-e}$. If $r = 0$, then $e(\bar{A}, X_2) = e(\bar{A}^{-e}, g_2)$, which means $x = -e$, and this gives us a signature on \mathbf{m} .

ZERO-KNOWLEDGE. The protocol is \mathcal{L} -HVZK, for \mathcal{L} which, on input g_1, x , outputs (g_1^r, g_1^{rx}) for $r \leftarrow \mathbb{Z}_p^*$, i.e., a random pair of form (U, U^x) . The simulator then computes $\bar{A} \leftarrow \mathbb{G}_1^*$, and set $\bar{C} = \bar{A}^x \in \mathbb{G}_1$ – this can be done by re-randomizing the leakage (U, U^x) . Then, the simulator picks a random challenge $c \leftarrow \mathbb{Z}_p$, as well as random $s, t \leftarrow \mathbb{Z}_p$, and sets $U = C(\mathbf{m})^s \bar{A}^t \bar{B}^{-c}$.

The fact that the simulator needs a sample (U, U^x) , and cannot simulate solely given the parameters and the verification key $X_2 = g_2^x$ is a technical oddity inherited from the use of type-3 pairings, and was also present in prior protocols [15]. Indeed, it is hard to compute g_1^x from the verification key g_2^x . However, this additional leakage is not really harmful. For example, *any* signature (A, e) on a message \mathbf{m} already satisfies $A^x = C(\mathbf{m}) \cdot A^{-e}$, and thus the protocol leaks no more than *any* valid message-signature pair. In particular, BBS remains secure given such leakage.

PARTIAL DISCLOSURE. For the case $\mathbf{m}' \subsetneq \mathbf{m}$, the components $\mathbf{m}[i]$ for which $\mathbf{m}'[i] = \star$ become parts of the witness. We let $I := \{i \in [\ell] : \mathbf{m}[i] = \star\}$ and $J = [n] \setminus I$. We also let $C_J(\mathbf{m}) = g_1 \prod_{i \in J} h_1[i]^{\mathbf{m}[i]}$, and note that $C_J(\mathbf{m})$ can be computed from the public input \mathbf{m}' by the verifier.

- Given a signature (A, e) for the message \mathbf{m} with $A \neq 1_{\mathbb{G}_1}$, the prover picks $r \leftarrow \mathbb{Z}_p^*$, and computes

$$\bar{A} \leftarrow A^r, \quad \bar{B} \leftarrow (C(\mathbf{m})A^{-e})^r = C_J(\mathbf{m})^r \cdot \left(\prod_{i \in I} h_1[i]^{\mathbf{m}[i]} \right)^r \cdot \bar{A}^{-e}.$$

It also picks $\alpha, \beta \leftarrow \mathbb{Z}_p$, and also $\delta_i \leftarrow \mathbb{Z}_p$ for every $i \in I$ and computes

$$U \leftarrow C_J(\mathbf{m})^\alpha \cdot \bar{A}^\beta \cdot \prod_{i \in I} h_1[i]^{\delta_i}$$

It sends (\bar{A}, \bar{B}, U) to the verifier.

- The verifier picks a random challenge $c \leftarrow \mathbb{Z}_p$, and sends it to the prover

- The prover responds with $(s, t, (u_i)_{i \in I})$, where

$$s \leftarrow \alpha + r \cdot c, \quad t \leftarrow \beta - e \cdot c, \quad u_i \leftarrow \delta_i + r \cdot \mathbf{m}[i] \cdot c \quad \forall i \in I.$$

- The verifier accepts if and only if

$$\mathbf{e}(\bar{A}, X_2) = \mathbf{e}(\bar{B}, g_2), \quad U \cdot \bar{B}^c = C_J(\mathbf{m})^s \bar{A}^t \prod_{i \in I} h_1[i]^{u_i}.$$

One can easily adapt the arguments for the above protocols for full disclosure to show special soundness and \mathcal{L} -HVZK.

NIZKs. Our protocols can be transformed into NIZKs in the random oracle model via the Fiat-Shamir transform [21] or Fischlin's transform [22]. For the Fiat-Shamir version, the prover computes \bar{A}, \bar{B}, U as above, then lets $c \leftarrow H(\mathbf{m}', \bar{A}, \bar{B}, U)$, and finally computes $s, t, (u_i)_{i \in I}$ as above. The final proof is

$$\pi = (\bar{A}, \bar{B}, c, s, t, (u_i)_{i \in I}).$$

Verification checks that $\mathbf{e}(\bar{A}, X_2) = \mathbf{e}(\bar{B}, g_2)$ and that $c = H(\mathbf{m}', \bar{A}, \bar{B}, U)$ with

$$U \leftarrow \bar{B}^{-c} C_J(\mathbf{m})^s \bar{A}^t \prod_{i \in I} h_1[i]^{u_i}.$$

Note that we could include U instead of c , but this leads to longer proofs for curves like BLS12-381, where elements in \mathbb{G}_1 have longer descriptions than scalars.

6 Signatures for Group Elements and Blind Issuance

One central property of BBS is its support of *blind issuance*, the setting where a user sends a commitment $C \in \mathbb{G}_1$ to the signer to obtain a pair $\sigma = (A, e)$ with $A = C^{\frac{1}{x+e}}$ —if $C = g_1 \prod_{i=1}^{\ell} h_1[i]^{\mathbf{m}[i]}$ for a message \mathbf{m} , then σ is a valid signature on \mathbf{m} , but crucially, the signer never learns \mathbf{m} . In fact, the user could make $\mathbf{m}[1]$ uniform, turning C into a perfectly-hiding (generalized) Pedersen commitment [31]. This approach is particularly important when σ acts as a credential, and we want to hide the actual attributes from the issuer. Blind issuance of BBS signatures is also part of an unofficial draft [3], which also requires the addition of a proof of knowledge for a representation of C , which consists of $O(\ell)$ scalars and can be expensive when ℓ is large. Here, we show that in the AGM the scheme is already sufficiently secure without such a proof. A suitable proof of knowledge is however still necessary if the user needs to reveal part of the attributes to the issuer, to prove these are consistent with the commitment. However, we note that this aspect would be orthogonal to our analysis below.

Game $\text{OMUF} +_{\text{GGen}, \text{eG}, \text{eS}}^A(\lambda)$: $\text{cnt} \leftarrow 0$ $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathbf{e}) \leftarrow \text{GGen}(1^\lambda)$ $g_1 \leftarrow \mathbb{G}_1^*, \mathbf{h} \leftarrow \mathbb{G}_1, g_2 \leftarrow \mathbb{G}_2^*$ $\text{par} \leftarrow (p, g_1, \mathbf{h}_1, g_2, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathbf{e})$ $\text{st}_e \leftarrow \text{eG}(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \mathbf{e})$ $x \leftarrow \mathbb{Z}_p; X_2 \leftarrow g_2^x; \text{sk} \leftarrow x; \text{vk} \leftarrow X_2$ $\{(\mathbf{m}_i^*, (A_i^*, e_i^*))\}_{i \in [q']} \leftarrow \mathcal{A}^{\text{SDH}}(\text{par}, \text{st}_e, \text{vk})$ For all $i \in [q']$ do $C_i^* \leftarrow g_1 \prod_{j \in [\ell]} \mathbf{h}_1[j]^{\mathbf{m}_i^*[j]}$ If $\forall i \in [q']$: $\mathbf{e}\left(A_i^*, X_2 g_2^{e_i^*}\right) = \mathbf{e}(C_i^*, g_2)$ then Return $(q' > \text{cnt})$ Return false	Oracle $\text{SDH}(C)$: $\text{cnt} \leftarrow \text{cnt} + 1$ $e_{\text{cnt}} \leftarrow \text{eS}(\text{st}_e, \text{cnt})$ $A_{\text{cnt}} \leftarrow C^{\frac{1}{x + e_{\text{cnt}}}}$ Return A_{cnt}
---	--

Fig. 7. One-more unforgeability of BBS. This game captures the one-more unforgeability of BBS when given an SDH oracle which returns $C^{1/(x+e_i)}$ for its i -th query, where e_i is generated via eS . We assume here that \mathcal{A} returns a set of q' distinct forgery attempts (i.e., no double entry are present in the list returned by \mathcal{A} .)

ONE-MORE UNFORGEABILITY. BBS can be thought as a signature scheme signing a group element $C \in \mathbb{G}_1$ as $\sigma = (A = C^{\frac{1}{x+e}}, e)$. However, it does not achieve unforgeability when signing group elements (as in the case of *structure-preserving signatures* (SPS) [5]). Indeed, the attacker, given $\sigma = (A, e)$, directly obtains other valid signatures, such as $\sigma' = (A^2, e)$, which is a valid signature for $C^2 \neq C$. Nonetheless, if $C = g_1 \prod_{i=1}^{\ell} \mathbf{h}_1[i]^{\mathbf{m}[i]}$, it is very unlikely that the attacker can exhibit a message \mathbf{m}' such that $C^2 = g_1 \prod_{i=1}^{\ell} \mathbf{h}_1[i]^{\mathbf{m}'[i]}$, i.e., such that (A^2, e) is valid for \mathbf{m}' .

We formalize this by showing BBS satisfies *one-more unforgeability* (OMUF), where given access q times to an oracle SDH that signs group elements as above—i.e., on input $C \in \mathbb{G}_1$ it returns $\sigma = (A, e)$ with $A = C^{\frac{1}{x+e}}$ —it is impossible for the attacker to come up with $q + 1$ valid BBS signatures. This property is defined via Game $\text{OMUF} +_{\text{GGen}, \text{eG}, \text{eS}}^A(\lambda)$ in Figure 7. Similar to Section 4, the game is parameterized by the group generator GGen and by a pair of algorithms eG, eS used to generate the e_i 's ahead of time. We also define

$$\text{Adv}_{\text{GGen}, \text{eG}, \text{eS}}^{\text{omuf}+}(\mathcal{A}, \lambda) = \Pr[\text{OMUF} +_{\text{GGen}, \text{eG}, \text{eS}}^A(\lambda)] .$$

We stress that we *could* define a general notion of signatures on commitment values, and require that upon obtaining q signatures on arbitrary elements from the commitment space, the attacker cannot come up with $q + 1$ valid signatures on commitments, along with their openings. However, we prefer the rather straightforward BBS-specific game as a better illustration of this property.

MAIN RESULT. We prove now that BBS satisfies one-more unforgeability in the AGM, and we do so via a reduction to its $\text{SUF}+$ security as defined in Section 4.

Theorem 3 (One-more unforgeability). *Let GGen be a group parameter generator, producing groups of order $p(\lambda)$, and let eG, eS as above. For every*

algebraic OMUF+ adversary \mathcal{A} issuing at most $q = q(\lambda)$ SDH queries, there exists an algebraic SUF+ adversary \mathcal{B} such that

$$\text{Adv}_{\text{GGen}, \text{eG}, \text{eS}}^{\text{omuf}+}(\mathcal{A}, \lambda) \leq \text{Adv}_{\text{GGen}, \text{eG}, \text{eS}}^{\text{suf}+}(\mathcal{B}, \lambda) + \delta_{\text{eG}, \text{eS}}(q, \lambda).$$

The adversary \mathcal{B} issues q SIGN queries, and runs in time equal that of running \mathcal{A} , plus the time needed to perform $O(q^3)$ operations in \mathbb{Z}_p .

The proof is deferred to the full version due to lack of space. The main challenge in the proof is to show how the signing oracle can be used to simulate signing a group element, given its representation. This is easy to do if the representation is only in terms of g_1 and \mathbf{h}_1 , but the challenge is that the representation can also depend on prior signatures.

Theorem 3 yields the following corollary when combined with Theorem 2.

Corollary 1 (One-more unforgeability). *Let GGen be a group parameter generator, producing groups of order $p(\lambda)$, and let eG, eS as above. For every algebraic OMUF+ adversary \mathcal{A} issuing at most q SDH queries, there exists a q -DL adversary \mathcal{C}_1 and a DL adversary \mathcal{C}_2 such that*

$$\text{Adv}_{\text{GGen}, \text{eG}, \text{eS}}^{\text{omuf}+}(\mathcal{A}, \lambda) \leq \text{Adv}_{\text{GGen}}^{q\text{-dl}}(\mathcal{C}_1, \lambda) + \text{Adv}_{\text{GGen}}^{\text{dl}}(\mathcal{C}_2, \lambda) + 2\delta_{\text{eG}, \text{eS}}(q, \lambda) + \frac{1}{p(\lambda)}.$$

The adversaries \mathcal{C}_1 and \mathcal{C}_2 are obtained by using \mathcal{B} from Theorem 3 within the adversaries of Theorem 2.

APPLICATIONS. As mentioned above, a typical application of BBS signatures is in the context of credentials. The above result validates the security of the canonical solution where the user obtains a credential for a vector of attributes \mathbf{m} by sending $C = g_1 \prod_{i=1}^{\ell} \mathbf{h}_1[i]^{\mathbf{m}[i]}$ to the authority, which in turn responds with the actual credential $(C^{\frac{1}{x+e}}, e)$ for a random e . The OMUF security from Theorem 3 and Corollary 1 implies that a malicious user (or any set of multiple such users) can only obtain q credentials by interacting with the authority q times. The user can then show the credential multiple times in an unlinkable way by using the zk-PoKs from Section 5, typically compiled via the Fiat-Shamir transform. These showings are then consistent with at most q attribute vectors. When issuing a credential, the user does not need to send any proof of knowledge along with C , unless the credential issuing needs to enforce some format on the values contained by C , in which case extra proofs need to be sent along.

We note that analyzing the security of the entire credential system is non-trivial, especially if we want to resort to PoKs compiled via the Fiat-Shamir transform, which are not online extractable. We believe that a security analysis of variants of this system is however possible, albeit very tedious, in the AGM, where one can resort to the online-extractability of the proposed PoKs from Section 5 in the AGM, along the lines of [25]. This goes however beyond the scope of this paper.

Acknowledgments

We wish to thank Christian Paquin and Greg Zaverucha for extensive discussions around BBS and for providing feedback throughout this project. We also thank the EUROCRYPT 2023 reviewers for their excellent comments and suggestions. This research was partially supported by NSF grants CNS-2026774, CNS-2154174, a JP Morgan Faculty Award, a CISCO Faculty Award, and a gift from Microsoft.

References

1. BBS+ implementation. <https://github.com/mattrglobal/bbs-signatures>, accessed: 2022-10-04
2. BBS+ implementation. <https://github.com/microsoft/bbs-node-reference>, accessed: 2022-10-04
3. Blind signatures extension of the BBS signature scheme. <https://identity.foundation/bbs-signature/draft-blind-bbs-signatures.txt>, accessed: 2022-10-04
4. Cheon's attack and its effect on the security of big trusted setups. <https://ethresear.ch/t/cheons-attack-and-its-effect-on-the-security-of-big-trusted-setups/6692>, accessed: 2022-10-04
5. Abe, M., Fuchsbauer, G., Groth, J., Haralambiev, K., Ohkubo, M.: Structure-preserving signatures and commitments to group elements. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 209–236. Springer, Heidelberg (Aug 2010). https://doi.org/10.1007/978-3-642-14623-7_12
6. Au, M.H., Susilo, W., Mu, Y.: Constant-size dynamic k-TAA. In: Prisco, R.D., Yung, M. (eds.) SCN 06. LNCS, vol. 4116, pp. 111–125. Springer, Heidelberg (Sep 2006). https://doi.org/10.1007/11832072_8
7. Barreto, P.S.L.M., Lynn, B., Scott, M.: Constructing elliptic curves with prescribed embedding degrees. In: Ciamato, S., Galdi, C., Persiano, G. (eds.) SCN 02. LNCS, vol. 2576, pp. 257–267. Springer, Heidelberg (Sep 2003). https://doi.org/10.1007/3-540-36413-7_19
8. Bauer, B., Fuchsbauer, G., Loss, J.: A classification of computational assumptions in the algebraic group model. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 121–151. Springer, Heidelberg (Aug 2020). https://doi.org/10.1007/978-3-030-56880-1_5
9. Boneh, D., Boyen, X.: Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology* **21**(2), 149–177 (Apr 2008). <https://doi.org/10.1007/s00145-007-9005-7>
10. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (Aug 2004). https://doi.org/10.1007/978-3-540-28628-8_3
11. Bowe, S.: BLS12-381: New zk-SNARK elliptic curve construction. <https://electriccoin.co/blog/new-snark-curve/> (2017)
12. Brickell, E., Li, J.: A pairing-based DAA scheme further reducing TPM resources. In: Acquisti, A., Smith, S.W., Sadeghi, A. (eds.) Trust and Trustworthy Computing, Third International Conference, TRUST 2010, Berlin, Germany, June 21–

- 23, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6101, pp. 181–195. Springer (2010). https://doi.org/10.1007/978-3-642-13869-0_12, https://doi.org/10.1007/978-3-642-13869-0_12
13. Brickell, E., Li, J.: Enhanced privacy ID from bilinear pairing for hardware authentication and attestation. *Int. J. Inf. Priv. Secur. Integr.* **1**(1), 3–33 (2011). <https://doi.org/10.1504/IJIPSI.2011.043729>, <https://doi.org/10.1504/IJIPSI.2011.043729>
14. Brown, D.R.L., Gallant, R.P.: The static Diffie-Hellman problem. *Cryptology ePrint Archive*, Report 2004/306 (2004), <https://eprint.iacr.org/2004/306>
15. Camenisch, J., Drijvers, M., Lehmann, A.: Anonymous attestation using the strong diffie hellman assumption revisited. In: Franz, M., Papadimitratos, P. (eds.) *Trust and Trustworthy Computing - 9th International Conference, TRUST 2016, Vienna, Austria, August 29-30, 2016, Proceedings. Lecture Notes in Computer Science*, vol. 9824, pp. 1–20. Springer (2016). https://doi.org/10.1007/978-3-319-45572-3_1, https://doi.org/10.1007/978-3-319-45572-3_1
16. Camenisch, J., Lysyanskaya, A.: A signature scheme with efficient protocols. In: Ciamato, S., Galdi, C., Persiano, G. (eds.) *SCN 02. LNCS*, vol. 2576, pp. 268–289. Springer, Heidelberg (Sep 2003). https://doi.org/10.1007/3-540-36413-7_20
17. Camenisch, J., Lysyanskaya, A.: Signature schemes and anonymous credentials from bilinear maps. In: Franklin, M. (ed.) *CRYPTO 2004. LNCS*, vol. 3152, pp. 56–72. Springer, Heidelberg (Aug 2004). https://doi.org/10.1007/978-3-540-28628-8_4
18. Chen, L.: A DAA scheme requiring less TPM resources. In: Bao, F., Yung, M., Lin, D., Jing, J. (eds.) *Information Security and Cryptology - 5th International Conference, Inscrypt 2009, Beijing, China, December 12-15, 2009. Revised Selected Papers. Lecture Notes in Computer Science*, vol. 6151, pp. 350–365. Springer (2009). https://doi.org/10.1007/978-3-642-16342-5_26, https://doi.org/10.1007/978-3-642-16342-5_26
19. Chen, S., Steinberger, J.P.: Tight security bounds for key-alternating ciphers. In: Nguyen, P.Q., Oswald, E. (eds.) *EUROCRYPT 2014. LNCS*, vol. 8441, pp. 327–350. Springer, Heidelberg (May 2014). https://doi.org/10.1007/978-3-642-55220-5_19
20. Cheon, J.H.: Security analysis of the strong Diffie-Hellman problem. In: Vaude- nay, S. (ed.) *EUROCRYPT 2006. LNCS*, vol. 4004, pp. 1–11. Springer, Heidelberg (May / Jun 2006). https://doi.org/10.1007/11761679_1
21. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) *CRYPTO’86. LNCS*, vol. 263, pp. 186–194. Springer, Heidelberg (Aug 1987). https://doi.org/10.1007/3-540-47721-7_12
22. Fischlin, M.: Communication-efficient non-interactive proofs of knowledge with online extractors. In: Shoup, V. (ed.) *CRYPTO 2005. LNCS*, vol. 3621, pp. 152–168. Springer, Heidelberg (Aug 2005). https://doi.org/10.1007/11535218_10
23. Fuchsbaauer, G., Kiltz, E., Loss, J.: The algebraic group model and its applications. In: Shacham, H., Boldyreva, A. (eds.) *CRYPTO 2018, Part II. LNCS*, vol. 10992, pp. 33–62. Springer, Heidelberg (Aug 2018). https://doi.org/10.1007/978-3-319-96881-0_2
24. Galbraith, S., Paterson, K., Smart, N.: Pairings for cryptographers. *Cryptology ePrint Archive*, Report 2006/165 (2006), <https://eprint.iacr.org/2006/165>
25. Ghoshal, A., Tessaro, S.: Tight state-restoration soundness in the algebraic group model. In: Malkin, T., Peikert, C. (eds.) *CRYPTO 2021, Part III. LNCS*, vol.

- 12827, pp. 64–93. Springer, Heidelberg, Virtual Event (Aug 2021). https://doi.org/10.1007/978-3-030-84252-9_3
26. Hoang, V.T., Tessaro, S.: Key-alternating ciphers and key-length extension: Exact bounds and multi-user security. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 3–32. Springer, Heidelberg (Aug 2016). https://doi.org/10.1007/978-3-662-53018-4_1
 27. Jao, D., Yoshida, K.: Boneh-Boyen signatures and the strong Diffie-Hellman problem. In: Shacham, H., Waters, B. (eds.) PAIRING 2009. LNCS, vol. 5671, pp. 1–16. Springer, Heidelberg (Aug 2009). https://doi.org/10.1007/978-3-642-03298-1_1
 28. Looker, T., Kalos, V., Whitehead, A., Lodder, M.: The BBS Signature Scheme. Internet-Draft draft-irtf-cfrg-bbs-signatures-01, Internet Engineering Task Force (Oct 2022), <https://datatracker.ietf.org/doc/draft-irtf-cfrg-bbs-signatures/01/>, work in Progress
 29. Maurer, U.M.: Unifying zero-knowledge proofs of knowledge. In: Preneel, B. (ed.) AFRICACRYPT 09. LNCS, vol. 5580, pp. 272–286. Springer, Heidelberg (Jun 2009)
 30. Patarin, J.: A proof of security in $O(2^n)$ for the Benes scheme. In: Vaudenay, S. (ed.) AFRICACRYPT 08. LNCS, vol. 5023, pp. 209–220. Springer, Heidelberg (Jun 2008)
 31. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) CRYPTO’91. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (Aug 1992). https://doi.org/10.1007/3-540-46766-1_9
 32. Pointcheval, D., Sanders, O.: Short randomizable signatures. In: Sako, K. (ed.) CT-RSA 2016. LNCS, vol. 9610, pp. 111–126. Springer, Heidelberg (Feb / Mar 2016). https://doi.org/10.1007/978-3-319-29485-8_7
 33. Sonnino, A., Al-Bassam, M., Bano, S., Meiklejohn, S., Danezis, G.: Coconut: Threshold issuance selective disclosure credentials with applications to distributed ledgers. In: NDSS 2019. The Internet Society (Feb 2019)