

On the Multi-User Security of Short Schnorr Signatures with Preprocessing

Jeremiah Blocki^[0000-0002-5542-4674] and Seunghoon Lee^[0000-0003-4475-5686]

Purdue University, West Lafayette, IN, 47906, USA
{jblocki, lee2856}@purdue.edu

Abstract. The Schnorr signature scheme is an efficient digital signature scheme with short signature lengths, i.e., $4k$ -bit signatures for k bits of security. A Schnorr signature σ over a group of size $p \approx 2^{2k}$ consists of a tuple (s, e) , where $e \in \{0, 1\}^{2k}$ is a hash output and $s \in \mathbb{Z}_p$ must be computed using the secret key. While the hash output e requires $2k$ bits to encode, Schnorr proposed that it might be possible to truncate the hash value without adversely impacting security.

In this paper, we prove that *short* Schnorr signatures of length $3k$ bits provide k bits of multi-user security in the (Shoup's) generic group model and the programmable random oracle model. We further analyze the multi-user security of key-prefixed short Schnorr signatures against preprocessing attacks, showing that it is possible to obtain secure signatures of length $3k + \log S + \log N$ bits. Here, N denotes the number of users and S denotes the size of the hint generated by our preprocessing attacker, e.g., if $S = 2^{k/2}$, then we would obtain secure $3.75k$ -bit signatures for groups of up to $N \leq 2^{k/4}$ users.

Our techniques easily generalize to several other Fiat-Shamir-based signature schemes, allowing us to establish analogous results for Chaum-Pedersen signatures and Katz-Wang signatures. As a building block, we also analyze the 1-out-of- N discrete-log problem in the generic group model, with and without preprocessing.

1 Introduction

The Schnorr signature scheme [Sch90] has been widely used due to its simplicity, efficiency and short signature size. In the Schnorr signature scheme, we start with a cyclic group $G = \langle g \rangle$ of prime order p and pick a random secret key $\text{sk} \in \mathbb{Z}_p$. To sign a message m , we pick $r \in \mathbb{Z}_p$ uniformly at random, compute $I = g^r$, $e = \text{H}(I||m)$, and $s = r + \text{sk} \cdot e \pmod p$. Then, the final signature is $\sigma = (s, e)$.

We recall that a signature scheme Π yields k bits of (multi-user) security if any attacker running in time at most t can forge a signature with probability at most $\varepsilon_t = t/2^k$ in the (multi-user) signature forgery game, and this should hold for all time bounds $t \leq 2^k$. To achieve k bits of security, we select a hash function H with $2k$ -bit outputs, and we select p to be a random $2k$ -bit prime so that the length of a signature is $4k$ bits.

In Schnorr’s original paper [Sch90], the author proposed the possibility of achieving even shorter Schnorr signatures by selecting a hash function H with k -bit outputs (or truncating to only use the first k bits) so that the final signature $\sigma = (s, e)$ can be encoded with $3k$ bits. We refer to this signature scheme as the *short* Schnorr signature scheme. In this paper, we investigate the following questions:

Does the short Schnorr signature scheme achieve k bits of (multi-user) security? If so, is the short Schnorr signature scheme also secure against preprocessing attacks?

Proving security for the Schnorr signatures has been a challenging task against the interactive attacks. Pointcheval and Stern [PS96] provided a reduction from the discrete-log problem in the random oracle model [BR93]. However, their reduction is not tight, i.e., they show that $\text{Adv}_{\text{sig}} \leq \text{Adv}_{\text{dlog}} \times q_H$ for any attacker making at most q_H queries to the random oracle. The loss of the factor q_H , which prevents us from concluding that the scheme provides k bits of security, seems to be unavoidable, e.g., see [Seu12, FJS14].

Neven et al. [NSW09] analyzed Schnorr signatures in the generic group model [Sho97], showing that the scheme provides k bits of security as long as the hash function satisfies two key properties: random-prefix preimage (rpp) and random-prefix second-preimage (rpsp) security. Interestingly, Neven et al. do not need to assume that H is a random oracle, though a random oracle H would satisfy both rpp and rpsp security. Neven et al. considered the short Schnorr signature scheme, but their upper bounds do not allow us to conclude that the short Schnorr signature scheme provides k bits of security¹.

An earlier paper of Schnorr and Jakobsson [SJ00] analyzed the security of the short Schnorr signature scheme in the random oracle model plus the generic group model. While they show that the scheme provides k bits of security, they also consider another version of the generic group model which is different from the definition proposed by Shoup [Sho97]. The reason is that the version they consider is not expressive enough to capture all known attacks, e.g., any attack that requires the ability to hash group elements including preprocessing attacks of Corrigan-Gibbs and Kogan [CK18] cannot be captured in their generic group model. See Appendix B for further discussion if interested.

Galbraith et al. [GMLS02] claimed to have a tight reduction showing that single-user security implies multi-user security of the regular Schnorr signature

¹ Neven et al. [NSW09, Section 7] present two upper bounds on the success probability of a Schnorr signature forgery attacker under plausible assumptions about rpp, rpsp and discrete log security. The first upper bound is $\mathcal{O}\left(\frac{t^2}{2^n} + \frac{t^2}{p}\right)$ where p is the size of the group and n denotes the length of a hash output (# bits). The second bound is $\mathcal{O}\left(\sqrt{\frac{t^3}{p}} + \frac{t}{2^n} + \frac{t^2}{p}\right)$. To achieve k -bit security with the second bound we can use shorter hashes $n = k$, but we must use a larger group $p \approx 2^{3k}$. By contrast, to achieve k -bit security with the first bound we can use a smaller group size $p \approx 2^{2k}$, but we require longer hash outputs $n = 2k$. See additional discussion in Appendix D.

scheme. However, Bernstein [Ber15] identified an error in the security proof in [GMLS02], proposed a modified “key-prefixed” version of the original Schnorr signature scheme (including the public key as a hash input), and proved that the “key-prefixed” version does provide multi-user security. Derler and Slamanig [DS19] later showed a tight reduction from single-user security to “key-prefixed” multi-user security for a class of key-homomorphic signature schemes including Schnorr signatures. The Internet Engineering Task Force (IETF) adopted the key-prefixed modification of Schnorr signatures to ensure multi-user security [Hao17]. Kiltz et al. [KMP16] later gave a tight security reduction establishing multi-user security of regular Schnorr signatures in the programmable random oracle model plus (another version of) the generic group model without key-prefixing. Our results imply that key-prefixing is not even necessary to establish tight multi-user security of short Schnorr signatures². On the other hand key-prefixing is both necessary and sufficient to establish multi-user security of (short) Schnorr signatures against preprocessing attackers.

1.1 Our Contributions

We show that the *short* Schnorr Signature scheme provides k bits of security against an attacker in *both* the single and multi-user versions of the signature forgery game. Our results assume the programmable random oracle model and the (Shoup’s) generic group model. We further analyze the *multi-user* security of key-prefixed short Schnorr signatures against preprocessing attacks. The preprocessing attacker outputs a hint of size S after making as many as 2^{3k} queries to the random oracle and examining the entire generic group oracle. Later on, the online attacker can use the hint to help win the *multi-user* signature forgery game by forging a signature for *any one* of the N users. By tuning the parameters of the key-prefixed short Schnorr signature scheme appropriately, we can obtain $(3k + \log S + \log N)$ -bit signatures with k bits of security against a preprocessing attacker.

Single-User Security of Short Schnorr Signatures. As a warm-up, we first consider the single-user security of the *short* Schnorr Signature scheme without preprocessing, showing that short Schnorr signatures provide k bits of security in the (Shoup’s) generic group model and the random oracle model.

Theorem 1 (informal). *Any attacker making at most q queries wins the signature forgery game (chosen message attack) against the short Schnorr signature scheme with probability at most $\mathcal{O}(q/2^k)$ in the generic group model (of order $p \approx 2^{2k}$) plus programmable random oracle model (See [Definition 1](#) and [Theorem 4](#)).*

² The authors of [KMP16] pointed out that their analysis can be adapted to demonstrate multi-user security of short Schnorr signatures (private communication) though the paper itself never discusses short Schnorr signatures. Furthermore, their proof is in a different version of the generic group model which is not suitable for analyzing preprocessing attacks. See discussion in [Remark 1](#).

Theorem 1 tells us that the short Schnorr signature obtained by truncating the hash output by half would yield the same k bits of security level with the signature length $3k$, instead of $4k$. A 25% reduction in signature length is particularly significant in contexts where space/bandwidth is limited, e.g., on the blockchain.

Multi-User Security of Short Schnorr Signatures. We show that our proofs can be extended to the multi-user case *even* in the so-called “1-out-of- N ” setting, i.e., if the attacker is given N public keys $\text{pk}_1, \dots, \text{pk}_N$, s/he can forge a signature σ which is valid under any one of these public keys (it does not matter which).

Theorem 2 (informal). *Let N denote the number of distinct users/public keys. Then any attacker making at most q queries wins the multi-user signature forgery game (chosen message attack) against the short Schnorr signature scheme with probability at most $\mathcal{O}((q + N)/2^k)$ in the generic group model (of order $p \approx 2^{2k}$) plus programmable random oracle model (See [Definition 2](#) and [Theorem 6](#)).*

Theorem 2 guarantees that breaking multi-user security of short Schnorr signatures in the 1-out-of- N setting is *not easier* than breaking a single instance, as the winning probability is still in the same order as long as $N \leq q$, which is the typical case. A naïve reduction loses a factor of N , i.e., any attacker winning the multi-user forgery game with probability ε_{MU} can be used to win the single-user forgery game with probability $\varepsilon \geq \varepsilon_{\text{MU}}/N$. For example, suppose that $p \approx 2^{224}$ (i.e., $k = 112$), and there are $N = 2^{32}$ instances of short Schnorr signatures, which is more than the half of the entire world population. In the original single-user security game, an attacker wins with probability at most $\varepsilon \leq \mathcal{O}(t/2^k)$, so an attacker running in time $t = 2^{80}$ would succeed with probability at most $\varepsilon \approx 2^{-32}$. This only allows us to conclude that an attacker succeeds with probability at most $\varepsilon_{\text{MU}} \leq N\varepsilon \approx 1$ in the multi-user security game! Our security proof implies that the attacker will succeed with probability $\varepsilon_{\text{MU}} \approx \varepsilon \approx 2^{-32}$ in the above example. In particular, we don’t lose a factor of N in the security reduction.

Security of Key-Prefixed Short Schnorr Signatures against Preprocessing Attacks. We further show that *key-prefixed short Schnorr* signatures are also secure against preprocessing attacks. Here, we consider a key-prefixed version of Schnorr signatures, because regular Schnorr signatures are trivially vulnerable to preprocessing attacks, e.g., if a preprocessing attacker finds some message m and an integer r such that $e = \text{H}(g^r \| m) = 0$, then $\sigma = (r, 0)$ is *always* a valid signature for *any* public key $\text{pk} = g^{\text{sk}}$ since $g^{r - \text{sk} \cdot 0} = g^r$. We note that several standardized implementations of Schnorr signatures (i.e., BSI [[fIS18](#)] or ISO/IEC [[fSC18](#)]) slightly deviate from Schnorr’s original construction and explicitly disallowing $e = 0$ signatures which defends against our particular preprocessing attack – see [Remark 2](#).

We consider a preprocessing attacker who may query the random oracle at up to 2^{3k} points and may also examine the entire generic group oracles before

outputting an S -bit hint for the online attacker. We leave it as an interesting open question whether or not the restriction on the number of random oracle queries is necessary. However, from a practical standpoint, we argue that a preprocessing adversary will never be able to make 2^{2k} queries, e.g., if $k \geq 128$, then 2^{2k} operations is already far too expensive for even a nation-state attacker.

Theorem 3 (informal). *Let N denote the number of distinct users/public keys. Then any preprocessing attacker making at most q_{pre} queries and outputs an S -bit hint during the preprocessing phase and making at most q_{on} queries during the online phase wins the multi-user signature forgery game (chosen message attack) against the short Schnorr signature scheme with probability at most $\tilde{\mathcal{O}}(SN(q_{\text{on}} + N)^2/p + q_{\text{on}}/2^k + Nq_{\text{pre}}q_{\text{on}}/p^2)$ in the generic group model of order $p > 2^{2k}$ plus programmable random oracle model (see [Theorem 8](#)).*

[Theorem 3](#) tells us that with suitable parameter setting, key-prefixed short Schnorr signatures also achieve k bits of multi-user security even against preprocessing attacks. In particular, by setting $p \approx 2^{2k}SN$ and maintaining k -bit hash outputs, the short Schnorr signature scheme still maintains k bits of multi-user security against our preprocessing attacker. For example, if $S = 2^{k/2}$ and $N = 2^{k/4}$, then setting $p \approx 2^{2.75k}$ yields signatures of length $k + \log p = 3.75k$. Up to a factor N , the results from [Theorem 3](#) are tight as a preprocessing attacker can succeed with probability at least Sq_{on}^2/p .

Other Fiat-Shamir Signatures. Using similar reductions, we establish similar security bounds for the full-domain hash variant of (key-prefixed) Chaum-Pedersen signatures [[CP93](#)] and for Katz-Wang signatures [[KW03](#)] with truncated hash outputs. In particular, a preprocessing attacker wins the multi-user signature forgery game with probability at most $\mathcal{O}(SNq^2/p + q/2^k)$ – see [Theorem 10](#) and [Theorem 12](#) in [Section 6](#). Short Katz-Wang signatures [[KW03](#)] have the same length as short Schnorr signatures with equivalent security guarantees, while Chaum-Pedersen signatures are a bit longer.

1.2 Our Techniques

The Multi-User Bridge-Finding Game. We introduce an intermediate problem called the 1-out-of- N bridge-finding game. Oversimplifying a bit, in a cyclic group $G = \langle g \rangle$, the attacker is given N inputs g^{x_1}, \dots, g^{x_N} , and the goal of the attacker is to produce a non-trivial linear dependence, i.e., a_1, \dots, a_N and b such that $a_1x_1 + \dots + a_Nx_N = b$, and $a_i \neq 0$ for at least one $i \leq N$. We then show that in the generic group model, an attacker making at most q generic group queries can succeed with probability at most $\mathcal{O}(q^2/p + qN/p)$, where $p \approx 2^{2k}$ is the size of the group.

We also show that a preprocessing attacker wins the 1-out-of- N bridge-finding game with probability at most $\mathcal{O}(SNq^2 \log p/p)$ when given an arbitrary S -bit hint fixed a priori before x_1, \dots, x_N are chosen. Our proof adapts a compression argument of [[CK18](#)] which was used to analyze the security of

the regular discrete logarithm problem. In particular, if the probability that our preprocessing attacker wins is $\omega(SNq^2 \log p/p)$, then we could derive a contradiction by compressing our random injective map τ , mapping group elements to binary strings.

An interesting corollary of these results is that the 1-out-of- N discrete-log problem is hard even for a preprocessing attacker. Intuitively, if a discrete-log attacker can successfully compute x_i for any $i \leq N$, then s/he can also win the 1-out-of- N bridge-finding game.

Restricted Discrete-Log Oracle. In fact, we consider a stronger attacker \mathcal{A} who may query the usual generic group oracle, and is additionally given access to a restricted discrete-log oracle DLog . The oracle DLog will solve the discrete-log problem but only for “fresh” inputs, i.e., given N inputs g^{x_1}, \dots, g^{x_N} , if $h = g^{a_1x_1 + \dots + a_Nx_N}$ for known values of a_1, \dots, a_N , then this input would not be considered fresh.

We remark that by restricting the discrete-log oracle to fresh queries, we can rule out the trivial attack where the attacker simply queries $\text{DLog}(g^{x_i})$ for any $1 \leq i \leq N$. The restriction also rules out other trivial attacks, where the attacker simply queries $\text{DLog}(g^{x_i+r})$ after computing g^{x_i+r} for some $i \leq N$ and some known value r .

Security Reduction. We then give a reduction showing that any attacker \mathcal{A}_{sig} that breaks multi-user security of short Schnorr signatures can be used to win the 1-out-of- N bridge-finding game. In the reduction, we interpret the bridge inputs g^{x_1}, \dots, g^{x_N} as public signing keys, and we simulate the attacker \mathcal{A}_{sig} . The reduction uses the restricted discrete log oracle to ensure that any group element that is submitted as an input to the random oracle has the form $g^{b+a_1x_1 + \dots + a_Nx_N}$ for values a_1, \dots, a_N , and b that are known to the bridge attacker. The reduction also makes use of a programmable random oracle to forge signatures whenever \mathcal{A}_{sig} queries the signing oracle for a particular user $i \leq N$.

One challenge with carrying out this reduction in the preprocessing setting is that we need to ensure that the hint does not allow the attacker to detect when the random oracle has been programmed. We rely on the observation that the reduction programs the random oracle at random inputs which are unknown to the preprocessing a priori.

A similar reduction allows us to establish multi-user security of other Fiat-Shamir-based signatures such as Chaum-Pedersen signatures [CP93] and Katz-Wang signatures [KW03], with and without preprocessing.

1.3 Related Work

Security Proofs in the Generic Group Model. The generic group model goes back to Nechaev [Nec94] and Shoup [Sho97]. One motivation for analyzing cryptographic protocols in the generic group model is that for certain elliptic curve groups, the best known attacks are all generic [JMV01, FST10, WZ11, BL12, GWZ15].

It is well known that in Shoup’s generic group model [Sho97], an attacker requires $\Omega(\sqrt{p})$ queries to solve the discrete-log problem in a group of prime order p and the same lower bound holds for other classical problems like Computational Diffie-Hellman (CDH) and Decisional Diffie-Hellman (DDH). This bound is tight as discrete-log algorithms such as the Baby-Step Giant-Step algorithm by Shanks [Sha71], Pollard’s Rho and Kangaroo algorithms [Pol78], and the Pohlig-Hellman algorithm [PH06] can all be described generically in Shoup’s model. However, there are some exceptions for other elliptic curves and subgroups of \mathbb{Z}_p^* , where the best discrete-log algorithms are not generic and are much more efficient than any generic discrete log algorithms, e.g., see [GHS02,MVO91,Sma99].

Dent [Den02] showed that there are protocols which are provably secure in the generic group model but which are trivially insecure when the generic group is replaced with any (efficiently computable) real one. However, these results were artificially crafted to provide a counterexample. Similar to the random oracle model, experience suggests that protocols with security proofs in the generic group model do not have inherent structural weaknesses, and will be secure as long as we instantiate with a reasonable elliptic curve group. See [KM07,Fis00,JS08] for additional discussion of the strengths/weaknesses of proofs in the generic group model.

Corrigan-Gibbs and Kogan [CK18] analyzed the security of several key cryptographic problems (e.g., discrete-log, computational/decisional Diffie-Hellman, etc.) against preprocessing attacks in the generic group model. We extend their analysis to analyze the multi-user security of key-prefixed short Schnorr signatures against preprocessing attacks. See Section 5 for the further details.

Schnorr Signatures and Multi-Signatures. Bellare and Dai [BD20] recently showed that the (single-user) security of Schnorr signatures could be based on the Multi-Base Discrete Logarithm problem which in turn is similar in flavor to the One More Discrete Log Problem [BNPS03]. There has also been an active line of work on adapting the Schnorr signature scheme to design compact multi-signature schemes, e.g., see [BN06,BCJ08,DEF⁺19,MPSW19]. The goal is for multiple parties to collaborate to generate a single Schnorr signature which is signed using an aggregate public key that can be (publicly) derived from the individually public keys of each signing party. A very recent line of work has reduced the interaction to generate a Schnorr multi-signature to two-rounds without pairings [NRSW20,NRS20,AB20].

Other Short Signatures. Boneh et al. [BLS04] proposed even shorter signatures called BLS signatures, which is as short as $2k$ bits to yield k bits of security in the random oracle model, assuming that the Computational Diffie-Hellman (CDH) problem is hard on certain elliptic curves over a finite field. While BLS signatures yield even shorter signature length than Schnorr signatures, the computation costs for the BLS verification algorithm is several orders of magnitude higher, due to the reliance on bilinear pairings. If we allow for “heavy” cryptographic solutions such as indistinguishability obfuscation [GGH⁺13] (practically

infeasible at the moment), then it becomes possible to achieve k -bit signatures with k bits of security [SW14,RW14,LM17].

2 Preliminaries

Let \mathbb{N} be the set of positive integers, and we define $[N] := \{1, \dots, N\}$ for $N \in \mathbb{N}$. Throughout the paper, we denote the security parameter by k . We say that $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$ is an N -dimensional vector over \mathbb{Z}_p^N , and for each $i \in [N]$, we define \hat{u}_i to be the i^{th} N -dimensional unit vector, i.e., the i^{th} element of \hat{u}_i is 1, and all other elements are 0 elsewhere. For simplicity, we let $\log(\cdot)$ be a log with base 2, i.e., $\log x := \log_2 x$. The notation \leftarrow_s denotes a uniformly random sampling, e.g., we say $x \leftarrow_s \mathbb{Z}_p$ when x is sampled uniformly at random from \mathbb{Z}_p .

2.1 The Generic Group Model

The generic group model is an idealized cryptographic model proposed by Shoup [Sho97]. Let $G = \langle g \rangle$ be a multiplicative cyclic group of prime order p . In the generic group model, since G is isomorphic to \mathbb{Z}_p , we select a random injective map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$, where \mathbb{G} is the set of bit strings of length ℓ (with $2^\ell \geq p$) and we encode the discrete log of a group element instead of the group element itself.

The key idea in the generic group model is that the map τ does not need to be a group homomorphism, because any adversary against a cryptographic scheme is only available to see a randomly chosen encoding of the discrete log of each group element, not the group element itself. Hence, the generic group model assumes that an adversary has no access to the concrete representation of the group elements. Instead, the adversary is given access to an oracle parametrized by τ , which computes the group operation indirectly in G as well as the encoding of the discrete log of the generator g given as $\mathbf{g} = \tau(1)$. More precisely, for an input $(\mathbf{a}, \mathbf{b}) \in \mathbb{G} \times \mathbb{G}$, the oracles $\text{Mult}(\mathbf{a}, \mathbf{b})$, and $\text{Inv}(\mathbf{a})$ act as following:

$$\begin{aligned} \text{Mult}(\mathbf{a}, \mathbf{b}) &= \tau(\tau^{-1}(\mathbf{a}) + \tau^{-1}(\mathbf{b})), \text{ and} \\ \text{Inv}(\mathbf{a}) &= \tau(-\tau^{-1}(\mathbf{a})), \end{aligned}$$

if $\tau^{-1}(\mathbf{a}), \tau^{-1}(\mathbf{b}) \in G$. We remark that the adversary has no access to the map τ itself and does not know what is going on in a group G . Hence, the adversary has no sense that which element in G maps to $\text{Mult}(\mathbf{a}, \mathbf{b})$ in \mathbb{G} even if s/he sees the oracle output.

For convenience, we will use the notation $\text{Pow}(\mathbf{a}, n) = \tau(n\tau^{-1}(\mathbf{a}))$. Without loss of generality, we do not allow the attacker to directly query Pow as an oracle, since the attacker can efficiently evaluate this subroutine using the Mult oracle. In particular, one can evaluate $\text{Pow}(\mathbf{a}, n)$ using just $\mathcal{O}(\log n)$ calls to Mult using the standard modular exponentiation algorithm.

2.2 The Schnorr Signature Scheme

The Schnorr signature scheme is a digital signature scheme, which consists of a tuple of probabilistic polynomial-time algorithms $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$, where $\text{Kg}(1^k)$ is a key-generation algorithm to generate a secret key $\text{sk} \in \mathbb{Z}_p$ and a public key $\text{pk} = \text{Pow}(\mathbf{g} = \tau(1), \text{sk}) = \tau(\text{sk})$. The size of the prime number p will be tied to the security parameter k , e.g., $p \approx 2^{2k}$. $\text{Sign}(\text{sk}, m)$ is a signing algorithm which generates a signature σ on a message $m \in \{0, 1\}^*$, and $\text{Vfy}(\text{pk}, m, \sigma)$ is a verification algorithm which outputs 1 if the signature is valid, and 0 otherwise.

Throughout the paper, we will consider the notion of the generic group model in the Schnorr signature scheme as described in Figure 1. We remark that verification works for a correct signature $\sigma = (s, e)$, because $R = \text{Mult}(\tau(s), \text{Pow}(\text{Inv}(\text{pk}), e)) = \tau(s - \text{sk} \cdot e) = \tau(r) = I$ if the signature is valid.

$\text{Kg}(1^k)$:	$\text{Sign}(\text{sk}, m)$:	$\text{Vfy}(\text{pk}, m, \sigma)$:
1: $\text{sk} \leftarrow_{\$} \mathbb{Z}_p$	1: $r \leftarrow_{\$} \mathbb{Z}_p$	1: Parse $\sigma = (s, e)$
2: $\text{pk} \leftarrow \text{Pow}(\mathbf{g}, \text{sk})$	2: $I \leftarrow \text{Pow}(\mathbf{g}, r)$	2: $R \leftarrow \text{Mult}(\text{Pow}(\mathbf{g}, s), \text{Pow}(\text{Inv}(\text{pk}), e))$
3: return (pk, sk)	3: $e \leftarrow \text{H}(I m)$	3: if $\text{H}(R m) = e$ then
	4: $s \leftarrow r + \text{sk} \cdot e \pmod p$	4: return 1
	5: return $\sigma = (s, e)$	5: else return 0

Figure 1. The Schnorr signature scheme in the generic group model. Note that instead of the direct group operation in G , we use the encoding by the map τ and the generic group oracle queries.

Short Schnorr Signatures. Typically, it is assumed that the random oracle $\text{H}(I||m)$ outputs a uniformly random element $e \in \mathbb{Z}_p$, where p is a random $2k$ -bit prime. Thus, we would need $2k$ bits to encode e . To produce a shorter signature, we can assume that $\text{H}(I||m)$ outputs a uniformly random integer $e \in \mathbb{Z}_{2^k}$ with just k bits. In practice, the shorter random oracle is *easier* to implement, since we do not need to worry about rounding issues when converting a binary string to \mathbb{Z}_{2^k} , i.e., we can simply take the first k bits of our random binary string. The result is a signature $\sigma = (s, e)$, which can be encoded in $3k$ bits – $2k$ bits to encode $s \in \mathbb{Z}_p$ plus k bits to encode e . This natural modification is straightforward and is not new to our paper. The key question we investigate is whether or not short Schnorr signatures can provide k bits of security.

3 Single-User Security of Short Schnorr Signatures

As a warm-up to our main result, we first prove that *short* Schnorr Signatures (of length $3k$) achieve k bits of security. We first describe the standard signature forgery experiment $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau}(k)$ in the generic group model and the random

oracle model. Here, an attacker is given the public key $\mathbf{pk} = \tau(\mathbf{sk})$ along with $\mathbf{g} = \tau(1)$ (the encoding of the group generator 1 of \mathbb{Z}_p). The attacker is given oracle access to the signing oracle $\text{Sign}(\cdot)$, as well as the generic group oracles $\mathbf{GO} = (\text{Mult}(\cdot, \cdot), \text{Inv}(\cdot))$, and the random oracle $\mathbf{H}(\cdot)$. The attacker’s goal is to eventually output a forgery $(m, \sigma = (s, e))$ for a *fresh* message m that has not previously been submitted to the signing oracle.

Generic Signature Forgery Game. Fixing an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$, $\mathbf{g} = \tau(1)$ and the random oracle \mathbf{H} , and an adversary \mathcal{A} , consider the following experiment defined for a signature scheme $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$:

The Generic Signature Forgery Game $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, \mathbf{H}}(k)$:

- (1) $\text{Kg}(1^k)$ is run to obtain the public and the secret keys $(\mathbf{pk}, \mathbf{sk})$. Here, \mathbf{sk} is chosen randomly from the group \mathbb{Z}_p where p is a $2k$ -bit prime, and $\mathbf{pk} = \text{Pow}(\mathbf{g}, \mathbf{sk}) = \tau(\mathbf{sk})$.
- (2) Adversary \mathcal{A} is given $(\mathbf{g} = \tau(1), \mathbf{pk}, p)$ and access to the generic group oracles $\mathbf{GO} = (\text{Mult}(\cdot, \cdot), \text{Inv}(\cdot))$, the random oracle $\mathbf{H}(\cdot)$, and the signing oracle $\text{Sign}(\cdot)$. After multiple access to these oracles, the adversary outputs $(m, \sigma = (s, e))$.
- (3) We define $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, \mathbf{H}}(k) = \text{Vfy}(\mathbf{pk}, m, \sigma)$, i.e., the output is 1 when \mathcal{A} succeeds, and 0 otherwise.

Definition 1 formalizes this argument in the sense that an attacker forges a signature if and only if $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, \mathbf{H}}(k) = 1$.

Definition 1. Consider the generic group model with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$. A signature scheme $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ is said to be $(q_{\mathbf{H}}, q_{\mathbf{G}}, q_{\mathbf{S}}, \varepsilon)$ -UF-CMA secure (unforgeable against chosen message attack) if for every adversary \mathcal{A} making at most $q_{\mathbf{H}}$ (resp. $q_{\mathbf{G}}$, $q_{\mathbf{S}}$) queries to the random oracle (resp. generic group, signing oracles), the following bound holds:

$$\Pr \left[\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, \mathbf{H}}(k) = 1 \right] \leq \varepsilon,$$

where the randomness is taken over the selection of τ , the random coins of \mathcal{A} , the random coins of Kg , and the selection of random oracle \mathbf{H} .

3.1 Discrete Log Problem with Restricted Discrete Log Oracle

Restricted Discrete-Log Oracle in the Generic Group Model. In the discrete log problem we pick a random $x \in \mathbb{Z}_p$ and the attacker is challenged to recover x given $\mathbf{g} = \tau(1)$ and $\mathbf{h} = \text{Pow}(\mathbf{g}, x)$ after making queries to the generic group oracles Mult and Inv . As we mentioned in [Section 1.2](#), we analyze the discrete log problem in a stronger setting where the attacker is additionally given access to a restricted discrete-log oracle DLog . Given the map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ and $y \in \mathbb{Z}_p$, $\text{DLog}(\tau(y))$ will output y as long as $\tau(y)$ is a “fresh” group element. More

specifically, we say $\tau(y)$ is “fresh” if (1) $\tau(y)$ is not equal to \mathfrak{h} , and (2) $\tau(y)$ has not been the output of a previous generic group query.

The requirement that $\tau(y)$ is fresh rules out trivial attacks where the attacker picks $a, b \in \mathbb{Z}_p$, computes $\tau(ax + b) = \text{Mult}(\text{Pow}(\mathfrak{h}, a), \text{Pow}(\mathfrak{g}, b))$ and queries $\text{DLog}(\tau(ax + b))$ and solves for $x = a^{-1}(\text{DLog}(\tau(ax + b)) - b) \pmod p$.

The Generic Discrete-Log Game. The formal definition of the discrete log experiment $\text{DLogChal}_{\mathcal{A}}^{\tau}(k)$ is given below:

The Generic Discrete-Log Game $\text{DLogChal}_{\mathcal{A}}^{\tau}(k)$:

- (1) The adversary \mathcal{A} is given $(\mathfrak{g} = \tau(1), \tau(x))$ for a random value of $x \in \mathbb{Z}_p$. Here, $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ is a map from \mathbb{Z}_p to a generic group \mathbb{G} with a $2k$ -bit prime p .
- (2) \mathcal{A} is allowed to query the usual generic group oracles (Mult, Inv) and is additionally allowed to query $\text{DLog}(\tau(y))$, but *only* if $\tau(y)$ is “fresh”, i.e., $\tau(y)$ is not $\tau(x)$, and $\tau(y)$ has not been the output of a previous random generic group query.
- (3) After multiple queries, \mathcal{A} outputs x' .
- (4) The output of the game is defined to be $\text{DLogChal}_{\mathcal{A}}^{\tau}(k) = 1$ if $x' = x$, and 0 otherwise.

Lemma 1 upper bounds the probability that an attacker wins the generic discrete-log game $\text{DLogChal}_{\mathcal{A}}^{\tau}(k)$. Intuitively, the proof works by maintaining a list \mathcal{L} of tuples $(\tau(y), a, b)$ such that $y = ax + b$ for every oracle output $\tau(y)$.

Initially, the list \mathcal{L} contains two items $(\tau(x), 1, 0)$ and $(\tau(1), 0, 1)$, and the list is updated after every query to the generic group oracles, e.g., if $(\tau(y_1), a_1, b_1) \in \mathcal{L}$ and $(\tau(y_2), a_2, b_2) \in \mathcal{L}$, then querying $\text{Mult}(\tau(y_1), \tau(y_2))$ will result in the addition of $(\tau(y_1 + y_2), a_1 + a_2, b_1 + b_2)$ into \mathcal{L} . If \mathcal{L} already contained a tuple of the form $(\tau(y_1 + y_2), a', b')$ with $a' \neq a_1 + a_2$ or $b' \neq b_1 + b_2$, then we say that the event **BRIDGE** occurs.

We can use the restricted discrete log oracle to maintain the invariant that every output of our generic group oracles Mult and Inv can be added to \mathcal{L} . In particular, if we every encounter an input $\eta = \tau(b)$ that does not already appear in \mathcal{L} , then η is fresh and we can simply query the restricted discrete log oracle to extract $b = \text{DLog}(\eta)$, ensuring that the tuple $(\eta, 0, b)$ is added to \mathcal{L} before the generic group query is processed.

The key component of the proof is to upper bound the probability of the event **BRIDGE**. This is sufficient as any attacker that can recover x will also be able to ensure that $(\tau(x), 0, x)$ is added to \mathcal{L} , which would immediately cause the event **BRIDGE** to occur, since we already have $(\tau(x), 1, 0) \in \mathcal{L}$. We defer the full proof of **Lemma 1** to [Appendix C.1](#) for readers who are interested.

Lemma 1. *The probability the attacker making at most $q_{\mathbb{G}}$ generic group oracle queries wins the generic discrete-log game $\text{DLogChal}_{\mathcal{A}}^{\tau}(k)$ (even with access to*

the restricted DLog oracle) is at most

$$\Pr[\text{DLogChal}_{\mathcal{A}}^{\tau}(k) = 1] \leq \frac{6q_{\mathbb{G}}(q_{\mathbb{G}} + 1) + 12}{4p - (3q_{\mathbb{G}} + 2)^2},$$

in the generic group model of prime order p , where the randomness is taken over the selection of τ , the challenge x , as well as any random coins of \mathcal{A} .

3.2 Security Reduction

Given [Lemma 1](#), we are now ready to describe our security reduction for short Schnorr signatures of length $3k$. As in our security proof for the discrete-log problem, we will ensure that for every output $\tau(y)$ of a generic group query, we can express $y = ax + b$ for known constants a and b – here x is the secret key that is selected in the security game, i.e., any time \mathcal{A}_{sig} makes a query involving a fresh element $\tau(y)$, we will simply query $\text{DLog}(\tau(y))$ so that we can add $\tau(y)$ to the list \mathcal{L} .

[Theorem 4](#) provides the first rigorous proof of the folklore claim that short ($3k$ -bit) Schnorr signatures can provide k bits of security. The formal security proof uses *both* the generic group model and the random oracle model.

Theorem 4. *The short Schnorr signature scheme $\Pi_{\text{short}} = (\text{Kg}, \text{Sign}, \text{Vfy})$ of length $3k$ is $(q_{\text{H}}, q_{\mathbb{G}}, q_{\text{S}}, \varepsilon)$ -UF-CMA secure with*

$$\varepsilon = \frac{6q_{\mathbb{G}}(q_{\mathbb{G}} + 1) + 12}{4p - (3q_{\mathbb{G}} + 2)^2} + \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p} + \frac{q_{\text{H}} + q_{\text{S}}}{p - (3q_{\mathbb{G}} + 2)} + \frac{q_{\text{H}} + 1}{2^k} = \mathcal{O}\left(\frac{q}{2^k}\right),$$

in the generic group model of prime order $p \approx 2^{2k}$ and the programmable random oracle model, where q denotes the total number of queries made by an adversary.

Proof Sketch of [Theorem 4](#): Here, we only give the intuition of how the proof works. The full proof of [Theorem 4](#) is similar to that of [Theorem 6](#) and we defer it to [Appendix C.1](#).

We give the proof by reduction, i.e., given an adversary \mathcal{A}_{sig} attacking short Schnorr signature scheme, we construct an efficient algorithm $\mathcal{A}_{\text{dlog}}$ which solves the discrete-log problem. During the reduction, we simulate the signature signing process without secret key x by programming the random oracle, i.e., to sign a message m we can pick s and e randomly, compute $I = \tau(s - xe)$ by querying adequate generic group oracles³, and see if the random oracle has been previously queried at $\text{H}(I||m)$. If not, then we can program the random oracle as $\text{H}(I||m) := e$ and output the signature (s, e) . Otherwise, the reduction simply outputs \perp for failure. Since s and e are selected randomly, we can argue that the probability that we output \perp because $\text{H}(I||m)$ is already defined is small, i.e., $\approx q_{\text{H}}/p$.

We can use the oracle DLog to maintain the invariant that before processing any random oracle query of the form $\text{H}(\eta||\cdot)$ that we know a, b such that $\tau(ax +$

³ We can compute I without knowledge of x because $\tau(x)$ is given as public key.

$b) = \eta$. In particular, if η is a fresh string that has not previously been observed, then we can simply set $a = 0$ and query the restricted discrete log oracle to find b such that $\tau(b) = \eta$. We say that the random oracle query is lucky if $H(\eta\|m) = -a$, or $H(\eta\|m) = 0$. Assuming that the event BRIDGE does not occur, it is straightforward to upper bound the probability of a lucky query as $\mathcal{O}(q_H/2^k)$. Similarly, it is straightforward to show that the probability that \mathcal{A}_{sig} gets lucky and guesses a valid signature (s, e) for a message m without first querying $H(\tau(s - xe)\|m)$ is $\mathcal{O}(2^{-k})$. Assuming that there are no lucky queries or guesses but the attacker still outputs a successful signature forgery (s, e) . In this case we have $I = \tau(s - xe) = \tau(ax + b)$, which allows for us to solve for x using the equation $(a + e)x = (s - b)$. Thus, we can argue that the probability $\mathcal{A}_{\text{dlog}}$ solves the discrete-log challenge correctly is lower bounded by the probability that \mathcal{A}_{sig} forges a signature minus $\mathcal{O}(q/2^k)$. Finally, by applying Lemma 1 we can upper bound the probability \mathcal{A}_{sig} wins the generic signature forgery game $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, H}(k)$ to be $\mathcal{O}(q/2^k)$. \square

4 Multi-User Security of Short Schnorr Signatures

In this section, we prove that short Schnorr signatures also provide k bits of security in the multi-user setting. The reduction uses similar ideas, but requires us to introduce and analyze a game called the *1-out-of- N generic BRIDGE ^{N} -finding game*. We first define the *1-out-of- N generic signature forgery game*, where an adversary is given N independent public keys $(\text{pk}_1, \dots, \text{pk}_N) = (\tau(\text{sk}_1), \dots, \tau(\text{sk}_N))$ along with oracle access to the signing oracles $\text{Sign}(\text{sk}_1, \cdot), \dots, \text{Sign}(\text{sk}_N, \cdot)$, the random oracle H , and the generic group oracles. The attacker can succeed if s/he can output a forgery (σ, m) which is valid under *any one* public key, e.g., for some public key pk_j we have $\text{Vfy}(\text{pk}_j, m, \sigma) = 1$, while the query m was never submitted to the j th signing oracle $\text{Sign}(\text{sk}_j, \cdot)$. In our reduction, we show that any attacker that wins the 1-out-of- N generic signature forgery game can be used to win the 1-out-of- N generic BRIDGE ^{N} -finding game. We separately upper bound the probability that a generic attacker can win the 1-out-of- N generic BRIDGE ^{N} -finding game.

1-out-of- N Generic Signature Forgery Game. Fixing the injective mapping $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$, a random oracle H , and an adversary \mathcal{A} , consider the following experiment defined for a signature scheme $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$:

The 1-out-of- N Generic Signature Forgery Game $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, H, N}(k)$:

- (1) $\text{Kg}(1^k)$ is run N times to obtain the public and the secret keys $(\text{pk}_i, \text{sk}_i)$ for each $i \in [N]$. Here, for each $i \in [N]$, sk_i is chosen randomly from the group \mathbb{Z}_p , where p is a $2k$ -bit prime, and $\text{pk}_i = \tau(\text{sk}_i)$.
- (2) Adversary \mathcal{A} is given $(\mathbf{g} = \tau(1), \text{pk}_1, \dots, \text{pk}_N, p)$, and access to the generic group oracles $\mathbf{GO} = (\text{Mult}(\cdot, \cdot), \text{Inv}(\cdot))$, the random oracle $H(\cdot)$,

and the signing oracles $\text{Sign}(\text{sk}_1, \cdot), \dots, \text{Sign}(\text{sk}_N, \cdot)$. The experiment ends when the adversary outputs $(m, \sigma = (s, e))$.

- (3) \mathcal{A} succeeds to forge a signature if and only if there exists some $j \in [N]$ such that $\text{Vfy}(\text{pk}_j, m, \sigma) = 1$ and the query m was never submitted to the oracle $\text{Sign}(\text{sk}_j, \cdot)$. The output of the experiment is $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, \text{H}, N}(k) = 1$ when \mathcal{A} succeeds; otherwise $\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, \text{H}, N}(k) = 0$.

Definition 2. Consider the generic group model with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$. A signature scheme $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ is $(N, q_{\text{H}}, q_{\text{G}}, q_{\text{S}}, \varepsilon)$ -MU-UF-CMA secure (multi-user unforgeable against chosen message attack) if for every adversary \mathcal{A} making at most q_{H} (resp. $q_{\text{G}}, q_{\text{S}}$) queries to the random oracle (resp. generic group, signing oracles), the following bound holds:

$$\Pr \left[\text{SigForge}_{\mathcal{A}, \Pi}^{\tau, \text{H}, N}(k) = 1 \right] \leq \varepsilon,$$

where the randomness is taken over the selection of τ , the random coins of \mathcal{A} , the random coins of Kg , and the selection of random oracle H .

The Discrete-Log Solution List \mathcal{L} in a Multi-User Setting. As before, we will maintain the invariant that for every output η of a generic group query that we have recorded a tuple (η, \vec{a}, b) in a list \mathcal{L} where $\text{DLog}(\eta) = \vec{a} \cdot \vec{x} + b$ (here, $\vec{a} = (a_1, \dots, a_N)$, $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$). Note that the restricted oracle $\text{DLog}(\cdot)$ will solve $\text{DLog}(\eta)$ for any fresh group element η such that $\eta \notin \{\tau(1), \tau(x_1), \dots, \tau(x_N)\}$, and η has not been the output of a prior generic group query.

- Initially, \mathcal{L} contains $(\tau(1), \vec{0}, 1)$ and $(\tau(x_i), \hat{u}_i, 0)$ for $1 \leq i \leq N$.
- If the attacker ever submits a fresh group element η which was not previously an output of a generic group oracle query, then we can query $b = \text{DLog}(\eta)$, and add $(\eta, \vec{0}, b)$ to our list. Thus, without loss of generality, we can assume that all query inputs to Mult , Inv were first added to \mathcal{L} .
- If $(\eta_1, \vec{a}_1, b_1), (\eta_2, \vec{a}_2, b_2) \in \mathcal{L}$, and the attacker queries $\text{Mult}(\eta_1, \eta_2)$, then add $(\text{Mult}(\eta_1, \eta_2), \vec{a}_1 + \vec{a}_2, b_1 + b_2)$ to \mathcal{L} .
- If $(\eta, \vec{a}, b) \in \mathcal{L}$, and $\text{Inv}(\eta)$ is queried, then add $(\text{Inv}(\eta), -\vec{a}, -b)$ to \mathcal{L} .

4.1 The Multi-User Bridge-Finding Game

We establish the multi-user security of short Schnorr signatures via reduction from a new game we introduce called the *1-out-of- N generic BRIDGE^N-finding game*. As in the 1-out-of- N discrete-log game, the attacker is given $\tau(1)$, as well as $\tau(x_1), \dots, \tau(x_N)$ for N randomly selected values $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$. The key difference between this game and the 1-out-of- N discrete-log game is that the attacker’s goal is simply to ensure that the “bridge event” BRIDGE^N occurs, whether or not the attacker is able to solve any of the discrete-log challenges. As in Section 3, we will assume that we have access to $\text{DLog}(\cdot)$, and we will maintain the invariant that for every output $\tau(y)$ of some generic group query, we have $y = \vec{a} \cdot \vec{x} + b$ for known values $\vec{a} = (a_1, \dots, a_N) \in \mathbb{Z}_p^N$ and $b \in \mathbb{Z}_p$, i.e., by querying the restricted oracle $\text{DLog}(\tau(y))$ whenever we encounter a fresh input.

The 1-out-of- N Generic BRIDGE ^{N} -Finding Game $\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k, \vec{x})$:

- (1) The challenger initializes the list $\mathcal{L} = \{(\tau(1), \vec{0}, 1), (\tau(x_1), \hat{u}_1, 0), \dots, (\tau(x_N), \hat{u}_N, 0)\}$, and $\vec{x} = (x_1, \dots, x_N)$.
- (2) The adversary \mathcal{A} is given $\mathbf{g} = \tau(1)$ and $\tau(x_i)$ for each $i \in [N]$.
- (3) \mathcal{A} is allowed to query the usual generic group oracles (**Mult**, **Inv**).
 - (a) If the challenger ever submits any fresh element η which does not appear in \mathcal{L} as input to a generic group oracle, then the challenger immediately queries $b_y = \text{DLog}(\eta)$, and adds the tuple $(\eta, \vec{0}, b_y)$ to the list \mathcal{L} .
 - (b) Whenever \mathcal{A} submits a query η_1, η_2 to **Mult**(\cdot, \cdot), we are ensured that there exist tuples $(\eta_1, \vec{a}_1, b_1), (\eta_2, \vec{a}_2, b_2) \in \mathcal{L}$. The challenger adds the tuple $(\text{Mult}(\eta_1, \eta_2), \vec{a}_1 + \vec{a}_2, b_1 + b_2)$ to the list \mathcal{L} .
 - (c) Whenever \mathcal{A} submits a query η to **Inv**(\cdot), we are ensured that some tuple $(\eta, \vec{a}_y, b_y) \in \mathcal{L}$. The challenger adds the tuple $(\text{Inv}(\eta), -\vec{a}_y, -b_y)$ to the list \mathcal{L} .
- (4) If at any point in time we have a collision, i.e., two distinct tuples $(\eta, \vec{a}_1, b_1), (\eta, \vec{a}_2, b_2) \in \mathcal{L}$ with $(\vec{a}_1, b_1) \neq (\vec{a}_2, b_2)$, then the event **BRIDGE ^{N}** occurs, and the output of the game is 1. If **BRIDGE ^{N}** never occurs, then the output of the game is 0.

We further define the game $\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k)$ in which $\vec{x} = (x_1, \dots, x_N)$ are first sampled uniformly at random, and then we run $\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k, \vec{x})$. Thus, we have

$$\Pr_{\mathcal{A}, \tau}[\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k) = 1] := \Pr_{\mathcal{A}, \tau, \vec{x}}[\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k, \vec{x}) = 1].$$

As long as the event **BRIDGE ^{N}** has not occurred, we can (essentially) view x_1, \dots, x_N as uniformly random values that are yet to be selected. More precisely, the values x_1, \dots, x_N are selected subject to a few constraints, e.g., if we know $f_1 = \tau(\vec{a}_1 \cdot \vec{x} + b_1) \neq f_2 = \tau(\vec{a}_2 \cdot \vec{x} + b_2)$ then we have the constraint that $\vec{a}_1 \cdot \vec{x} + b_1 \neq \vec{a}_2 \cdot \vec{x} + b_2$.

Theorem 5. *For any attackers \mathcal{A} making at most $q_{\mathbf{G}} := q_{\mathbf{G}}(k)$ queries to the generic group oracles,*

$$\Pr[\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k) = 1] \leq \frac{q_{\mathbf{G}}N + 3q_{\mathbf{G}}(q_{\mathbf{G}} + 1)/2}{p - (N + 3q_{\mathbf{G}} + 1)^2 - N},$$

in the generic group model of prime order p where the randomness is taken over the selection of x_1, \dots, x_N, τ as well as any random coins of \mathcal{A} .

Proof. Consider the output η_i of the i^{th} generic group query. We first analyze the probability that this query results in the event **BRIDGE ^{N}** conditioning on the event $\overline{\text{BRIDGE}_{<i}^N}$ that the event has not yet occurred, i.e., the event **BRIDGE ^{N}** has not been occurred until the $(i-1)^{\text{th}}$ query. Before we even receive the output

η_i , we already know the values \vec{a}_i, b_i such that the tuple (η_i, \vec{a}_i, b_i) will be added to \mathcal{L} . If \mathcal{L} does already contain this *exact* tuple, then outputting η_i will not produce the event BRIDGE^N . If \mathcal{L} does not already contain this tuple (η_i, \vec{a}_i, b_i) , then we are interested in the event B_i that some other tuple $(\eta_i, \vec{a}'_i, b'_i)$ has been recorded with $(\vec{a}'_i, b'_i) \neq (\vec{a}_i, b_i)$. Observe that B_i occurs if and only if there exists a tuple of the form (\cdot, \vec{a}, b) with $(\vec{a} - \vec{a}_i) \cdot \vec{x} = b_i - b$ and $(\vec{a}, b) \neq (\vec{a}_i, b_i)$. If we pick \vec{x} randomly, the probability that $(\vec{a} - \vec{a}_i) \cdot \vec{x} = b_i - b$ would be $1/p$. However, we cannot quite view \vec{x} as random due to the restrictions, i.e., because we condition on the event $\overline{\text{BRIDGE}}_{<i}^N$, we know that for any distinct pair (η_i, \vec{a}_i, b_i) and (η_j, \vec{a}_j, b_j) we know that $\vec{a}_i \cdot \vec{x} + b_i \neq \vec{a}_j \cdot \vec{x} + b_j$.

Consider sampling \vec{x} uniformly at random subject to this restriction. Let $r \leq N$ be an index such that $\vec{a}[r] - \vec{a}_i[r] \neq 0$ and suppose that $x_r = \vec{x}[r]$ is the last value sampled. At this point, we can view x_r as being drawn uniformly at random from a set of at least $p - |\mathcal{L}|^2 - (N - 1)$ remaining values, subject to all of the restrictions. We also observe that $|\mathcal{L}| \leq N + 3q_G + 1$ since each generic group oracle query adds *at most* three new tuples to \mathcal{L} — exactly three in the case that we query $\text{Mult}(\eta_1, \eta_2)$ on two fresh elements. Thus, the probability that $(\vec{a} - \vec{a}_i) \cdot \vec{x} = b_i - b$ is at most $\frac{1}{p - (N + 3q_G + 1)^2 - (N - 1)}$. Union bounding over all tuples $(\cdot, \vec{a}, b) \in \mathcal{L}$, we have

$$\Pr \left[B_i : \overline{\text{BRIDGE}}_{<i}^N \right] \leq \frac{N + 3i}{p - (N + 3q_G + 1)^2 - N}.$$

To complete the proof, we observe that

$$\begin{aligned} \Pr \left[\text{BridgeChal}_{\mathcal{A}}^{\text{GO}, N}(k) = 1 \right] &= \sum_{i \leq q_G} \Pr \left[B_i : \overline{\text{BRIDGE}}_{<i}^N \right] \\ &\leq \sum_{i \leq q_G} \frac{N + 3i}{p - (N + 3q_G + 1)^2 - N} = \frac{q_G N + 3q_G(q_G + 1)/2}{p - (N + 3q_G + 1)^2 - N}. \quad \square \end{aligned}$$

As an immediate corollary of [Theorem 5](#), we can show that an attacker wins the 1-out-of- N discrete log game with (approximately) the same probability as in the multi-user bridge-finding game. In particular, given any attacker \mathcal{A}' in the 1-out-of- N discrete-log game $\text{1ofNDLog}_{\mathcal{A}}^{\tau, N}(k)$, where the attacker's goal is to output *any* $x \in \{x_1, \dots, x_N\}$ given input $\tau(1), \tau(x_1), \dots, \tau(x_N)$, we can construct an attacker \mathcal{A} in the game $\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k)$. \mathcal{A} simply runs \mathcal{A}' to obtain an output x , and then computes $\tau(x)$ using at most $2 \log p$ queries to the $\text{Mult}(\cdot, \cdot)$ oracle. If $x \in \{x_1, \dots, x_N\}$, then the bridge event BRIDGE^N must have occurred at some point, since we have $(\tau(x), \vec{0}, x) \in \mathcal{L}$ and $(\tau(x), \hat{u}_i, 0) \in \mathcal{L}$ for some $i \in [N]$.

Corollary 1. *For any attacker \mathcal{A} making at most $q_G + 2 \log p$ queries,*

$$\Pr \left[\text{1ofNDLog}_{\mathcal{A}}^{\tau, N}(k) = 1 \right] \leq \frac{q_G N + 3q_G(q_G + 1)/2}{p - (N + 3q_G + 1)^2 - N},$$

in the generic group model of prime order p , where the randomness is taken over the selection of τ , the challenges x_1, \dots, x_N , and any random coins of \mathcal{A} .

4.2 Security Reduction

Theorem 6. *The short Schnorr signature scheme $\Pi_{\text{short}} = (\text{Kg}, \text{Sign}, \text{Vfy})$ of length $3k$ is $\left(N, q_{\text{H}}, q_{\text{G}}, q_{\text{S}}, \varepsilon = \mathcal{O}\left(\frac{q_{\text{H}} + N}{2^k}\right)\right)$ -MU-UF-CMA secure with*

$$\varepsilon = \frac{q_{\text{G}}N + 3q_{\text{G}}(q_{\text{G}} + 1)/2}{p - (N + 3q_{\text{G}} + 1)^2 - N} + \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p} + \frac{q_{\text{H}} + q_{\text{S}}}{p - (N + 3q_{\text{G}} + 1)} + \frac{q_{\text{H}} + 1}{2^k},$$

in the generic group model of prime order $p \approx 2^{2k}$ and the programmable random oracle model, where q denotes the total number of queries made by an adversary.

Proof. Given an adversary \mathcal{A}_{sig} attacking short Schnorr signature scheme, we construct the following efficient algorithm $\mathcal{A}_{\text{bridge}}$ which tries to succeed in the 1-out-of- N generic BRIDGE ^{N} -finding game $\text{BridgeChal}_{\mathcal{A}_{\text{bridge}}}^{\tau, N}(k)$:

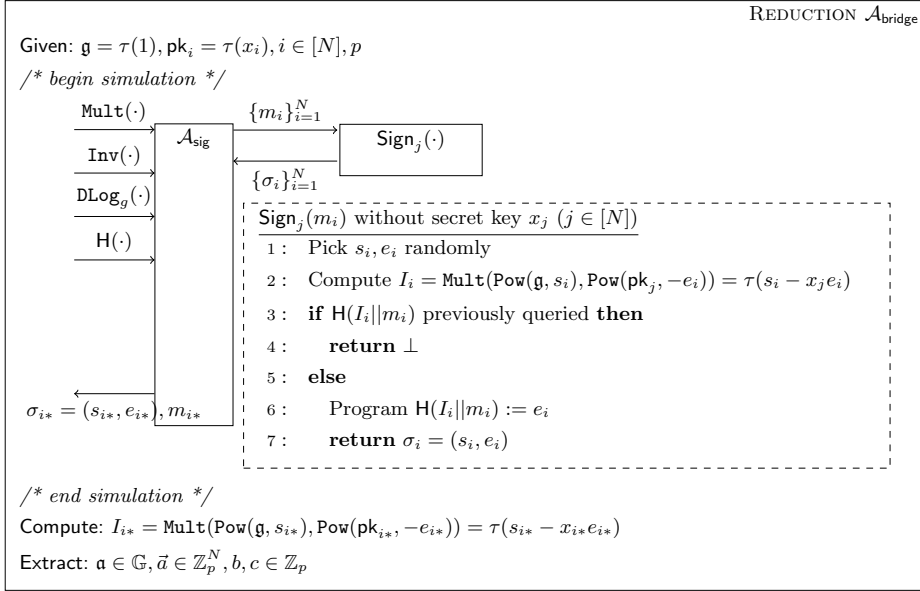


Figure 2. A reduction to the $\text{BridgeChal}_{\mathcal{A}_{\text{bridge}}}^{\tau, N}(k)$ attacker $\mathcal{A}_{\text{bridge}}$ from the short Schnorr signature attacker \mathcal{A}_{sig} .

Algorithm $\mathcal{A}_{\text{bridge}}$:

The algorithm is given $p, \mathbf{g} = \tau(1), \tau(x_i), 1 \leq i \leq N$ as input.

1. Initialize the list $\mathcal{L} = \{(\tau(1), \vec{0}, 1), (\tau(x_i), \hat{u}_i, 0) \text{ for each } i \in [N]\}$, and $\text{H}_{\text{resp}} = \{\}$, where H_{resp} stores the random oracle queries.
2. Run \mathcal{A}_{sig} with a number of access to the generic oracles $\text{GO} = (\text{Mult}(\cdot, \cdot), \text{Inv}(\cdot), \text{DLog}_g(\cdot), \text{Sign}_i(\cdot))$ for $1 \leq i \leq N$, and $\text{H}(\cdot)$. The signing oracle without a secret key is described in Figure 2. Now we consider the following cases:

- (a) Whenever \mathcal{A}_{sig} submits a query w to the random oracle H :
 - If there is a pair $(w, R) \in \mathsf{H}_{\text{resp}}$ for some string R , then return R .
 - Otherwise, select $R \leftarrow_{\mathfrak{s}} \mathbb{Z}_{2^k}$, and add (w, R) to the set H_{resp} .
 - If w has the form $w = (\mathbf{a} \| m_i)$, where the value \mathbf{a} has not been observed previously (i.e., is not in the list \mathcal{L}), then we query $b = \text{DLog}(\mathbf{a})$, and add $(\mathbf{a}, \vec{0}, b)$ to \mathcal{L} .
 - (b) Whenever \mathcal{A}_{sig} submits a query \mathbf{a} to the generic group oracle $\text{Inv}(\mathbf{a})$:
 - If \mathbf{a} is not in \mathcal{L} then we immediately query $b = \text{DLog}(\mathbf{a})$ and add $(\mathbf{a}, \vec{0}, b)$ to \mathcal{L} .
 - Otherwise, $(\mathbf{a}, \vec{a}, b) \in \mathcal{L}$. Then we query $\text{Inv}(\mathbf{a}) = \tau(-\vec{a} \cdot \vec{x} - b)$, output the result and add the result $(\tau(-\vec{a} \cdot \vec{x} - b), -\vec{a}, -b) \in \mathcal{L}$.
 - (c) Whenever \mathcal{A}_{sig} submits a query \mathbf{a}, \mathbf{b} to the generic group oracle $\text{Mult}(\mathbf{a}, \mathbf{b})$:
 - If the element \mathbf{a} (resp. \mathbf{b}) is not in \mathcal{L} , then query $b_0 = \text{DLog}(\mathbf{a})$ (resp. $b_1 = \text{DLog}(\mathbf{b})$), and add the element $(\mathbf{a}, \vec{0}, b_0)$ (resp. $(\mathbf{b}, \vec{0}, b_1)$) to \mathcal{L} .
 - Otherwise, both elements $(\mathbf{a}, \vec{a}_0, b_0), (\mathbf{b}, \vec{a}_1, b_1) \in \mathcal{L}$. Then we return $\text{Mult}(\mathbf{a}, \mathbf{b}) = \tau((\vec{a}_0 + \vec{a}_1) \cdot \vec{x} + b_0 + b_1)$, and add $(\tau((\vec{a}_0 + \vec{a}_1) \cdot \vec{x} + b_0 + b_1), \vec{a}_0 + \vec{a}_1, b_0 + b_1) \in \mathcal{L}$.
 - (d) Whenever \mathcal{A}_{sig} submits a query m_i to the signing oracle $\text{Sign}(x_j, \cdot)$:
 - We use the procedure Sign_j described in Figure 2 to forge a signature without knowledge of the secret key x_i . Intuitively, the forgery procedure relies on our ability to program the random oracle.
 - We remark that a side effect of querying the Sign_j oracle is the addition of the tuples $(\tau(s_i), \vec{0}, s_i), (\tau(x_j e_i), e_i \hat{u}_i, 0)$ and $(\tau(s_i - x_j e_i), -e_i \hat{u}_i, s_i)$ to \mathcal{L} , since these values are computed using the generic group oracles Inv and Mult .
 - (e) If at any point we find some string η such that $(\eta, \vec{a}, b) \in \mathcal{L}$ and $(\eta, \vec{c}, d) \in \mathcal{L}$ for $(\vec{a}, b) \neq (\vec{c}, d)$, then we can immediately have a BRIDGE^N instance $(\tau((\vec{a} - \vec{c}) \cdot \vec{x}), \vec{a} - \vec{c}, 0) \in \mathcal{L}$ and $(\tau(d - b), \vec{0}, d - b) \in \mathcal{L}$ since $\tau((\vec{a} - \vec{c}) \cdot \vec{x}) = \tau(d - b)$.⁴ Thus, without loss of generality, we can assume that each string η occurs at most once in the list \mathcal{L} .
3. After \mathcal{A}_{sig} outputs $\sigma_{i^*} = (s_{i^*}, e_{i^*})$ and m_{i^*} , identify the index $i^* \in [N]$ such that $\text{Vfy}(\text{pk}_{i^*}, m_{i^*}, \sigma_{i^*}) = 1$.
 4. Compute $\tau(-e_{i^*} x_{i^*}) = \text{Inv}(\text{Pow}(\tau(x_{i^*}), e_{i^*}))$ and $\mathfrak{s}_{i^*} = \text{Pow}(\mathfrak{g}, s_{i^*})$. This will ensure that the elements $(\tau(-e_{i^*} x_{i^*}), -e_{i^*} \hat{u}_{i^*}, 0)$ and $(\tau(e_{i^*} x_{i^*}), e_{i^*} \hat{u}_{i^*}, 0)$, and $(\mathfrak{s}_{i^*}, \vec{0}, s_{i^*})$ are all added to \mathcal{L} .
 5. Compute $I_{i^*} = \text{Mult}(\mathfrak{s}_{i^*}, \tau(-e_{i^*} x_{i^*})) = \tau(s_{i^*} - x_{i^*} e_{i^*})$ which ensures that $(I_{i^*}, -e_{i^*} \hat{u}_{i^*}, s_{i^*}) \in \mathcal{L}$. Finally, we can check to see if we previously had any tuple of the form $(I_{i^*}, \vec{a}, b) \in \mathcal{L}$.

Analysis. We first remark that if the signature is valid then we must have $e_{i^*} = \mathsf{H}(I_{i^*} \| m_{i^*})$ and $\text{DLog}(I_{i^*}) = s_{i^*} - x_{i^*} e_{i^*} = \vec{a} \cdot \vec{x} + b$.

We now define failure events $\text{FailtoFind}(I_{i^*})$ and BadQuery . $\text{FailtoFind}(I_{i^*})$ denotes the event that we find that the signature is valid, but I_{i^*} was not previously recorded in our list \mathcal{L} before we computed $\text{Mult}(\mathfrak{s}_{i^*}, \tau(-e_{i^*} x_{i^*}))$ in the

⁴ Note that $(\vec{a}, b) \neq (\vec{c}, d)$ implies $\vec{a} \neq \vec{c}$ since if $\vec{a} = \vec{c}$ then $\vec{a} \cdot \vec{x} + b = \vec{a} \cdot \vec{x} + d$ implies $b = d$ as $b, d \in \mathbb{Z}_p$.

last step. Similarly, let **BadQuery** denote the event that the signature is valid but for the only prior tuple $(I_{i^*}, \vec{a}, b) \in \mathcal{L}$ recorded in \mathcal{L} we have that $\vec{a} = -e_{i^*} \hat{u}_{i^*}$. If the signature is valid and neither of the events **FailtoFind** (I_{i^*}) and **BadQuery** occur, then the bridge event BRIDGE^N must have occurred and we immediately win the game since $(I_{i^*}, \vec{a}, b) \in \mathcal{L}$, $(I_{i^*}, -e_{i^*} \hat{u}_{i^*}, s_{i^*}) \in \mathcal{L}$ and $\vec{a} \neq -e_{i^*} \hat{u}_{i^*}$.

We additionally consider then the event **FailtoSign** where our reduction outputs \perp in Step 2.(d) due signing oracle failure i.e., because $H(I_i \| m_i)$ has been queried previously. Intuitively, the attacker will output a valid signature forgery with probability at least $\Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}}, \Pi_{\text{short}}}^{\tau, N}(k) = 1] - \Pr[\text{FailtoSign}]$ after we replace the signing oracle with the procedure **Sign_j** described in [Figure 2](#).

[Claim 1](#), [Claim 2](#), and [Claim 3](#) upper bound the probability of our events **FailtoSign**, **FailtoFind** and **BadQuery** respectively. We defer the proofs to [Appendix C.2](#).

$$\textbf{Claim 1. } \Pr[\text{FailtoSign}] \leq \frac{q_s(q_H + q_s)}{p}.$$

$$\textbf{Claim 2. } \Pr[\text{FailtoFind}(I_{i^*})] \leq \frac{q_H + q_s}{p - |\mathcal{L}|} + \frac{1}{2^k}.$$

$$\textbf{Claim 3. } \Pr[\text{BadQuery}] \leq \frac{q_H}{2^k}.$$

Since we have $|\mathcal{L}| \leq N + 3q_G + 1$, we can apply [Theorem 5](#) to conclude that

$$\begin{aligned} & \Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}}, \Pi_{\text{short}}}^{\tau, N}(k) = 1] \\ & \leq \Pr[\text{BridgeChal}_{\mathcal{A}_{\text{bridge}}}^{\tau, N}(k) = 1] + \Pr[\text{FailtoSign}] + \Pr[\text{FailtoFind}(I_{i^*})] + \Pr[\text{BadQuery}] \\ & \leq \frac{q_G N + 3q_G(q_G + 1)/2}{p - (N + 3q_G + 1)^2 - N} + \frac{q_s(q_H + q_s)}{p} + \frac{q_H + q_s}{p - (N + 3q_G + 1)} + \frac{q_H + 1}{2^k} \\ & = \mathcal{O}\left(\frac{q + N}{2^k}\right). \quad \square \end{aligned}$$

Remark 1. Kiltz et al. [\[KMP16\]](#) proved that multi-user security of *regular* Schnorr signatures is tightly equivalent to the Q -IDLOG problem in the random oracle model. Here, the Q -IDLOG problem is defined as follows: an adversary is given g^x and access to a challenge oracle CH , where $\text{CH}(g^{r_i})$ returns a uniformly random element $h_i \in \mathbb{Z}_p$. After up to Q queries to CH , the adversary wins if it returns a value s from the set $\{xh_i + r_i \mid i \in [Q]\}$. Kiltz et al. [\[KMP16, Theorem A.1\]](#) also proved that a generic attacker making q generic group queries and Q queries to CH wins with probability proportional to $(q^2 + Q)/p$.

In private communication, the authors of [\[KMP16\]](#) pointed out that [Theorem A.1](#) could be updated to give the bound $q^2/p + Q/2^k$ when CH returns a random element $h_i \in \{0, 1\}^k$ as in short Schnorr signatures. This would also imply tight multi-user security of short Schnorr signatures in the random oracle model and (their version of) the generic group model. However, short Schnorr signatures are not discussed in [\[KMP16\]](#). We prefer to work in Shoup's generic group model as it is more suitable for analyzing protocol composition and preprocessing attacks.

◁

5 Multi-User Security of Short Schnorr Signatures with Key-Prefixing against Preprocessing Attacks

In this section, we analyze the security of short Schnorr signatures against a preprocessing attacker who first outputs an S -bit hint after making (a very large number of) preprocessing queries to the generic group oracles `Mult` and `Inv`, as well as the random oracle `H`. After the public/secret keys are chosen, the signature forgery attacker will try use the hint to help win the signature forgery game. The hint must be fixed *before* the public/secret keys for our signature scheme are selected, otherwise the preprocessing attacker can generate forged signatures and embed them in the hint.

We first observe that Schnorr signatures are trivially broken against a preprocessing attack, e.g., if the preprocessing attacker finds some message m and an integer r such that $e = H(\tau(r)||m) = 0$, then the attacker can simply include the tuple (m, r) as part of the S -bit hint. Observe that the hint is completely independent of the public key \mathbf{pk} . In fact, for any public key \mathbf{pk} , we have that $\sigma' = (s = r, e = 0)$ is a valid signature for the message m ! To see this, note that $R = \tau(s - \mathbf{sk} \cdot 0) = \tau(r) = 0$ and that, by assumption, $H(R||m) = H(\tau(r)||m) = 0 = e$.

The above attack can easily be addressed with key-prefixing, i.e., to sign a message m , we pick an integer r , compute $e = H(\mathbf{pk}||\tau(r)||m)$, and output the signature $\sigma = (s = r + \mathbf{sk} \cdot e, e)$. Intuitively, since the preprocessing attacker does not know the public key \mathbf{pk} in advance, s/he is unlikely to have stored a tuple of the form $(\mathbf{pk}, m, r, \tau(r))$. The key question is whether or not *short Schnorr signatures with key-prefixing* are secure against *any* preprocessing attack. To prove that short Schnorr signatures with key-prefixing are secure against preprocessing attacks, we revisit the 1-out-of- N bridge-finding game in the preprocessing setting.

Remark 2. Several standardized implementations of Schnorr signatures, such as BSI-TR-03111 [fIS18] or ISO/IEC 14888-3 [fSC18], slightly deviate from Schnorr’s original construction by explicitly disallowing $e = 0$ signatures. Disallowing $e = 0$ signatures defends against the particular preprocessing attack we described above, and it is possible that the particular preprocessing attack was one of the motivations for excluding $e = 0$ signatures in these standards. Investigating the (in)security of the BSI/ISO version of short Schnorr signatures against preprocessing attacks remains an interesting open research challenge. \triangleleft

5.1 Security of BRIDGE^N-Finding Game with Preprocessing

We analyze the BRIDGE^N-finding game in the setting with preprocessing attacks. In particular, an attacker consists of a pair of algorithms $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$. The basic idea is that we split the attack into two phases, *preprocessing* and *online* phase, so that the attacker has (exponential) time to make preprocessing queries before playing the bridge game. Specifically,

- Algorithm \mathcal{A}_{pre} runs a preprocessing phase, where it takes as input $\mathbf{g} = \tau(1)$ and outputs a hint str_τ , which is a binary string after making queries to the generic group oracles $\mathbf{GO} = (\text{Mult}(\cdot, \cdot), \text{Inv}(\cdot))$. Without loss of generality, we can assume that \mathcal{A}_{pre} is deterministic, and we simply use str_τ to refer to the hint when the random mapping τ is fixed.
- The online attacker \mathcal{A}_{on} attempts to win the BRIDGE^N -finding game. The online attacker \mathcal{A}_{on} is given the hint str_τ (which was produced in the preprocessing phase), as well as $(\tau(x_1), \dots, \tau(x_N))$. However, the challenger picks $(x_1, \dots, x_N) \in \mathbb{Z}_p^N$ after the hint str_τ is fixed. For convenience, we will write $\mathcal{A}_{\text{on}, \text{str}_\tau}$ to denote the online attacker with the hint str_τ hardcoded.

We are interested in the setting where the preprocessing algorithm \mathcal{A}_{pre} can make $q_{\mathbb{G}}^{\text{pre}} \geq 2^{2k}$ queries to the generic group oracles. In other words, the preprocessing algorithm \mathcal{A}_{pre} can examine the entire input/output table of the mapping τ . However, the length of the hint str_τ given to the online attacker is bounded by S , and the online attacker can make at most $q_{\mathbb{G}}^{\text{on}} < 2^k$ queries to the generic group oracles. [Theorem 7](#) says that the probability of a successful preprocessing attack is at most $\tilde{\mathcal{O}}(SN(q_{\mathbb{G}}^{\text{on}})^2/p)$.

Theorem 7. *Let $p > 2^{2k}$ be a prime number and $N \in \mathbb{N}$ be a parameter. Let $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that \mathcal{A}_{pre} outputs an S -bit hint and \mathcal{A}_{on} makes at most $q_{\mathbb{G}}^{\text{on}} := q_{\mathbb{G}}^{\text{on}}(k)$ queries to the generic group oracles. Then*

$$\Pr \left[\text{BridgeChal}_{\mathcal{A}_{\text{on}, \text{str}_\tau}}^{\tau, N}(k) = 1 \right] \leq \tilde{\mathcal{O}} \left(\frac{SN(q_{\mathbb{G}}^{\text{on}} + N)(q_{\mathbb{G}}^{\text{on}} + 2N)}{p} \right),$$

where the randomness is taken over the selection of τ , the random coins of \mathcal{A}_{on} , and the random coins used by the challenger in the bridge game (the hint $\text{str}_\tau = \mathcal{A}_{\text{pre}}^{\text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger). In particular, if $q_{\mathbb{G}}^{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then

$$\Pr \left[\text{BridgeChal}_{\mathcal{A}_{\text{on}, \text{str}_\tau}}^{\tau, N}(k) = 1 \right] \leq \frac{12SN(q_{\mathbb{G}}^{\text{on}})^2 \log p}{p}.$$

Remark 3. The upper bound is essentially tight as a preprocessing attacker can solve a random 1-out-of- N discrete-log challenge with probability $\tilde{\Omega}((q_{\mathbb{G}}^{\text{on}})^2 S/p)$ which would trivially allow the attacker to win the bridge-finding game. In particular, even when $N = 1$, there is a preprocessing with success probability $\Omega((q_{\mathbb{G}}^{\text{on}})^2 S/p)$, e.g., see [[CK18](#), Section 7.1]. Thus, our upper bound is tight up to a factor of N . \triangleleft

The proof of [Theorem 7](#) closely follows [[CK18](#), Theorem 2] with a few minor modifications, and the full proof can be found in [Appendix C.3](#). One small difference is that we need to extend the proofs of [[CK18](#)] to handle queries to the inverse oracle $\text{Inv}(\cdot)$, and the restricted discrete log oracle DLog . The proof of [Theorem 7](#) relies on [Lemma 2](#), which is similar to [[CK18](#), Lemma 4]. Intuitively, if the preprocessing attack is too successful, then one can derive a contradiction by compressing the random mapping τ .

Lemma 2. *Let \mathbb{G} be the set of binary strings of length ℓ such that $2^\ell \geq p$ for a prime p . Let $\mathcal{T} = \{\tau_1, \tau_2, \dots\}$ be a subset of the labeling functions from \mathbb{Z}_p to \mathbb{G} . Let $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ be a pair of generic algorithms for \mathbb{Z}_p on \mathbb{G} such that for every $\tau \in \mathcal{T}$ and every $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$, \mathcal{A}_{pre} outputs an S -bit advice string, \mathcal{A}_{on} makes at most q_{on} oracle queries, and $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ satisfy $\Pr_{\mathcal{A}_{\text{on}}} \left[\text{BridgeChal}_{\mathcal{A}_{\text{on}}, \text{str}_\tau}^{\tau, N}(k, \vec{x}) = 1 \right] \geq \varepsilon$, where $\text{str}_\tau = \mathcal{A}_{\text{pre}}^{\text{GO}}(\tau(1))$. Then, there exists a randomized encoding scheme that compresses elements of \mathcal{T} to bitstrings of length at most*

$$\log \frac{|\mathbb{G}|!}{(|\mathbb{G}| - p)!} + S + 1 - \frac{\varepsilon p}{6q_{\text{on}}(q_{\text{on}} + N)(N \log p + 1)},$$

and succeeds with probability at least $1/2$.

The full proof of [Lemma 2](#) can be found in [Appendix C.3](#). Here, we only give the brief idea as follows. To compress τ , our encoding algorithm first runs \mathcal{A}_{pre} to extract an S -bit hint str_τ . We then execute $\mathcal{A}_{\text{on}}(\text{str}_\tau, \tau(x_1), \dots, \tau(x_N))$ multiple times with different challenges x_1, \dots, x_N . During each execution we record the responses to the new generic group oracle queries, so that the decoder can also execute $\mathcal{A}_{\text{on}}(\text{str}_\tau, \tau(x_1), \dots, \tau(x_N))$. Intuitively, whenever the BRIDGE^N event occurs, the decoder can save a few bits by simply recording the index of prior query involved in the collision. This requires just $\log q_{\text{on}}$ bits to encode instead of $\log p$ bits.

5.2 Multi-User Security of Key-Prefixed Short Schnorr Signatures with Preprocessing

[Theorem 7](#) upper bounds the probability that a preprocessing attacker wins the multi-user bridge-finding game. In this setting, we observe that the hint $\text{str} := \text{str}_{\tau, \mathbb{H}}$ that the preprocessing attacker outputs may depend both on the random oracle \mathbb{H} as well as the encoding map τ . We show how to adapt our prior reduction to establish the multi-user security of key-prefixed short Schnorr signatures against preprocessing attackers. Recall that in our reduction, we simulated a signature forgery attacker for (non key-prefixed) short Schnorr signatures responding to queries to the signing oracle by programming the random oracle. In the preprocessing setting without key-prefixing, the reduction breaks down immediately. For example, the probability of a lucky random oracle query $\mathbb{H}(\tau(r) \| m) = 0$ is no longer $\approx q_{\mathbb{H}}^{\text{on}}/2^k$, since the preprocessing attacker can simply hardcode the pair (r, m) as part of the hint $\text{str} := \text{str}_{\tau, \mathbb{H}}$. Similarly, the hint $\text{str} := \text{str}_{\tau, \mathbb{H}}$ may be correlated with particular input/output pairs from the random oracle, making it infeasible to program those points.

We address this challenge by considering a model where a preprocessing attacker is *time-bounded*, i.e., the preprocessing attacker can look at the entire generic group oracles but only allowed to query the random oracle at up to $q_{\mathbb{H}}^{\text{pre}} = 2^{3k}$ points during the preprocessing phase. We leave it as an interesting theoretical challenge whether or not the bounds can be extended to unbounded

preprocessing attacks. However, we would argue that in practice, 2^{3k} greatly overestimates the running time of any preprocessing attacker, e.g., if $k = 112$, then $2^{3k} = 2^{336}$. Intuitively, the signing oracle for key-prefixed short Schnorr signatures involves two random points: a public key $\mathbf{pk} \in \mathbb{G}$ and a random value $r \leftarrow_s \mathbb{Z}_p$. The probability that a preprocessing attacker submitted a query of the form $H(\mathbf{pk}, \tau(r), \cdot)$ is at most $q_H^{\text{pre}} p^{-2} \leq 2^{-k}$, since the q_H^{pre} random oracle queries are fixed before \mathbf{pk} and r are sampled.

In our analysis, we consider the bad event that the signing oracle queries the random oracle at a point $H(\mathbf{pk}_i \| \tau(r_j) \| m)$, which was previously queried by the preprocessing attacker. Note that if this bad event never occurs, then we can view $H(\mathbf{pk}_i \| \tau(r_j) \| m)$ as a uniformly random string that is uncorrelated with the attacker's state. The probability of this bad event occurring on any single query to the signing oracle is at most $N q_H^{\text{pre}} / p^2$. In particular, fixing an arbitrary set of q_H^{pre} random oracle queries and then sampling $\mathbf{pk}_1, \dots, \mathbf{pk}_N \in \mathbb{G}$ and $r \in \mathbb{Z}_p$, we can apply union bounds to argue that the probability that the preprocessing attacker previously submitted some query of the form $H(\mathbf{pk}_i, \tau(r), \cdot)$ for any i is at most $N q_H^{\text{pre}} / p^2$. Union bounding over the q_S^{on} online queries to the signing oracle, the probability of the bad event ever occurring on any query to the signing oracle is at most $N q_H^{\text{pre}} q_S^{\text{on}} / p^2$. Assuming that the bad event never occurs, we can safely program the random oracle to simulate queries to the signing oracle when we simulate our signature forgery attacker.

The other challenge that arises in the preprocessing setting is upper bounding the probability of the bad event that the attacker forges a signature without causing the bridge event to occur. Previously, our argument relied on the observation that for “fresh” group elements $r \in \mathbb{Z}_p$, we can effectively view $\tau(r)$ as random bit string that is yet to be fixed. This intuition does not carry over into the preprocessing setting, as the hint str might be correlated with $\tau(r)$. We address these challenges by applying a random oracle compression argument. In particular, if the attacker can generate forged signatures without causing the bridge event to occur, we can use this attacker to predict random oracle outputs, allowing us to derive a contraction by compressing the random oracle.

Theorem 8. *Let $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ be a key-prefixed Schnorr signature scheme and $p > 2^{2k}$ be a prime number. Let $N \in \mathbb{N}$ be a parameter and $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that $\mathcal{A}_{\text{sig}}^{\text{pre}}$ makes at most q_H^{pre} queries to the random oracle $H : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$ and outputs an S -bit hint $\text{str}_{\tau, H}$, and $\mathcal{A}_{\text{sig}}^{\text{on}}$ makes at most $q_G^{\text{on}} := q_G^{\text{on}}(k)$ queries to the generic group oracles and at most q_H^{on} queries to the random oracle. Then $\Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}}^{\text{on}}, \text{str}_{\tau, H}, \Pi}^{\tau, N}(k) = 1 \right] \leq \varepsilon$, with*

$$\varepsilon = \tilde{\mathcal{O}} \left(\frac{SN(q_G^{\text{on}} + N)(q_G^{\text{on}} + 2N)}{p} \right) + \frac{N q_H^{\text{pre}} q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} + \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p},$$

where q_S^{on} denotes the number of queries to the signing oracle and the randomness is taken over the selection of τ and the random coins of $\mathcal{A}_{\text{sig}}^{\text{on}}$ (the hint $\text{str}_{\tau, H} = \mathcal{A}_{\text{sig}}^{\text{pre}, \text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger).

In particular, if $q_G^{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then

$$\varepsilon = \frac{12SN(q_G^{\text{on}})^2 \log p}{p} + \frac{Nq_H^{\text{pre}} q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} + \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}.$$

Remark 4. The upper bound in [Theorem 8](#) is essentially tight (up to a factor of N), because of the following observations:

- Making the reasonable assumption that $Nq_H^{\text{pre}} q_S^{\text{on}} < 2^{3k}$ and $q_G^{\text{on}} > \sqrt{N}$, the dominating terms in ε are $\tilde{\mathcal{O}}(SN(q_G^{\text{on}})^2/p)$ and/or $\mathcal{O}(q_H^{\text{on}}/2^{k_1})$.
- A preprocessing attacker can simply solve one of the discrete-log challenges with probability at least $\Omega(S(q_G^{\text{on}})^2/p)$ which would recover a secret key and make it trivial to forge a signature.
- Any attacker who makes $q_H^{\text{on}} \geq q_S^{\text{on}}$ queries to the random oracle can fix an arbitrary message m and pick random numbers $r_1, \dots, r_{q_H^{\text{on}}}$ hoping that $H(\text{pk}_1 \| \tau(r_j) \| m) = 0$ for some $j \leq q_H^{\text{on}}$. In this case, $(r_j, 0)$ is a valid forged signature for m under public key pk_1 . Thus, the attacker can succeed with probability $\approx q_H^{\text{on}}/2^{k_1}$. \triangleleft

The full proof of [Theorem 8](#) can be found in [Appendix C.3](#). The key idea is that we can repeat the essentially same reduction from [Section 4](#), i.e., we can build a bridge-finding game attacker $(\mathcal{A}_{\text{bridge}}^{\text{pre}}, \mathcal{A}_{\text{bridge}}^{\text{on}})$ with preprocessing from the signature forgery attacker $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ with preprocessing, except that when we program a random oracle, we define an additional bad event that we program a random oracle at a point the attacker has already queried the point during the preprocessing phase. We observe that such probability is negligibly small. As long as the failure event does not occur we can program the random oracle and the attacker will not notice the difference.

Instantiating Key-Prefixed Short Schnorr Signatures. We would like to have the success probability in [Theorem 8](#) bounded by $\mathcal{O}(q/2^k)$ for any $q \leq 2^k$, where $q = q_G^{\text{on}} + q_H^{\text{on}} + q_S^{\text{on}}$ is the total number of *online* queries made by a preprocessing attacker. To achieve k bits of multi-user security for key-prefixed short Schnorr signatures with preprocessing, we can fix p such that $p \approx 2^{2k} SN \log p$, and set the length of our hash output to be $k_1 = k$. With these parameters, [Theorem 8](#) tells us that a preprocessing attacker wins the signature forgery game with probability at most $\varepsilon = \mathcal{O}((q_H^{\text{on}} + q_G^{\text{on}})/2^k)$. The length of the signatures we obtain will be $k + \log p = 3k + \log N + \log S + \log \log p$.

As a concrete example, if $N \leq 2^{k/4}$ and $S \leq 2^{k/2}$, then we obtain signatures of length $\approx 3.75k + \log 2.75k$. If we want $k \geq 128$ bits of security, then the assumption that $N < 2^{k/4}$ seems quite reasonable, since $2^{32} (\approx 4.3 \text{ billion})$ is over half of the current global population, and 2^{64} bits exceeds the storage capacity of Facebook’s data warehouse⁵. As a second example, if we take $S \leq 2^{80}$ as an upper bound on the storage capacity of any nation state and $N \approx 2^{40}$, then we obtain signatures of length $\approx 3k + 120 + \log(2k + 120)$.

⁵ See the link: <https://engineering.fb.com/2014/04/10/core-data/scaling-the-facebook-data-warehouse-to-300-pb/> (Retrieved 2/20/2021)

6 Multi-User Security of Other Fiat-Shamir Signatures

In this section, we show that our techniques from [Section 4](#) and [Section 5](#) apply to other Fiat-Shamir-based signature schemes. We apply our reductions to analyze the multi-user security of the full-domain hash variant of Chaum-Pedersen signatures [[CP93](#)], and (short) Katz-Wang signatures [[KW03](#)], with and without preprocessing. In practice, the full-domain hash variant of Chaum-Pedersen would be used to ensure that our signature scheme supports the message space $m \in \{0, 1\}^*$ instead of requiring that m is a group element. We begin by introducing regular Chaum-Pedersen signatures in the next paragraph before describing the full-domain hash variant (Chaum-Pedersen-FDH) that we analyze.

Security Analysis of Chaum-Pedersen-FDH Signatures. The Chaum-Pedersen signature scheme [[CP93](#)] is obtained by applying the Fiat-Shamir transform [[FS87](#)] to the Chaum-Pedersen identification scheme and works as follows.

- Given a cyclic group $G = \langle g \rangle$ of prime order p , the key generation algorithm picks $\text{sk} \leftarrow \mathbb{Z}_p$ and sets $\text{pk} = g^{\text{sk}}$.
- To sign a message $m \in G$ with the secret key sk , we sample $r \leftarrow \mathbb{Z}_p$ and compute $y = m^{\text{sk}}$, $a = g^r$, $b = m^r$, and $e = \text{H}(m\|y\|a\|b)$. Finally, we output a signature $\sigma = (y, a, b, s)$, where $s := r + \text{sk} \cdot e \pmod p$.
- The verification algorithm takes as inputs a signature $\sigma' = (y', a', b', s')$ and computes $e' = \text{H}(m\|y'\|a'\|b')$, $A = g^{s'}$, $B = a'g^{\text{sk} \cdot e'}$, $C = m^{s'}$ and $D = b'y'^{e'}$. Finally, we verify that $(A = B)$ and $(C = D)$ before accepting the signature.

The *full-domain hash variant* of Chaum-Pedersen signature, say *Chaum-Pedersen-FDH signature*, is obtained by hashing a message m into a group element so that we can perform generic group operations when signing the message. That is, in the generic group model, we compute $h = \text{H}'(\text{pk}\|m) := \text{Pow}(\mathbf{g}, \text{H}(\text{pk}\|m))$ and compute $\mathbf{h} = \text{Pow}(h, \text{sk})$ and $\mathbf{b} = \text{Pow}(h, r)$ (which corresponds to $y = h^{\text{sk}}$ and $b = h^r$ when instantiated with a cyclic group $G = \langle g \rangle$) during the signing procedure. Note that key-prefixing is necessary as otherwise an attacker can always forge a signature for a message m , e.g., simply find $m \neq m'$ such that $\text{H}(m) = \text{H}(m')$. The full description for each of these algorithms can be found in [Figure 4](#) in [Appendix A](#).

Our reduction in [Section 4](#) naturally extends to Chaum-Pedersen-FDH signature scheme by using signing oracle in [Figure 3](#). The signing oracle is able to generate valid signatures *without* the secret key by programming the random oracle. This allows us to prove [Theorem 9](#). We remark that a Chaum-Pedersen-FDH signature with k bits of security has length $8k$ — each group element requires $2k$ bits to encode since $p \approx 2^{2k}$. Note that reducing the length of the hash output does not have any effect on Chaum-Pedersen-FDH signature length. Thus, we assume that H is a random oracle with $2k$ -bit outputs. The proof of [Theorem 9](#) can be found in [Appendix C.4](#).

Theorem 9. *The Chaum-Pedersen-FDH signature scheme is $(N, q_{\text{H}}, q_{\text{G}}, q_{\text{S}}, \mathcal{O}\left(\frac{q+N}{2^k}\right))$ -MU-UF-CMA secure under the generic group model of prime order $p \approx 2^{2k}$ and*

the programmable random oracle model, where q denotes the total number of queries made by an adversary.

We can also show that the *key-prefixed* Chaum-Pedersen-FDH signature scheme is secure against preprocessing attacks. That is, we apply key-prefixing when computing e , i.e. $e \leftarrow \mathsf{H}(\mathsf{pk}\|h\|\eta\|\mathbf{a}\|\mathbf{b})$ during the signing procedure and $e' \leftarrow \mathsf{H}(\mathsf{pk}\|h\|\eta'\|\mathbf{a}'\|\mathbf{b}')$ during the verification (see Figure 4). During the online phase we can request a signature σ for m and output $\sigma' = \sigma$ as our forgery for m' . We defer the full proof of Theorem 10 to Appendix C.4.

Theorem 10. *Let $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ be a key-prefixed Chaum-Pedersen-FDH signature scheme and $p > 2^{2k}$ be a prime number. Let $N \in \mathbb{N}$ be a parameter and $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that $\mathcal{A}_{\text{sig}}^{\text{pre}}$ makes at most $q_{\text{H}}^{\text{pre}} < 2^{3k}$ queries to the random oracle $\mathsf{H} : \{0, 1\}^* \rightarrow \{0, 1\}^{2k}$ and outputs an S -bit hint $\text{str}_{\tau, \text{H}}$, and $\mathcal{A}_{\text{sig}}^{\text{on}}$ makes at most $q_{\text{G}}^{\text{on}} := q_{\text{G}}^{\text{on}}(k)$ queries to the generic group oracles and at most q_{H}^{on} queries to the random oracle. Then $\Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}}^{\text{pre}}, \text{str}_{\tau, \text{H}}, \Pi}^{\tau, N}(k) = 1 \right] \leq \varepsilon$, with*

$$\varepsilon = \tilde{\mathcal{O}} \left(\frac{SN(q_{\text{G}}^{\text{on}} + N)(q_{\text{G}}^{\text{on}} + 2N)}{p} \right) + \frac{Nq_{\text{H}}^{\text{pre}}q_{\text{S}}^{\text{on}}}{p^2} + \frac{q_{\text{S}}^{\text{on}}(q_{\text{S}}^{\text{on}} + q_{\text{H}}^{\text{on}})}{p} + \frac{4(q_{\text{H}}^{\text{on}} + \tilde{q}_{\text{on}}^2 + 1)}{2^{2k}} + \frac{3N^2(S + 2k)}{2p},$$

where q_{S}^{on} denotes the number of queries to the signing oracle, $\tilde{q}_{\text{on}} = q_{\text{H}}^{\text{on}} + 2q_{\text{S}}^{\text{on}}$, and the randomness is taken over the selection of τ and the random coins of $\mathcal{A}_{\text{sig}}^{\text{on}}$ (the hint $\text{str}_{\tau, \text{H}} = \mathcal{A}_{\text{sig}}^{\text{pre}, \text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger).

Applying Theorem 10, we can fix p such that $p \approx 2^{2k}SN \log p$ to achieve k bits of multi-user security. The final signature size would be $\approx 8k + 4 \log S + 4 \log N + 4 \log(2k + \log SN)$.

Security Analysis of Katz-Wang Signatures. The Katz-Wang signature scheme [KW03] is a double generator version of Schnorr signature scheme. In the generic group model on a cyclic group G of prime order p , we have two generators $p_1, p_2 \in \mathbb{Z}_p$ so that we can associate with g^{p_1} and g^{p_2} to the generators of the group G . Here, the message space for m is arbitrary, i.e., $m \in \{0, 1\}^*$.

Given our encoding $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ and $\mathbf{g} = \tau(1)$, our key generation algorithm picks $\text{sk} \leftarrow_{\$} \mathbb{Z}_p$ and sets $\mathsf{pk} = (p_1, p_2, \mathbf{h}_1, \mathbf{h}_2)$, where $\mathbf{h}_i = \text{Pow}(\tau(p_i), \text{sk})$ for $i = 1, 2$. To sign a message $m \in \{0, 1\}^*$ with the secret key sk , we sample $r \leftarrow_{\$} \mathbb{Z}_p$, and compute $\mathbf{a}_i = \text{Pow}(\tau(p_i), r)$ for $i = 1, 2$, $e = \mathsf{H}(\mathsf{pk}\|\mathbf{a}_1\|\mathbf{a}_2\|m)$, and $s = r + \text{sk} \cdot e \pmod p$. Finally, we output $\sigma = (s, e)$. The verification algorithm takes as inputs a signature $\sigma' = (s', e')$, $\mathsf{pk} = (p_1, p_2, \mathbf{h}_1, \mathbf{h}_2)$ and the message m , and compute $\mathbf{a}'_i = \text{Mult}(\text{Pow}(\tau(p_i), s'), \text{Pow}(\text{Inv}(\mathbf{h}_i), e'))$ for $i = 1, 2$. Finally, we verify that $e' = \mathsf{H}(\mathsf{pk}\|\mathbf{a}'_1\|\mathbf{a}'_2\|m)$ before accepting the signature. The pseudocode for each of these algorithms can be found in Figure 5 in Appendix A.

We remark that the length of a regular Katz-Wang signature is $4k$ bits when $p \approx 2^{2k}$. Similar to short Schnorr signatures, one can shorten the length of the

Sign_j(m_i) without secret key x_j, j ∈ [N] (Chaum-Pedersen-FDH)

```

1: Pick si and ei ∈ ℤp randomly
2: Compute si = Pow(g, si) and hij = Pow(g, H(pkj||mi))
3: Compute ηi = Pow(pkj, H(pkj||mi))
4: Compute ai = Mult(si, Pow(Inv(pkj), ei))
5: Compute bi = Mult(Pow(hij, si), Pow(Inv(ηi), ei))
6: if H(hij||ηi||ai||bi) ∈ prior query then
7:   return ⊥
8: else Program H(hij||ηi||ai||bi) := ei
9:   return σi = (ηi, ai, bi, si)

```

Sign_j(m_i) without secret key x_j, j ∈ [N] (Katz-Wang)

```

1: Pick si, ei ∈ ℤp randomly
2: Compute a1,i = Mult(Pow(τ(p1), si), Pow(Inv(Pow(τ(xj), p1)), ei))
3: Compute a2,i = Mult(Pow(τ(p2), si), Pow(Inv(Pow(τ(xj), p2)), ei))
4: if H(pkj||a1,i||a2,i||mi) ∈ prior query then
5:   return ⊥
6: else Program H(pkj||a1,i||a2,i||mi) := ei
7:   return σi = (si, ei)

```

Figure 3. The signing oracle without secret key in the Chaum-Pedersen-FDH scheme (top) and the Katz-Wang scheme (bottom). Note that $\text{pk}_j = \tau(x_j)$ is public in both schemes while the signing oracle has no information about x_j . We further remark that in the key-prefixed Chaum-Pedersen-FDH scheme, the only difference is to do a key-prefixing $\text{pk}_j = \tau(x_j)$ to the input of the random oracle (line 6 and 8).

hash output to k bits to obtain $3k$ bit signature. Essentially the same reduction can be used to demonstrate the multi-user security of (short) Katz-Wang signatures, while we use the signing oracle in Figure 3 without the secret key.

We observe that Katz-Wang signature is already key-prefixed. The security bounds in Theorem 11 and Theorem 12 are equivalent to our bounds for short Schnorr signatures with and without pre-processing. Thus, we obtain $3k$ (resp. $3k + \log N + \log S + \log(2k + \log NS)$)-bit signatures with k bits of security in the multi-user setting without preprocessing (resp. with preprocessing). As before in the preprocessing setting we select our prime number $p \approx 2^{2k} NS \log(2k + \log NS)$ and we fix the length of the hash output to be $k_1 = k$. We defer the full proof of Theorem 11 and Theorem 12 to Appendix C.4.

Theorem 11. *The (short) Katz-Wang signature scheme is $(N, q_H, q_G, q_S, \mathcal{O}\left(\frac{q+N}{2^k}\right))$ -MU-UF-CMA secure under the generic group model of prime order $p \approx 2^{2k}$ and the programmable random oracle model, where q denotes the total number of queries made by an adversary.*

Kiltz et al. [KMP16] showed that if the decisional Diffie-Hellman problem is (t, ϵ) -hard then an adversary who tries to forge one out of N (regular) Katz-

Wang signatures running at most time t' can succeed with the probability $\varepsilon' \leq t'(4\varepsilon/t + q_S/p + 1/2^k)$. While their result is similar to [Theorem 11](#), our bounds apply to (short) Katz-Wang signatures, with and without preprocessing.

Theorem 12. *Let $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ be a Katz-Wang signature scheme and $p > 2^{2k}$ be a prime number. Let $N \in \mathbb{N}$ be a parameter and $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that $\mathcal{A}_{\text{sig}}^{\text{pre}}$ makes at most $q_H^{\text{pre}} < 2^{3k}$ queries to the random oracle at most $q_H^{\text{pre}} < 2^{3k}$ queries to the random oracle $H : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$ and outputs an S -bit hint $\text{str}_{\tau, H}$, and $\mathcal{A}_{\text{sig}}^{\text{on}}$ makes at most $q_G^{\text{on}} := q_G^{\text{on}}(k)$ queries to the generic group oracles and at most q_H^{on} queries to the random oracle. Then $\Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, H}}^{\tau, N}, \Pi}(k) = 1 \right] \leq \varepsilon$, with*

$$\varepsilon = \tilde{\mathcal{O}} \left(\frac{SN(q_G^{\text{on}} + N)(q_G^{\text{on}} + 2N)}{p} \right) + \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} + \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p},$$

where q_S^{on} denotes the number of queries to the signing oracle and the randomness is taken over the selection of τ and the random coins of $\mathcal{A}_{\text{sig}}^{\text{on}}$ (the hint $\text{str}_{\tau, H} = \mathcal{A}_{\text{sig}}^{\text{pre}, \text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger).

Acknowledgements

Jeremiah Blocki was supported in part by the National Science Foundation under NSF CAREER Award CNS-2047272 and NSF Awards CNS-1704587 and CNS-1755708 and CCF-1910659. Seunghoon Lee was supported in part by NSF Award CNS-1755708 and by the Center for Science of Information (NSF CCF-0939370). The opinions in this paper are those of the authors and do not necessarily reflect the position of the National Science Foundation.

References

- AB20. Handan Kilinc Alper and Jeffrey Burdges. Two-round trip schnorr multi-signatures via delinearized witnesses. Cryptology ePrint Archive, Report 2020/1245, 2020. <https://eprint.iacr.org/2020/1245>.
- BCJ08. Ali Bagherzandi, Jung Hee Cheon, and Stanislaw Jarecki. Multisignatures secure under the discrete logarithm assumption and a generalized forking lemma. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *ACM CCS 2008*, pages 449–458. ACM Press, October 2008.
- BD20. Mihir Bellare and Wei Dai. The multi-base discrete logarithm problem: Tight reductions and non-rewinding proofs for Schnorr identification and signatures. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *INDOCRYPT 2020*, volume 12578 of *LNCS*, pages 529–552. Springer, Heidelberg, December 2020.
- Ber15. Daniel J. Bernstein. Multi-user Schnorr security, revisited. Cryptology ePrint Archive, Report 2015/996, 2015. <http://eprint.iacr.org/2015/996>.

- BL12. Daniel J. Bernstein and Tanja Lange. Two grumpy giants and a baby. Cryptology ePrint Archive, Report 2012/294, 2012. <http://eprint.iacr.org/2012/294>.
- BL19. Jeremiah Blocki and Seunghoon Lee. On the multi-user security of short schnorr signatures with preprocessing. Cryptology ePrint Archive, Report 2019/1105, 2019. <https://ia.cr/2019/1105>.
- BLS04. Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. *Journal of Cryptology*, 17(4):297–319, September 2004.
- BN06. Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 390–399. ACM Press, October / November 2006.
- BNPS03. Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-RSA-inversion problems and the security of Chaum’s blind signature scheme. *Journal of Cryptology*, 16(3):185–215, June 2003.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- CK18. Henry Corrigan-Gibbs and Dmitry Kogan. The discrete-logarithm problem with preprocessing. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 415–447. Springer, Heidelberg, April / May 2018.
- CP93. David Chaum and Torben P. Pedersen. Wallet databases with observers. In Ernest F. Brickell, editor, *CRYPTO’92*, volume 740 of *LNCS*, pages 89–105. Springer, Heidelberg, August 1993.
- DEF⁺19. Manu Drijvers, Kasra Edalatnejad, Bryan Ford, Eike Kiltz, Julian Loss, Gregory Neven, and Igors Stepanovs. On the security of two-round multi-signatures. In *2019 IEEE Symposium on Security and Privacy*, pages 1084–1101. IEEE Computer Society Press, May 2019.
- Den02. Alexander W. Dent. Adapting the weaknesses of the random oracle model to the generic group model. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 100–109. Springer, Heidelberg, December 2002.
- DS19. David Derler and Daniel Slamanig. Key-homomorphic signatures: definitions and applications to multiparty signatures and non-interactive zero-knowledge. *Designs, Codes and Cryptography*, 87(6):1373–1413, Jun 2019.
- Fis00. Marc Fischlin. A note on security proofs in the generic model. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 458–469. Springer, Heidelberg, December 2000.
- fis18. Federal Office for Information Security. Elliptic curve cryptography, version 2.1. *Technical Guideline BSI TR-03111*, Jun 2018.
- FJS14. Nils Fleischhacker, Tibor Jager, and Dominique Schröder. On tight security proofs for Schnorr signatures. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 512–531. Springer, Heidelberg, December 2014.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.

- fSC18. International Organization for Standardization and International Electrotechnical Commission. It security techniques – digital signatures with appendix – part 3: Discrete logarithm based mechanisms. *ISO/IEC 14888-3*, Nov 2018.
- FST10. David Freeman, Michael Scott, and Edlyn Teske. A taxonomy of pairing-friendly elliptic curves. *Journal of Cryptology*, 23(2):224–280, April 2010.
- GGH⁺13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- GHS02. Pierrick Gaudry, Florian Hess, and Nigel P. Smart. Constructive and destructive facets of Weil descent on elliptic curves. *Journal of Cryptology*, 15(1):19–46, January 2002.
- GMLS02. S. Galbraith, J. Malone-Lee, and N. P. Smart. Public key signatures in the multi-user setting. *Inf. Process. Lett.*, 83(5):263–266, September 2002.
- GWZ15. Steven D. Galbraith, Ping Wang, and Fangguo Zhang. Computing elliptic curve discrete logarithms with improved baby-step giant-step algorithm. Cryptology ePrint Archive, Report 2015/605, 2015. <http://eprint.iacr.org/2015/605>.
- Hao17. Feng Hao. Schnorr Non-interactive Zero-Knowledge Proof. RFC 8235, September 2017.
- JMV01. Don Johnson, Alfred Menezes, and Scott Vanstone. The elliptic curve digital signature algorithm (ecdsa). *International Journal of Information Security*, 1(1):36–63, Aug 2001.
- JS08. Tibor Jager and Jörg Schwenk. On the equivalence of generic group models. In Joonsang Baek, Feng Bao, Kefei Chen, and Xuejia Lai, editors, *ProvSec 2008*, volume 5324 of *LNCS*, pages 200–209. Springer, Heidelberg, October / November 2008.
- KM07. Neal Koblitz and Alfred Menezes. Another look at generic groups, 2007.
- KMP16. Eike Kiltz, Daniel Masny, and Jiaxin Pan. Optimal security proofs for signatures from identification schemes. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 33–61. Springer, Heidelberg, August 2016.
- KW03. Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM CCS 2003*, pages 155–164. ACM Press, October 2003.
- LM17. Bei Liang and Aikaterini Mitrokotsa. Fast and adaptively secure signatures in the random oracle model from indistinguishability obfuscation. Cryptology ePrint Archive, Report 2017/969, 2017. <http://eprint.iacr.org/2017/969>.
- MPSW19. Gregory Maxwell, Andrew Poelstra, Yannick Seurin, and Pieter Wuille. Simple schnorr multi-signatures with applications to bitcoin. *Designs, Codes and Cryptography*, 87(9):2139–2164, 2019.
- MVO91. Alfred Menezes, Scott A. Vanstone, and Tatsuaki Okamoto. Reducing elliptic curve logarithms to logarithms in a finite field. In *23rd ACM STOC*, pages 80–89. ACM Press, May 1991.
- Nec94. V. I. Nechaev. Complexity of a determinate algorithm for the discrete logarithm. *Math Notes*, 55:165, 1994.

- NRS20. Jonas Nick, Tim Ruffing, and Yannick Seurin. MuSig2: Simple two-round schnorr multi-signatures. Cryptology ePrint Archive, Report 2020/1261, 2020. <https://eprint.iacr.org/2020/1261>.
- NRSW20. Jonas Nick, Tim Ruffing, Yannick Seurin, and Pieter Wuille. MuSig-DN: Schnorr multi-signatures with verifiably deterministic nonces. Cryptology ePrint Archive, Report 2020/1057, 2020. <https://eprint.iacr.org/2020/1057>.
- NSW09. Gregory Neven, Nigel Smart, and Bogdan Warinschi. Hash function requirements for schnorr signatures. *Journal of Mathematical Cryptology*, 3, 05 2009.
- PH06. S. Pohlig and M. Hellman. An improved algorithm for computing logarithms over $GF(p)$ and its cryptographic significance (corresp.). *IEEE Trans. Inf. Theor.*, 24(1):106–110, September 2006.
- Pol78. John M. Pollard. Monte Carlo methods for index computation mod p . *Mathematics of Computation*, 32:918–924, 1978.
- PS96. David Pointcheval and Jacques Stern. Security proofs for signature schemes. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 387–398. Springer, Heidelberg, May 1996.
- RW14. Kim Ramchen and Brent Waters. Fully secure and fast signing from obfuscation. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014*, pages 659–673. ACM Press, November 2014.
- Sch90. Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 239–252. Springer, Heidelberg, August 1990.
- Seu12. Yannick Seurin. On the exact security of Schnorr-type signatures in the random oracle model. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 554–571. Springer, Heidelberg, April 2012.
- Sha71. D. Shanks. Class number, a theory of factorization, and genera. In *1969 Number Theory Institute (Proc. Sympos. Pure Math., Vol. XX, State Univ. New York, Stony Brook, N.Y., 1969)*, pages 415–440, 1971.
- Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997.
- SJ00. Claus-Peter Schnorr and Markus Jakobsson. Security of signed ElGamal encryption. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 73–89. Springer, Heidelberg, December 2000.
- Sma99. Nigel P. Smart. The discrete logarithm problem on elliptic curves of trace one. *Journal of Cryptology*, 12(3):193–196, June 1999.
- SW14. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.
- WZ11. Ping Wang and Fangguo Zhang. Computing elliptic curve discrete logarithms with the negation map. Cryptology ePrint Archive, Report 2011/008, 2011. <http://eprint.iacr.org/2011/008>.

A Full Description of Other Fiat-Shamir-Based Signatures

A.1 (Key-Prefixed) Chaum-Pedersen-FDH Signature Scheme

$\text{Kg}(1^k)$:	$\text{Sign}(\text{sk}, m)$:	$\text{Vfy}(\text{pk}, m, \sigma)$:
1: $\text{sk} \leftarrow_{\mathcal{S}} \mathbb{Z}_p$	1: $r \leftarrow_{\mathcal{S}} \mathbb{Z}_p$	1: Parse $\sigma = (\eta', a', b', s')$
2: $\text{pk} \leftarrow \text{Pow}(\mathfrak{g}, \text{sk})$	2: $h \leftarrow \text{Pow}(\mathfrak{g}, \text{H}(\text{pk} m))$	2: $h \leftarrow \text{Pow}(\mathfrak{g}, \text{H}(\text{pk} m))$
3: return (pk, sk)	3: $\eta \leftarrow \text{Pow}(h, \text{sk})$	3: $e' \leftarrow \text{H}(h \eta' a' b')$
	4: $a \leftarrow \text{Pow}(\mathfrak{g}, r)$	4: $A \leftarrow \text{Pow}(\mathfrak{g}, s')$
	5: $b \leftarrow \text{Pow}(h, r)$	5: $B \leftarrow \text{Mult}(a', \text{Pow}(\text{pk}, e'))$
	6: $e \leftarrow \text{H}(h \eta a b)$	6: $C \leftarrow \text{Pow}(h, s')$
	7: $s \leftarrow r + \text{sk} \cdot e \pmod p$	7: $D \leftarrow \text{Mult}(b', \text{Pow}(\eta', e'))$
	8: return $\sigma = (\eta, a, b, s)$	8: if $(A = B) \wedge (C = D)$ then
		9: return 1
		10: else return 0

Figure 4. The Chaum-Pedersen-FDH signature scheme in the generic group model. Note that in the *key-prefixed* Chaum-Pedersen-FDH signature scheme, we include the public key in the input to the random oracle, i.e., $e \leftarrow \text{H}(\text{pk}||h||\eta||a||b)$ on line 6 in $\text{Sign}(\text{sk}, m)$ and $e' \leftarrow \text{H}(\text{pk}||h||\eta'||a'||b')$ on line 3 in $\text{Vfy}(\text{pk}, m, \sigma)$.

A.2 Katz-Wang Signature Scheme

$\text{Kg}(1^k)$:	$\text{Sign}(\text{sk}, m)$:	$\text{Vfy}(\text{pk}, m, \sigma)$:
1: $\text{sk} \leftarrow_{\mathcal{S}} \mathbb{Z}_p$	1: $r \leftarrow_{\mathcal{S}} \mathbb{Z}_p$	1: Parse $\sigma = (s', e')$
2: $\mathfrak{h}_1 \leftarrow \text{Pow}(\tau(p_1), \text{sk})$	2: $a_1 \leftarrow \text{Pow}(\tau(p_1), r)$	2: Parse $\text{pk} = (p_1, p_2, \mathfrak{h}_1, \mathfrak{h}_2)$
3: $\mathfrak{h}_2 \leftarrow \text{Pow}(\tau(p_2), \text{sk})$	3: $a_2 \leftarrow \text{Pow}(\tau(p_2), r)$	3: $a'_1 \leftarrow \text{Mult}(\text{Pow}(\tau(p_1), s'), \text{Pow}(\text{Inv}(\mathfrak{h}_1), e'))$
4: $\text{pk} \leftarrow (p_1, p_2, \mathfrak{h}_1, \mathfrak{h}_2)$	4: $e \leftarrow \text{H}(\text{pk} a_1 a_2 m)$	4: $a'_2 \leftarrow \text{Mult}(\text{Pow}(\tau(p_2), s'), \text{Pow}(\text{Inv}(\mathfrak{h}_2), e'))$
5: return (pk, sk)	5: $s \leftarrow r + \text{sk} \cdot e \pmod p$	5: if $e' = \text{H}(\text{pk} a'_1 a'_2 m)$ then
	6: return $\sigma = (s, e)$	6: return 1
		7: else return 0

Figure 5. The Katz-Wang Signature Scheme in the generic group model.

B Comparing Different Generic Group Models

In our analysis we show that short Schnorr signatures provide k bits of security in Shoup's generic group model [Sho97]. Schnorr and Jakobsson's model [SJ00]

previously established the security of short Schnorr signatures in their version of the generic group model using programmable random oracles. Similarly, Kiltz et al. [KMP16] proved that *regular* Schnorr signatures provide k bits of security in the multi-user setting using a different version of the generic group model. In this section, we review the different variants of the generic group model, and motivate why we chose to conduct our analysis in Shoup’s model [Sho97]. As a motivating example, we consider a recent generic preprocessing attack on the squared Decisional Diffie-Hellman problem (sqDDH) [?] due to Corrigan-Gibbs and Kogan [CK18]. While it is straightforward to describe the attack in Shoup’s generic group model [Sho97] it does appear to be not possible to describe the attack in other variations of the generic group model.

B.1 The Generic Group Model based on Collisions

In the generic group model of Shoup, the attacker is given a handle $\tau(h)$ for any group element that is the output of any generic group query. In the model of Schnorr and Jakobsson [SJ00], the attacker is not directly given a handle. In particular, if f_i denotes the output of the i^{th} generic group query the attacker is simply informed whether or not f_i is a new group element or whether f_i collided with a prior query. The attacker may *indirectly* reference previously computed group elements by submitting a query $(a_1, \dots, a_{i-1}) \in \mathbb{Z}_p^{i-1}$ to the generic group oracle. The attacker is then informed whether or not the group element $f_i := \prod_{j=1}^{i-1} f_j^{a_j}$ is new or not. If $f_i \in \{f_1, \dots, f_{i-1}\}$ then the attacker is given the index $j < i$ of any group element such that $f_j = f_i$. The formal definition of a *generic algorithm*, as defined by Schnorr and Jakobsson [SJ00], is given in Definition 3.

Definition 3. [SJ00] A generic algorithm (of Schnorr and Jakobsson) is a sequence of t generic steps; for time $1 \leq t' < t$, the algorithm takes inputs as $f_1, \dots, f_{t'} \in G$, and computes $f_i = \prod_{j=1}^{i-1} f_j^{a_j}$ for $i = t' + 1, \dots, t$, where $(a_1, \dots, a_{i-1}) \in \mathbb{Z}_p^{i-1}$ depends arbitrarily on i , the non-group element and the set $\mathcal{CO}_{i-1} := \{(j, k) \mid f_j = f_k, 1 \leq j < k \leq i-1\}$ of previous “collisions” of group elements.

Remark 5. While many natural attacks can be modeled using Definition 3, there are several limitations due to the lack of a direct handle on group elements. In Definition 3 the only way to obtain new group elements is by executing another generic step. By contrast, in Shoup’s model, the attacker can pick a binary string x at any time expecting that $x = \tau(h)$ for some group element $h \in G$ and submit x as the input to generic group oracles. Similarly, in Shoup’s model, the attacker can easily “mark” or partition the group elements G . In fact, this can be done *before* the attacker ever queries the generic group oracle(s). For example, we could define $G_0 = \{h : 0 = \tau(h) \bmod 2\}$ and $G_1 = \{h : 1 = \tau(h) \bmod 2\}$. Later when we are given the handle $\tau(h)$ it is trivial to test whether $h \in G_b$. Similarly, if we define $G_0 = \{h : \tau(h) = 0 \bmod \sqrt{p}\}$ then it is easy to sample elements in G_0 in Shoup’s model, i.e., pick a random x such that $x = 0 \bmod \sqrt{p}$ expecting that $x = \tau(h)$ for some $h \in G$. \triangleleft

Preprocessing Attacks on the sqDDH Problems. As an illustrative example of a generic attack which cannot be described in Schnorr and Jakobsson’s generic group model, we discuss the preprocessing attack on the sqDDH problem [?] proposed by Corrigan-Gibbs and Kogan [CK18]. The sqDDH problem requires to distinguish tuples of the form (g, g^x, g^y) from $(g, g^x, g^{(x^2)})$ for random $x, y \in \mathbb{Z}_p$. Corrigan-Gibbs and Kogan [CK18] introduced a sqDDH distinguisher $\mathcal{D}_{\text{sqDDH}}$ using preprocessing. The preprocessing attack generates a hint of size at most s (bits) in the offline phase after arbitrary interaction with the generic group oracles. In the online phase the attacker is given the pair $(\tau(h_1), \tau(h_2))$ and must guess whether the pair is a valid sqDDH pair. The attack of Corrigan-Gibbs and Kogan [CK18] runs in time t and achieves advantage ϵ provided that $st^2 = \Omega(p\epsilon^2)$. Interestingly, this attack matches the lower bound for the regular DDH problem, i.e., any preprocessing attack which achieves distinguishing advantage ϵ for DDH must have $st^2 = \tilde{\Omega}(p\epsilon^2)$.

Intuitively, the preprocessing phase takes advantage of the ability of the attacker to “mark” and/or “color” exponentially large subsets of pairs $(\mathbf{u}_1, \mathbf{u}_2) \in \mathbb{G}^2$ — a capability that the attacker does not have in Definition 3. In particular, the attack relies on several random functions $H_m : \mathbb{G}^2 \rightarrow \{1, \dots, t\}$ and $H_c : \mathbb{G}^2 \rightarrow \{1, \dots, s\}$ to “mark” and “color” vertices, i.e., the set of marked vertices is $\mathcal{M} = \{(\mathbf{u}_1, \mathbf{u}_2) \in \mathbb{G}^2 : H_m(\mathbf{u}_1, \mathbf{u}_2) = 1\}$ and any marked node $(\mathbf{u}_1, \mathbf{u}_2) \in \mathcal{M}$ is assigned the color $H_c(\mathbf{u}_1, \mathbf{u}_2)$. Next Corrigan-Gibbs and Kogan define a random walk using a random function $f : \tau(G)^2 \rightarrow \mathbb{Z}_p$, i.e., starting at the node $(\tau(h_0), \tau(h_1)) \in \tau(G)^2$ we compute $\alpha \leftarrow f(\tau(h_0), \tau(h_1))$ and move to $(\tau(h_0^\alpha), \tau(h_1^{\alpha^2})) \in \tau(G)^2$. Letting $\mathcal{Y} = \{(\tau(g^x), \tau(g^{x^2})) : x \in \mathbb{Z}_p\} \subset \tau(G)^2$ denote the “yes” instance of the sqDDH problem it is easy to observe that walk that starts inside (resp. outside) \mathcal{Y} remains inside (resp. outside) \mathcal{Y} . The random walk is used to select $\Omega(p/(3t^2))$ marked nodes $\mathcal{T} \subseteq \mathcal{Y}$. Then for each color $c \in \{1, \dots, s\}$ the preprocessing algorithm computes the advice string $w_c \in \{0, 1\}^{\log p}$ such that

$$w_c = \arg \max_{w \in \{0, 1\}^{\log p}} \sum_{\substack{(\mathbf{m}, c_m) \in \mathcal{T}: \\ c_m = c}} H(w, \mathbf{m}),$$

where $H : \{0, 1\}^{\log p} \times \mathbb{G}^2 \rightarrow \{0, 1\}$ is a random function. Please see [CK18] for the complete discussion.

We now discuss some of the inherent challenges in modeling the attack why the non-standard generic group model cannot capture the sqDDH distinguishing attack illustrated above. When it comes to sampling a set of p^2/t “marked” points $\mathcal{M} := \{(\tau(y_0), \tau(y_1)) : H_m(\tau(y_0), \tau(y_1)) = 1\}$ we required an explicit handle $\tau(\cdot)$ in addition to the random hash function $H_m : \mathbb{G}^2 \rightarrow \{1, \dots, t\}$. Similarly, coloring each of the marked points using the random function H_m requires an explicit handle. Finally, we also require an explicit handle to compute each of the advice strings w_c using the hash function $H : \{0, 1\}^{\log p} \times \mathbb{G}^2 \rightarrow \{0, 1\}$.

B.2 The Generic Group Model using Incrementing Counters

While version of the generic group model used by Kiltz et al. [KMP16] is different from Schnorr and Jakobsson [SJ00], it is also not equivalent to the generic group model of Shoup [Sho97]. In fact, we show that any generic attack in the model of Kiltz et al. [KMP16] can be simulated within the model of Schnorr and Jakobsson [SJ00]. Thus, the model used by Kiltz et al. [KMP16] would also fail to capture attacks that require explicit handles on group elements such as the sqDDH preprocessing attack [CK18].

In the model of Kiltz et al. [KMP16], the generic group oracle maintains a global counter i which is incremented every time a new group element is produced. The generic group oracle also maintains a list of tuples which consists of the group element along with the corresponding counter, i.e., $(y, C_y) \in \mathbb{Z}_p \times \mathbb{N}$ where \mathbb{N} denotes the set of positive integers. The counter C_y for each group element y is used as a “handle”, e.g., if g and $h = g^x$ are the “public” group elements then the attacker is initially given the counters $C_g = 1$ and $C_h = 2$ and the next group element r that is generated will be assigned counter value $C_r = 3$.

Formally, the generic oracle works as following:

Oracle \mathcal{O}_G :

The oracle \mathcal{O}_G takes input of two counters and output the resulting counter.

1. (Initialization) Let $g = g^1, h_1 = g^{x_1}, \dots, h = g^{x_n}$ be the initial public elements. We add $(1, C_1 = 1)$ (the generator), and $(x_i, C_{x_i} = i + 1)$ for each $x_i \in \mathbb{Z}_p$ to our table. We also set our global counter $i = n + 1$ to count the number of group elements observed so far.
2. On input of two counters (C_a, C_b) , the oracle searches the internal values (a, C_a) and (b, C_b) , and computes $z = a + b \pmod p$.
3. If the tuple (z, C_z) already is in the list, then output the counter C_z .
4. Otherwise, the counter i is increased by 1, the tuple $(z, C_z := i)$ is stored in the list, and the oracle outputs the counter C_z .

Consider a generic attack in the model above. We can translate this attack to an attack within model of Schnorr and Jakobsson [SJ00] by defining $\alpha_{C_a, C_b} \in \mathbb{Z}_p^*$ such that $\alpha_{C_a, C_b}[j] = 0$ for all $j \neq C_a, C_b$. If $C_a = C_b$ then $\alpha_{C_a, C_b}[C_a] = 2$ otherwise $\alpha_{C_a, C_b}[C_a] = \alpha_{C_a, C_b}[C_b] = 1$. If we let f_i be the group element corresponding to the counter i (following the notation of Schnorr and Jakobsson [SJ00]) then the response to the query α_{C_a, C_b} is $\prod_j f_j^{\alpha_{C_a, C_b}[j]} = f_{C_a} f_{C_b}$. Whenever our generic attack submits the query (C_a, C_b) to the oracle \mathcal{O}_G we intercept this query and submit the query α_{C_a, C_b} to the Schnorr/Jakobsson oracle. We are not given the output $f_{C_a} f_{C_b}$ directly, but we are informed whether or not a collision occurred (and where) which allows us to either increment the counter or (in the case of a collision) return the original counter.

Thus, any attack which cannot be described by the model of Schnorr and Jakobsson [SJ00] also cannot be described using the model of Kiltz et al. [KMP16].

Remark 6. Another (minor) limitation is that there is no `Inv` oracle in the Kiltz et al. [KMP16]. Thus, computing h^{-1} requires us to compute h^{p-1} using $\mathcal{O}(\log p)$ queries to the generic group model. Thus, to simulate an attacker who makes t queries to the `Mult`, `Inv` oracles we might require *up to* $t \log p$ queries if we only have the `Mult` oracle. Such a reduction increases running time by a factor of $\mathcal{O}(\log p)$ so we lose a factor of $\log p$ in the security reduction if there is an efficient algorithm to compute h^{-1} directly. \triangleleft

C Missing Proofs

C.1 Missing Proofs from Section 3

Reminder of Lemma 1. *The probability the attacker making at most q_G generic group oracle queries wins the generic discrete-log game $\text{DLogChal}_{\mathcal{A}}^{\tau}(k)$ (even with access to the restricted `DLog` oracle) is at most*

$$\Pr[\text{DLogChal}_{\mathcal{A}}^{\tau}(k) = 1] \leq \frac{6q_G(q_G + 1) + 12}{4p - (3q_G + 2)^2},$$

in the generic group model of prime order p , where the randomness is taken over the selection of τ , the challenge x , as well as any random coins of \mathcal{A} .

Proof of Lemma 1: Without loss of generality, we can assume that for every generic group query involving a “fresh” $\tau(y)$ that the attacker queries `DLog`($\tau(y)$) before making the query (We say `WLOG` because the attacker can always ignore the result.) Thus, we will build up the known sets \mathcal{K} (elements for which we *know* the discrete-log solution and it initially contains $(\tau(1), 0, 1)$) and the partially known set \mathcal{PK}_x (elements for which the discrete-log solution is *partially known* and it initially contains $(\tau(x), 1, 0)$) which are the subsets of the list \mathcal{L} . Then we have the following observations for those sets:

- Initially, \mathcal{K} contains $(\tau(1), 0, 1)$ and \mathcal{PK}_x contains $(\tau(x), 1, 0)$.
- If $(\mathbf{a}_1, 0, c_1), (\mathbf{a}_2, 0, c_2) \in \mathcal{K}$, then $(\text{Mult}(\mathbf{a}_1, \mathbf{a}_2), 0, c_1 + c_2)$ can be added to \mathcal{K} .
- If $(\mathbf{a}_1, a_1, b_1), (\mathbf{a}_2, a_2, b_2) \in \mathcal{PK}_x$, then $(\text{Mult}(\mathbf{a}_1, \mathbf{a}_2), a_1 + a_2, b_1 + b_2)$ can be added to \mathcal{PK}_x except for a special case when $a_1 + a_2 = 0$ in which case the tuple $(\text{Mult}(\mathbf{a}_1, \mathbf{a}_2), 0, b_1 + b_2)$ should be added to \mathcal{K} instead.
- If $(\mathbf{a}, 0, c) \in \mathcal{K}$ and $(\mathbf{b}, a, b) \in \mathcal{PK}_x$, then $(\text{Mult}(\mathbf{a}, \mathbf{b}), a, b + c)$ can be added to \mathcal{PK}_x .
- If $(\mathbf{a}, a, b) \in \mathcal{K}$ (resp. \mathcal{PK}_x), then $(\text{Inv}(\mathbf{a}), -a, -b)$ can be added to \mathcal{K} (resp. \mathcal{PK}_x).

Now consider the event `BRIDGE` which is the event that one of the following occurs when querying `Mult`, `Inv` or `DLog`:

- (1) The output of a query which could be added to \mathcal{K} is already found in \mathcal{PK}_x , and
- (2) The output of a query which could be added to \mathcal{PK}_x is already found in \mathcal{K} .

Intuitively, the attacker wants the event **BRIDGE** to occur, e.g., in either case we have found an element $\tau(y)$ which can be written in two ways; (1) $\tau(y) = \tau(r)$ for some known $r \in \mathbb{Z}_p$ and (2) $\tau(y) = \tau(ax + b)$ for known $a, b \in \mathbb{Z}_p$ with $a \neq 0$. This allows us to solve for $x = (r - b)a^{-1}$.

If the event **BRIDGE** does not occur then after all queries have finished the attacker can still view $x \notin \mathcal{K}$ as an element yet to be sampled from a uniform distribution over a set of size at least $p - |\mathcal{K}| \times |\mathcal{PK}_x|$. To see this note that each pair of distinct elements $\tau(ax + b)$ in \mathcal{PK}_x and $\tau(r)$ in \mathcal{K} eliminates at most one possible value of x , i.e., since $\tau(ax + b)$ and $\tau(r)$ are distinct we have $x \neq (r - b)a^{-1}$.

We further note that $|\mathcal{K}| + |\mathcal{PK}_x| \leq 3q_G + 2$ since each query to the generic group oracles adds at most 3 elements to $\mathcal{K} \cup \mathcal{PK}_x$ — equality holds when both inputs to $\text{Mult}(\cdot, \cdot)$ are fresh. The probability the attacker guesses x correctly is at most

$$\begin{aligned} \Pr [\text{DLogChal}_{\mathcal{A}}^r(k) = 1 \mid \overline{\text{BRIDGE}}] &\leq \frac{1}{p - |\mathcal{K}| \times |\mathcal{PK}_x|} \\ &< \frac{1}{p - (3q_G + 2)^2/4}. \end{aligned}$$

To compute the probability that **BRIDGE** occurs we use our prior observation that the combined size of $|\mathcal{K}|$ and $|\mathcal{PK}_x|$ is at most $3q_G + 2$ and the AM-GM inequality.

Consider the i th query to the generic group oracle and let $\overline{\text{BRIDGE}}_{<i}$ be the event that we have not yet seen a bridge query. Conditioning on this event we can view x as a yet to be selected value uniformly sampled from a set of size at least $p - (3q_G + 2)^2/4$. Let η_i be the output of query i and let (η_i, a_i, b_i) be the tuple that is added to $\mathcal{PK}_x \cup \mathcal{K}$. For each tuple (η, a, b) in $\mathcal{K} \cup \mathcal{PK}_x$ the probability that $x = (b_i - b)(a - a_i)^{-1}$ is at most $\frac{1}{p - (3q_G + 2)^2/4}$. Union bounding over all $(\leq 3i + 2)$ such tuples in $\mathcal{PK}_x \cup \mathcal{K}$ the probability of the event B_i that the i th query bridges is at most

$$\Pr [B_i \mid \overline{\text{BRIDGE}}_{<i}] \leq \frac{3i + 2}{p - (3q_G + 2)^2/4}.$$

Therefore, the probability of **BRIDGE** is upper bounded by

$$\begin{aligned} \Pr[\text{BRIDGE}] &= \sum_{i \leq q_G} \Pr [B_i \mid \overline{\text{BRIDGE}}_{<i}] \\ &\leq \sum_{i \leq q_G} \frac{3i + 2}{p - (3q_G + 2)^2/4} \\ &\leq \frac{3(q_G + 1)q_G/2 + 2}{p - (3q_G + 2)^2/4}. \end{aligned}$$

Thus, the probability the attacker succeeds is upper bounded by⁶

$$\begin{aligned} \Pr [\text{DLogChal}_{\mathcal{A}}^{\tau}(k) = 1] &\leq \Pr [\text{BRIDGE}] + \Pr [\text{DLogChal}_{\mathcal{A}}^{\tau}(k) = 1 \mid \overline{\text{BRIDGE}}] \\ &\leq \frac{6q_{\mathbb{G}}(q_{\mathbb{G}} + 1) + 12}{4p - (3q_{\mathbb{G}} + 2)^2}, \end{aligned}$$

i.e., we would need $q_{\mathbb{G}} = \mathcal{O}(\sqrt{p})$ queries to succeed with constant probability. \square

Reminder of Theorem 4. *The short Schnorr signature scheme $\Pi_{\text{short}} = (\text{Kg}, \text{Sign}, \text{Vfy})$ of length $3k$ is $(q_{\text{H}}, q_{\mathbb{G}}, q_{\text{S}}, \varepsilon)$ -UF-CMA secure with*

$$\varepsilon = \frac{6q_{\mathbb{G}}(q_{\mathbb{G}} + 1) + 12}{4p - (3q_{\mathbb{G}} + 2)^2} + \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p} + \frac{q_{\text{H}} + q_{\text{S}}}{p - (3q_{\mathbb{G}} + 2)} + \frac{q_{\text{H}} + 1}{2^k} = \mathcal{O}\left(\frac{q}{2^k}\right),$$

in the generic group model of prime order $p \approx 2^{2k}$ and the programmable random oracle model, where q denotes the total number of queries made by an adversary.

Proof of Theorem 4: Given an adversary \mathcal{A}_{sig} attacking Schnorr signature scheme, we construct the following efficient algorithm $\mathcal{A}_{\text{dlog}}$ which solves the discrete-logarithm problem relative to the generic group \mathbb{G} using \mathcal{A}_{sig} as a subroutine:

Algorithm $\mathcal{A}_{\text{dlog}}$:

The algorithm is given $\mathbf{g} = \tau(1), \tau(x), p$ as input. The goal of $\mathcal{A}_{\text{dlog}}$ is to output the discrete log of $\tau(x)$, i.e., find x .

1. Initialize the list $\mathcal{L} = \{(\tau(x), 1, 0), (\mathbf{g}, 0, 1)\}$ and $\text{H}_{\text{resp}} = \{\}$ where H_{resp} stores the random oracle queries.
2. Run \mathcal{A}_{sig} with a number of access to the generic oracles $\text{GO} = (\text{Mult}(\cdot, \cdot), \text{Inv}(\cdot), \text{Sign}(\cdot), \text{H}(\cdot))$. The signing oracle without a secret key is described in Figure 6. Now we consider the following cases:
 - (a) Whenever \mathcal{A}_{sig} submits a query w to the random oracle H :
 - If there is a pair $(w, R) \in \text{H}_{\text{resp}}$ for some string R then return R .
 - Otherwise, select $R \in \mathbb{Z}_{2^k}$ uniformly at random and add (w, R) to the set H_{resp} .
 - If w has the form $w = (\mathbf{a}||m)$ where the value \mathbf{a} has not been observed previously (i.e., is not in the list \mathcal{L}) then we query $b = \text{DLog}(\mathbf{a})$ and add $(\tau(b), 0, b)$ to \mathcal{L} .
 - (b) Whenever \mathcal{A}_{sig} submits a query \mathbf{a} to the generic group oracle $\text{Inv}(\mathbf{a})$:
 - First check if \mathbf{a} is not in the list \mathcal{L} . If so we immediately query $b = \text{DLog}(\mathbf{a})$ and add $(\tau(b), 0, b)$ to \mathcal{L} .
 - Otherwise $(\mathbf{a}, a, b) \in \mathcal{L}$ (i.e., $\mathbf{a} = \tau(ax + b)$). In this case we query $\text{Inv}(\mathbf{a}) = \tau(-ax - b)$, output the result and add the result $(\tau(-ax - b), -a, -b) \in \mathcal{L}$.
 - (c) Whenever \mathcal{A}_{sig} submits a query \mathbf{a}, \mathbf{b} to the generic group oracle $\text{Mult}(\mathbf{a}, \mathbf{b})$:

⁶ To see this note that for the events A and B we have that $\Pr[A] = \Pr[A \wedge B] + \Pr[A \wedge \overline{B}] \leq \Pr[B] + \Pr[A|\overline{B}]$.

- First, if the element \mathbf{a} (resp. \mathbf{b}) is not in the list \mathcal{L} then query $b_0 = \text{DLog}(\mathbf{a})$ (resp. $b_1 = \text{DLog}(\mathbf{b})$) and add the element $(\mathbf{a}, 0, b_0)$ (resp. $(\mathbf{b}, 0, b_1)$) to \mathcal{L} .
 - Otherwise both elements $(\mathbf{a}, a_0, b_0), (\mathbf{b}, a_1, b_1) \in \mathcal{L}$. Then we compute $\text{Mult}(\mathbf{a}, \mathbf{b}) = \tau((a_0 + a_1)x + b_0 + b_1)$ and add the element $(\tau((a_0 + a_1)x + b_0 + b_1), a_0 + a_1, b_0 + b_1) \in \mathcal{L}$.
- (d) Whenever \mathcal{A}_{sig} submits a query m to the signing oracle $\text{Sign}(\cdot)$:
- We use the procedure described in Figure 6 to forge a signature without knowledge of the secret key x . Intuitively, the forgery procedure relies on our ability to program the random oracle.
 - We remark that a side effect of querying the Sign oracle is the addition of the tuples $(\tau(s), 0, s)$, $(\tau(xe), e, 0)$ and $(\tau(s - xe), -e, s)$ to \mathcal{L} since these values are computed using the generic group oracles Inv, Mult .
- (e) If at any point we have some string η such that $(\eta, a, b) \in \mathcal{L}$ and $(\eta, c, d) \in \mathcal{L}$ for $(a, b) \neq (c, d)$ then we can immediately return $x = (d - b)(a - c)^{-1}$.⁷ Thus, without loss of generality, we can assume that each string η occurs at most once in the list \mathcal{L} .
3. After \mathcal{A}_{sig} outputs $\sigma = (s, e)$ and m we first compute $I_\sigma = \text{Mult}(\text{Pow}(\mathbf{g}, s), \text{Inv}(\text{Pow}(\tau(x), e))) = \tau(s - xe)$ and then check to see if we previously had any tuple of the form $(I_\sigma, a, b) \in \mathcal{L}$.
- (a) If no such tuple exists we return \perp .
 - (b) Otherwise, we let a, b be given such that $I_\sigma = \tau(ax + b)$.
 - If $a + e = 0$ then we return \perp .
 - Otherwise, we return $x = (s - b)(a + e)^{-1}$.

Analysis. We first remark that if the signature is valid then we must have $e = \text{H}(I_\sigma \| m)$ and $\text{DLog}(I_\sigma) = s - xe = ax + b$ or equivalently $s - b = (a + e)x$. Thus, as long as $(a + e)$ is invertible we will have $x = (s - b)(a + e)^{-1}$.

Now consider the upper bound of the probability that our algorithm outputs \perp for failure before \mathcal{A}_{sig} outputs a signature as well as the probability our algorithm outputs \perp after \mathcal{A}_{sig} outputs a valid signature.

- (1) One way to output failure is during the signing oracle if $\text{H}(I \| m)$ has been queried previously (Algorithm 2.(d)). We define this event as FailtoSign . Note that every time the attacker queries the signing oracle, we would generate a new query to the random oracle H . Thus, we would have at most $q_{\text{H}} + q_{\text{S}}$ input/output pairs recorded for the random oracle. Since $I = \tau(s - xe)$ represents a fresh/randomly selected group element of size p , the probability that (I, m) is one of the inputs is at most $(q_{\text{H}} + q_{\text{S}})/p$. Applying union bound over q_{S} queries to the signing oracle, we conclude that

$$\Pr[\text{FailtoSign}] \leq \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p}.$$

⁷ Note that $(a, b) \neq (c, d)$ implies $a \neq c$ since if $a = c$ then $ax + b = ax + d$ implies $b = d$ as $a, b, c, d \in \mathbb{Z}_p$.




Figure 6. A reduction to the discrete-log attacker $\mathcal{A}_{\text{dlog}}$ from the Schnorr signature attacker \mathcal{A}_{sig} .

- (2) The second way to output failure is if the value I_σ does not previously appear in either set \mathcal{PK}_x or \mathcal{K} (Algorithm 3.(a)). We use $\text{FailtoFind}^N(I_\sigma)$ to denote the event that the signature is valid but the value I_σ does not appear in \mathcal{L} . If $I_\sigma \notin \mathcal{L}$ we can view $I_\sigma = \tau(s - xe)$ as a uniformly random binary string from a set of size at least $p - |\mathcal{L}|$ which had not yet been selected at the time \mathcal{A}_{sig} output σ . Thus, the probability that the query $\text{H}(I_\sigma \| m)$ was previously recorded is at most $(q_H + q_S) / (p - |\mathcal{L}|)$. Observe that if the query $\text{H}(I_\sigma \| m)$ was not previously recorded then the probability of a successful forgery $\text{H}(I_\sigma \| m) = e$ is at most 2^{-k} since we can view $\text{H}(I_\sigma \| m)$ as a uniformly random k -bit string. Hence,

$$\Pr[\text{FailtoFind}(I_\sigma)] \leq \frac{q_H + q_S}{p - |\mathcal{L}|} + \frac{1}{2^k}.$$

- (3) Finally, we could output failure if $a = -e$ (Algorithm 3.(b)). We call this event BadQuery . We first note that by construction we ensure that the tuple (I, a, b) will always be recorded in \mathcal{K} or \mathcal{PK}_x before a query of the form $\text{H}(I \| m)$ is ever issued — if I is new then we call $\text{DLog}(I)$ before querying the random oracle. We call a random oracle query $x = I \| m$ “bad” if $\text{H}(x) = -a$ where the tuple (I, a, b) has already been recorded. Recall that if there were two recorded tuples (I, a, b) and (I, c, d) then our algorithm would have already found x . Thus, the probability each individual query is “bad” is at most 2^{-k} and we can use union bounds to upper bound the probability of any “bad” query as

$$\Pr[\text{BadQuery}] \leq \frac{q_H}{2^k}.$$

Now we have shown that

$$\Pr[\text{DLogChal}_{\mathcal{A}_{\text{dlog}}}^{\tau}(k) = 1]$$

$$\begin{aligned}
 &\geq \Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}}, \Pi}^\tau(k) = 1] - \Pr[\text{FailtoSign}] - \Pr[\text{FailtoFind}(I_\sigma)] - \Pr[\text{BadQuery}] \\
 &\geq \Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}}, \Pi}^\tau(k) = 1] - \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p} - \frac{q_{\text{H}} + q_{\text{S}}}{p - |\mathcal{L}|} - \frac{1}{2^k} - \frac{q_{\text{H}}}{2^k}.
 \end{aligned}$$

Since $|\mathcal{L}| \leq 3q_{\text{G}} + 2$, we can apply [Lemma 1](#) to conclude that

$$\begin{aligned}
 &\Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}}, \Pi}^\tau(k) = 1] \\
 &\leq \frac{6q_{\text{G}}(q_{\text{G}} + 1) + 12}{4p - (3q_{\text{G}} + 2)^2} + \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p} + \frac{q_{\text{H}} + q_{\text{S}}}{p - (3q_{\text{G}} + 2)} + \frac{q_{\text{H}} + 1}{2^k} \\
 &= \mathcal{O}\left(\frac{q}{2^k}\right),
 \end{aligned}$$

where q denotes the total number of queries made by the adversary. Note that for all $q_{\text{G}} < \frac{2^k - 2}{3}$ we have that $(3q_{\text{G}} + 2)^2 < 2^{2k} \approx p$ and therefore it clearly holds that $\frac{6q_{\text{G}}(q_{\text{G}} + 1) + 12}{4p - (3q_{\text{G}} + 2)^2} = \mathcal{O}\left(\frac{q}{2^k}\right)$. \square

C.2 Missing Proofs from [Section 4](#)

Reminder of [Claim 1](#). $\Pr[\text{FailtoSign}] \leq \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p}$.

Proof of [Claim 1](#): Note that every time the attacker queries the signing oracle, we would generate a new query to the random oracle H . Thus, we would have at most $q_{\text{H}} + q_{\text{S}}$ input/output pairs recorded for the random oracle. Since $I_i = \tau(s_i - x_i e_i)$ represents a fresh/randomly selected group element of size p , the probability that (I_i, m_i) is one of the inputs is at most $(q_{\text{H}} + q_{\text{S}})/p$. Applying union bound over q_{S} queries to the signing oracle, we conclude that $\Pr[\text{FailtoSign}] \leq \frac{q_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p}$. \square

Reminder of [Claim 2](#). $\Pr[\text{FailtoFind}(I_{i^*})] \leq \frac{q_{\text{H}} + q_{\text{S}}}{p - |\mathcal{L}|} + \frac{1}{2^k}$.

Proof of [Claim 2](#): If $I_{i^*} \notin \mathcal{L}$ then we can view $I_{i^*} = \tau(s_{i^*} - x_{i^*} e_{i^*})$ as a uniformly random binary string from a set of size at least $p - |\mathcal{L}|$ which had not yet been selected at the time \mathcal{A}_{sig} output σ_{i^*} . Thus, the probability that the query $\text{H}(I_{i^*} \| m_{i^*})$ was previously recorded is at most $(q_{\text{H}} + q_{\text{S}})/(p - |\mathcal{L}|)$. Observe that if the query $\text{H}(I_{i^*} \| m_{i^*})$ was not previously recorded then the probability of a successful forgery $\text{H}(I_{i^*} \| m_{i^*}) = e_{i^*}$ is at most 2^{-k} since we can view $\text{H}(I_{i^*} \| m_{i^*})$ as a uniformly random k -bit string. Hence, we have that $\Pr[\text{FailtoFind}(I_{i^*})] \leq \frac{q_{\text{H}} + q_{\text{S}}}{p - |\mathcal{L}|} + \frac{1}{2^k}$. \square

Reminder of [Claim 3](#). $\Pr[\text{BadQuery}] \leq \frac{q_{\text{H}}}{2^k}$.

Proof of [Claim 3](#): Recall that the event BadQuery happens if $\vec{a} = -e_{i^*} \hat{u}_{i^*}$. Note that by construction we ensure that the tuple (I, \vec{a}, b) will always be

recorded in \mathcal{L} before a query of the form $H(I||m)$ is ever issued — if I is new then we call $D\text{Log}(I)$ before querying the random oracle. Now define a subset $\widehat{\mathcal{L}} \subset \mathcal{L}$ as the set of tuples $(\mathbf{a}, \vec{a}, b) \in \mathbb{G} \times \mathbb{Z}_p^N \times \mathbb{Z}_p$ such that \vec{a} has exactly *one* nonzero element. Now we call a random oracle query $x = (I||m)$ “bad” if $H(x) = -\vec{a}$ where the tuple $(I, \vec{a}, b) \in \widehat{\mathcal{L}}$ has already been recorded and the nonzero element of \vec{a} is \tilde{a} (Recall that if there were two recorded tuples (I, \vec{a}, b) and (I, \vec{c}, d) then our algorithm would have already found a BRIDGE^N instance). Thus, the probability each individual query is “bad” is at most 2^{-k} and we can use union bounds to upper bound the probability of any “bad” query as $\Pr[\text{BadQuery}] \leq \frac{qn}{2^k}$. \square

C.3 Missing Proofs from Section 5

Reminder of Theorem 7. *Let $p > 2^{2k}$ be a prime number and $N \in \mathbb{N}$ be a parameter. Let $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that \mathcal{A}_{pre} outputs an S -bit hint and \mathcal{A}_{on} makes at most $q_{\mathbb{G}}^{\text{on}} := q_{\mathbb{G}}^{\text{on}}(k)$ queries to the generic group oracles. Then*

$$\Pr \left[\text{BridgeChal}_{\mathcal{A}_{\text{on}}, \text{str}_{\tau}}^{\tau, N}(k) = 1 \right] \leq \tilde{O} \left(\frac{SN(q_{\mathbb{G}}^{\text{on}} + N)(q_{\mathbb{G}}^{\text{on}} + 2N)}{p} \right),$$

where the randomness is taken over the selection of τ , the random coins of \mathcal{A}_{on} , and the random coins used by the challenger in the bridge game (the hint $\text{str}_{\tau} = \mathcal{A}_{\text{pre}}^{\text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger). In particular, if $q_{\mathbb{G}}^{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then

$$\Pr \left[\text{BridgeChal}_{\mathcal{A}_{\text{on}}, \text{str}_{\tau}}^{\tau, N}(k) = 1 \right] \leq \frac{12SN(q_{\mathbb{G}}^{\text{on}})^2 \log p}{p}.$$

Proof of Theorem 7: We follow the same idea from [CK18, Theorem 2], which shows the relationship between the size of the hint S and the probability ε that can win the multi-user bridge-finding game. Say that a map τ is *good* if $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ wins the multi-user bridge-finding game with probability at least $\varepsilon/2$ on τ . That is, τ is *good* if we have

$$\Pr \left[\text{BridgeChal}_{\mathcal{A}_{\text{on}}, \text{str}_{\tau}}^{\tau, N}(k) = 1 \right] \geq \frac{\varepsilon}{2},$$

where the probability is taken over the selection of $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$, the random coins of \mathcal{A}_{on} , and the random coins used by the challenger in the bridge game (the hint $\text{str}_{\tau} = \mathcal{A}_{\text{pre}}^{\text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger).

Let $\mathcal{T} = \{\tau_1, \tau_2, \dots\}$ be the set of good labeling maps. Then by a standard averaging argument ([?, Lemma A.12]), an $(\varepsilon/2)$ -fraction of injective mappings from \mathbb{Z}_p to \mathbb{G} are good. One could also observe that the number of injective mappings from \mathbb{Z}_p to \mathbb{G} is $|\mathbb{G}|!/(|\mathbb{G}| - p)!$, which implies that $|\mathcal{T}| \geq (\varepsilon/2) \cdot |\mathbb{G}|!/(|\mathbb{G}| - p)!$.

We introduce [Lemma 3](#) which has been adapted from Abadi et al. [?] that an average-case multi-user bridge-finding game implies a worst-case multi-user bridge-finding game as shown below.

Lemma 3. *For a prime p , let $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ be a pair of generic algorithms for \mathbb{Z}_p on \mathbb{G} with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that \mathcal{A}_{pre} outputs an S -bit hint str_τ and \mathcal{A}_{on} makes at most q_{on} generic group oracle queries. Then there exists a generic algorithm \mathcal{A}'_{on} that makes at most $q_{\text{on}} + 2N \log p + N$ generic group oracle queries, and, for every $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$, if $\Pr_{\vec{x}, \mathcal{A}_{\text{on}}}[\text{BridgeChal}_{\mathcal{A}_{\text{on}}, \text{str}_\tau}^{\tau, N}(k, \vec{x}) = 1] \geq \varepsilon$, then for every $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$, $\Pr_{\mathcal{A}'_{\text{on}}}[\text{BridgeChal}_{\mathcal{A}'_{\text{on}}, \text{str}_\tau}^{\tau, N}(k, \vec{x}) = 1] \geq \varepsilon$.*

Proof. Algorithm $\mathcal{A}'_{\text{on}}(\text{str}_\tau, \tau(x_1), \dots, \tau(x_N))$ executes the following:

- Sample a random $(r_1, \dots, r_N) \leftarrow_{\$} \mathbb{Z}_p^N$ and computes $\text{Mult}(\tau(x_i), \tau(r_i))$ for each $i \in [N]$, using at most $2N \log p + N$ group operations, since it would need at most $2 \log p$ group operations to compute $\tau(r_i)$ given each r_i for $i \in [N]$ and we have extra one operation of Mult to compute $\text{Mult}(\tau(x_i), \tau(r_i))$ for each $i \in [N]$.
- Run $\mathcal{A}_{\text{on}}(\text{str}_\tau, \tau(x_1 + r_1), \dots, \tau(x_N + r_N))$, which plays a multi-user bridge-finding game $\text{BridgeChal}_{\mathcal{A}_{\text{on}}, \text{str}_\tau}^{\tau, N}(k, \vec{x}')$, where $\vec{x}' = (x_1 + r_1, \dots, x_N + r_N)$.
- When \mathcal{A}_{on} outputs 1, then algorithm \mathcal{A}'_{on} also outputs 1, and vice versa.

Observe that during the execution of \mathcal{A}'_{on} , we run \mathcal{A}_{on} on $(\tau(x_1 + r_1), \dots, \tau(x_N + r_N))$, which is the image of a uniformly random point in \mathbb{Z}_p^N . Since we assume that the bridge-finding game with \mathcal{A}_{on} succeeds with probability at least ε over the random selection of $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$ and its coins, we can conclude that the bridge-finding game with \mathcal{A}'_{on} also succeeds with probability at least ε only over the selection of its coins. \square

Now we are back to the proof of [Theorem 7](#). [Lemma 3](#) implies that there exists a pair of generic algorithms $(\mathcal{A}_{\text{pre}}, \mathcal{A}'_{\text{on}})$ such that for every $\tau \in \mathcal{T}$ and every $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$, \mathcal{A}'_{on} making at most $q_{\mathbb{G}}^{\text{on}'} = q_{\mathbb{G}}^{\text{on}} + 2N \log p + N$ generic group oracle queries and we have $\Pr_{\mathcal{A}'_{\text{on}}}[\text{BridgeChal}_{\mathcal{A}'_{\text{on}}, \text{str}_\tau}^{\tau, N}(k, \vec{x}) = 1] \geq \varepsilon/2$.

Now from [Lemma 2](#), we can use $(\mathcal{A}_{\text{pre}}, \mathcal{A}'_{\text{on}})$ to compress any mapping $\tau \in \mathcal{T}$ to a binary string of length at most

$$\log \frac{|\mathbb{G}|!}{(|\mathbb{G}| - p)!} + S + 1 - \frac{(\varepsilon/2)p}{6q_{\mathbb{G}}^{\text{on}'}(q_{\mathbb{G}}^{\text{on}'} + N)(N \log p + 1)},$$

where the encoding scheme works with probability at least $1/2$. The incompressibility argument⁸ of De et al. [?] says that this length must be at least $\log |\mathcal{T}| - \log 2$. Hence, we have that

$$\log \frac{|\mathbb{G}|!}{(|\mathbb{G}| - p)!} + S + 1 - \frac{(\varepsilon/2)p}{6q_{\mathbb{G}}^{\text{on}'}(q_{\mathbb{G}}^{\text{on}'} + N)(N \log p + 1)} \geq \log \frac{|\mathbb{G}|!}{(|\mathbb{G}| - p)!} - \log \frac{4}{\varepsilon},$$

⁸ [?, Fact 8.1] says that “Suppose there is a randomized encoding procedure $\text{Enc} : \{0, 1\}^n \times \{0, 1\}^r \rightarrow \{0, 1\}^m$ and a decoding procedure $\text{Dec} : \{0, 1\}^m \times \{0, 1\}^r \rightarrow \{0, 1\}^n$ such that $\Pr_{r \in U_r}[\text{Dec}(\text{Enc}(x, r), r) = x] \geq \delta$. Then $m \geq n - \log(1/\delta)$.” Since $\{0, 1\}^n$ corresponds to the set of mapping \mathcal{T} , setting $\delta = 1/2$, we get the result.

which implies that

$$S \geq \frac{\varepsilon p}{6(q_{\mathbb{G}}^{\text{on}} + 2N \log p + N)(q_{\mathbb{G}}^{\text{on}} + 2N \log p + 2N)(N \log p + 1)} - \log \frac{8}{\varepsilon}.$$

Without loss of generality, we may assume that $\varepsilon \geq 1/p$, since the probability to find a BRIDGE^N instance is not lower than solving 1-out-of- N discrete log problem (see [Corollary 1](#)), which is also not lower than $1/p$ when guessing them randomly. Hence, we have that $\log \frac{8}{\varepsilon} \leq \log(8p)$ and we get

$$6(S + \log(8p))(q_{\mathbb{G}}^{\text{on}} + 2N \log p + N)(q_{\mathbb{G}}^{\text{on}} + 2N \log p + 2N)(N \log p + 1) \geq \varepsilon p,$$

which implies that

$$\begin{aligned} \varepsilon &\leq \frac{6}{p}(S + \log(8p))(q_{\mathbb{G}}^{\text{on}} + N + 2N \log p)(q_{\mathbb{G}}^{\text{on}} + 2N + 2N \log p)(N \log p + 1) \\ &= \tilde{O}\left(\frac{SN(q_{\mathbb{G}}^{\text{on}} + N)(q_{\mathbb{G}}^{\text{on}} + 2N)}{p}\right), \end{aligned}$$

i.e., if the size of the hint is bounded then the probability of winning the 1-out-of- N generic bridge-finding game is also bounded correspondingly. In particular, if $q_{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then we observe that

$$\begin{aligned} \varepsilon &\leq \frac{6}{p}(S + \log(8p))(q_{\mathbb{G}}^{\text{on}} + N + 2N \log p)(q_{\mathbb{G}}^{\text{on}} + 2N + 2N \log p)(N \log p + 1) \\ &\leq \frac{6}{p}(1.1S)(1.1q_{\mathbb{G}}^{\text{on}})^2(1.5N \log p) \leq \frac{12SN(q_{\mathbb{G}}^{\text{on}})^2 \log p}{p}, \end{aligned}$$

where we have $N \log p + 1 \leq 1.5N \log p$ because $p > 2^{2k}$ implies $N \log p > 2kN \geq 2k$ and $N \log p + 1 \leq (1 + 1/(2k))N \log p \leq 1.5N \log p$. \square

Reminder of Lemma 2. Let \mathbb{G} be the set of binary strings of length ℓ such that $2^\ell \geq p$ for a prime p . Let $\mathcal{T} = \{\tau_1, \tau_2, \dots\}$ be a subset of the labeling functions from \mathbb{Z}_p to \mathbb{G} . Let $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ be a pair of generic algorithms for \mathbb{Z}_p on \mathbb{G} such that for every $\tau \in \mathcal{T}$ and every $\vec{x} = (x_1, \dots, x_N) \in \mathbb{Z}_p^N$, \mathcal{A}_{pre} outputs an S -bit advice string, \mathcal{A}_{on} makes at most q_{on} oracle queries, and $(\mathcal{A}_{\text{pre}}, \mathcal{A}_{\text{on}})$ satisfy $\Pr_{\mathcal{A}_{\text{on}}}[\text{BridgeChal}_{\mathcal{A}_{\text{on}}, \text{str}_\tau}^{\tau, N}(k, \vec{x}) = 1] \geq \varepsilon$, where $\text{str}_\tau = \mathcal{A}_{\text{pre}}^{\text{GO}}(\tau(1))$. Then, there exists a randomized encoding scheme that compresses elements of \mathcal{T} to bitstrings of length at most

$$\log \frac{|\mathbb{G}|!}{(|\mathbb{G}| - p)!} + S + 1 - \frac{\varepsilon p}{6q_{\text{on}}(q_{\text{on}} + N)(N \log p + 1)},$$

and succeeds with probability at least $1/2$.

Proof of Lemma 2: The proof works largely the same as it from [\[CK18\]](#) except for some modifications. In the proof of [\[CK18\]](#), they build an encoding and a decoding routine to construct a randomized encoding scheme. The encoding

routine takes as input an encoding function $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ and a random binary string r , and outputs a compression of τ . The decoding routine reverts this operation; it takes as input a compressed representation of τ and the same random string r as used in the encoding routine, and outputs the original τ . One modification from the proof of [CK18] is that in the encoding routine, we also need to handle $\text{Inv}(\cdot)$ while [CK18] only considers $\text{Mult}(\cdot, \cdot)$. Another difference is that the multi-user bridge-finding game $\text{BridgeChal}_{\mathcal{A}_{\text{on}}}^{\tau, N}(k, \vec{x})$ that \mathcal{A}_{on} plays does not aim for finding discrete-log solutions but for finding a BRIDGE^N instance.

Encoding Routine:

Input: an encoding function $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$, and parameters $d, R \in \mathbb{Z}^+$, where \mathbb{Z}^+ denotes the set of positive integers.

- (1) Compute and write the S -bit hint $\text{str}_\tau \leftarrow \mathcal{A}_{\text{pre}}(\mathfrak{g})$ into the encoding.
- (2) Write the image of τ using $\log \binom{|\mathbb{G}|}{p}$ bits into the encoding.
- (3) Initialize **Table** to an empty list. It will store the pairs $(y, \tau(y))$, i.e., an element in the image of τ and its discrete log value.
- (4) Repeat the following d times in total:
 - (a) Choose the first N strings in the lexicographical order of the image of τ that are not in the table. Call these strings $\tau(x_1), \dots, \tau(x_N)$ and add the pairs $(X_1, \tau(x_1)), \dots, (X_N, \tau(x_N))$ to the table, where for each $i \in [N]$, X_i is the indeterminate that represents discrete log value x_i of $\tau(x_i)$ that the decoder does not yet know.
 - (b) Run $\mathcal{A}_{\text{on}}(\text{str}_\tau, \tau(x_1), \dots, \tau(x_N))$ up to R times using independent randomness from the encoder's random string in each run. Write the index $r^* \in [R]$ of the successful execution (i.e., it finds a BRIDGE^N instance) into the encoding using $\log R$ bits. If \mathcal{A}_{on} fails on all R execution, then return \perp and abort the entire routine.
 - (c) Write a placeholder of $\log q_{\text{on}}$ zeros in the encoding, and it will be overwritten by the actual number of queries until it finds a BRIDGE^N instance during the next step.
 - (d) Rerun $\mathcal{A}_{\text{on}}(\text{str}_\tau, \tau(x_1), \dots, \tau(x_N))$ using the r^{*th} random tape. It processes each generic group oracle query (see **Handling Mult**(η_1, η_2) **and Handling Inv**(η) as described below) and as soon as it finds a BRIDGE^N instance it will mark the actual number of total queries $q^* \leq q_{\text{on}}$ into its placeholder. Note that it will count the number of queries altogether whenever it queries $\text{Mult}(\cdot, \cdot)$ or $\text{Inv}(\cdot)$.
 - (e) For the remaining indeterminates that are not yet resolved, simply replace them with the discrete log values (naïve encoding), since the encoder knows all of τ . If this is the i^{th} entry in **Table**, then it requires $\log(p - i + 1)$ bits since all the entries that were already in the table are constants.
 - (f) Write the remaining values that are not yet in the table to the encoding in lexicographical order.

Handling $\text{Mult}(\eta_1, \eta_2)$:

- (1) If either of η_1 or η_2 is not in the image of τ , reply \perp and continue to the next query.
- (2) If either of η_1 or η_2 is not in the table, then this is a *fresh* query input. For each such argument $\mathfrak{s} \in \{\eta_1, \eta_2\}$, add the pair $(\text{DLog}(\mathfrak{s}), \mathfrak{s})$ to the table and write \mathfrak{s} to the encoding using $\log(p-l)$ bits, where l is the number of labels already in the table.
- (3) Otherwise, look up the tuples $(f_1, \eta_1), (f_2, \eta_2)$ in the table where $f_1 = f_1(X_1, \dots, X_N)$ and $f_2 = f_2(X_1, \dots, X_N)$ are linear polynomials of N indeterminates X_1, \dots, X_N , and compute $f_1 + f_2$.
 - (a) If $(f_1 + f_2, \text{Mult}(\eta_1, \eta_2))$ is already in the table, simply reply with $\text{Mult}(\eta_1, \eta_2)$.
 - (b) If $(f_1 + f_2, \text{Mult}(\eta_1, \eta_2))$ is not in the table, then add $\text{Mult}(\eta_1, \eta_2)$ to the encoding and reply with $\text{Mult}(\eta_1, \eta_2)$. Writing $\text{Mult}(\eta_1, \eta_2)$ into the encoding requires $\log(p-l)$ bits, where l is the number of labels already in the table.
 - (c) If $\text{Mult}(\eta_1, \eta_2)$ is in the table but its corresponding discrete log value in the table is a linear polynomial $f = f(X_1, \dots, X_N)$ such that f is not identical to $f_1 + f_2$ (i.e., the coefficients for X_i 's are not all the same), then encode the reply to this query as a pointer to the table entry $(f, \text{Mult}(\eta_1, \eta_2))$ and add this pointer to the encoding. Then use the equation $f = f_1 + f_2$ to derive an equation $X_j = g(X_1, \dots, \bar{X}_j, \dots, X_N)$ for some j and a linear polynomial g of X_1, \dots, X_N except for X_j , and replace X_j by $g(X_1, \dots, \bar{X}_j, \dots, X_N)$ in the table.
In this case, we successfully found a BRIDGE^N instance, so stop the execution of the algorithm and indicate the *early stop* by writing the actual number of queries $q^* \leq q_{\text{on}}$ into its placeholder above.

Handling $\text{Inv}(\eta)$:

- (1) If η is not in the image of τ , reply \perp and continue to the next query.
- (2) If η is not in the table, then this is an *fresh* query input. Add the pair $(\text{DLog}(\eta), \eta)$ to the table.
- (3) Otherwise, look up the tuple (f, η) in the table where $f = f(X_1, \dots, X_N)$ is a linear polynomial of N indeterminates X_1, \dots, X_N , and compute $-f$.
 - (a) If $(-f, \text{Inv}(\eta))$ is already in the table, simply reply with $\text{Inv}(\eta)$.
 - (b) If $(-f, \text{Inv}(\eta))$ is not in the table, then add $\text{Inv}(\eta)$ to the encoding and reply with $\text{Inv}(\eta)$. Writing $\text{Inv}(\eta)$ into the encoding requires $\log(p-l)$ bits, where l is the number of labels already in the list.
 - (c) If $\text{Inv}(\eta)$ is in the table but its corresponding discrete log value in the table is a linear polynomial $\tilde{f} = \tilde{f}(X_1, \dots, X_N)$ such that f is

not identical to \tilde{f} (i.e., the coefficients for X_i 's are not all the same), then encode the reply to this query as a pointer to the table entry $(\tilde{f}, \text{Inv}(\eta))$ and add this pointer to the encoding.

Then use the equation $f = -\tilde{f}$ to derive an equation $X_j = g(X_1, \dots, \overline{X}_j, \dots, X_N)$ for some j and a linear polynomial g of X_1, \dots, X_N except for X_j , and replace X_j by $g(X_1, \dots, \overline{X}_j, \dots, X_N)$ in the table.

In this case, we successfully found a BRIDGE^N instance, so stop the execution of the algorithm and indicate the *early stop* by writing the actual number of queries $q^* \leq q_{\text{on}}$ into its placeholder above.

Decoding Routine:

The decoding routine is given the encoded string and recovers τ as follows:

- (1) Extract str_τ from the first S bits of the encoding.
- (2) Extract the image of τ from the next $\log \binom{|G|}{p}$ bits from the encoding.
- (3) Initialize **Table** (input/output pairs for τ) as an empty list.
- (4) Extract $\tau(1)$ from the encoding and add $(1, \tau(1))$ to table.
- (5) Repeat the following d times in total:
 - (a) Choose the first N strings in the lexicographical order of the image of τ that are not in **Table**. Call these strings η_1, \dots, η_N and add the pairs $(X_1, \eta_1), \dots, (X_N, \eta_N)$ to the table, where X_1, \dots, X_N denotes the indeterminates that represent the discrete log of η_i 's for $i = 1, \dots, N$.
 - (b) Decode $r^* \in [R]$ after reading $\log R$ bits from the encoding.
 - (c) Decode $q^* \in [q_{\text{on}}]$ after reading $\log q_{\text{on}}$ bits from the encoding.
 - (d) Run $\mathcal{A}_{\text{on}}(\text{str}_\tau, \eta_1, \dots, \eta_N)$ for q^* queries using the r^{*th} random tape allocated for this instance of \mathcal{A}_{on} .
 - i. If \mathcal{A}_{on} makes the query $\text{Mult}(\mathfrak{s}_1, \mathfrak{s}_2)$:
 - A. If either \mathfrak{s}_1 or \mathfrak{s}_2 is not in the image of τ , reply \perp .
 - B. If either \mathfrak{s}_1 or \mathfrak{s}_2 is in the image of τ , but not in the table, then add entries (S_i, \mathfrak{s}_1) and (S_j, \mathfrak{s}_2) to the table, where i and j are the smallest integers such that S_i and S_j are the fresh indeterminates in the table.
 - C. Search for the table and find the pairs (f_1, \mathfrak{s}_1) and (f_2, \mathfrak{s}_2) , and compute the linear polynomial $f_1 + f_2$.
 - (1) If $(f_1 + f_2, \mathfrak{w})$ is already in the table, then reply with \mathfrak{w} and continue to the next query.
 - (2) If there is no such \mathfrak{w} in the table, and if this is not the last (q^{*th}) query of this execution, then read $\log(p - l)$ bits from the encoding, where l is the number of labels in the table, and decode the bits to output \mathfrak{w} . Reply with \mathfrak{w} , and add $(f_1 + f_2, \mathfrak{w})$ to the table and continue to the next query.

- (3) If there is no such \mathfrak{w} in the table, and if this is the last (q^{*th}) query, then we know that we have a bridge instance here. Thus, read a $\log |\text{Table}|$ -bit pointer from the encoding and search for the entry in the table with that pointer, say (f, \mathfrak{w}) . Solve the equation $f = f_1 + f_2 \pmod p$ for the first indeterminate S in the equation. Replace every S in the polynomials in the table with the solution of the equation so that there is no longer S in the entire table.
- ii. If \mathcal{A}_{on} makes the query $\text{Inv}(\mathfrak{s})$:
- A. If \mathfrak{s} is not in the image of τ , reply \perp .
 - B. If \mathfrak{s} is in the image of τ , but not in the table, then add an entry (S_i, \mathfrak{s}) to the table, where i is the smallest integer such that S_i is the fresh indeterminate in the table.
 - C. Search for the table and find the pair (f, \mathfrak{s}) , and compute the linear polynomial $-f$.
 - (1) If $(-f, \mathfrak{w})$ is already in the table, then reply with \mathfrak{w} and continue to the next query.
 - (2) If there is no such \mathfrak{w} in the table, and if this is not the last (q^{*th}) query of this execution, then read $\log(p-l)$ bits from the encoding, where l is the number of labels in the table, and decode the bits to output \mathfrak{w} . Now reply with \mathfrak{w} , and add $(-f, \mathfrak{w})$ to the table and continue to the next query.
 - (3) If there is no such \mathfrak{w} in the table, and if this is the last (q^{*th}) query, then we know that we have a bridge instance here. Thus, read a $\log |\text{Table}|$ -bit pointer from the encoding and search for the entry in the table with this pointer, say $(\tilde{f}, \mathfrak{w})$. Solve the equation $\tilde{f} = -f \pmod p$ for the first indeterminate S in the equation. Replace every S in the polynomials in the table with the solution of the equation so that there is no longer S in the entire table.
- iii. Read the values of all the remaining indeterminates as they appear in the table.

The total encoding length will be calculated as follows:

- the hint τ (S bits),
- the encoding of the image of τ ($\log \binom{|\mathbb{G}|}{p}$ bits),
- the index r^* which denotes the successful execution of \mathcal{A}_{on} ($d \log R$ bits),
- the actual number of total queries q^* ($d \log q_{\text{on}}$ bits),
- for the l^{th} entry that is added to the table ($0 \leq l < |\text{Table}|$), if the entry
 - corresponds to an indeterminate that has been resolved by finding a BRIDGE^N instance (at most $\log |\text{Table}|$ bits),
 - otherwise, $\log(p-l)$ bits, and
- the remaining discrete log values, encoded using $\log(p-l)$ bits each, where $l \in \{|\text{Table}|, \dots, p-1\}$.

Note that at the end of each of the d executions of \mathcal{A}_{on} , we either learn the discrete log of one label $\tau(\cdot)$, or, we find a bridge instance in the table. In this case, we get $\log(p-l)$ bits of information on τ at a cost of at most $\log R + \log q_{\text{on}} + \log |\text{Table}|$ bits. Since each execution of \mathcal{A}_{on} adds at most $3q_{\text{on}} + N$ rows to the table (N inputs, q_{on} query replies, and at most $2q_{\text{on}}$ fresh query inputs), we observe that $|\text{Table}| \leq d(3q_{\text{on}} + N)$. Since $l \leq |\text{Table}|$, the net profit is lower bounded by

$$\log \frac{p-l}{\log R + \log q_{\text{on}} + \log |\text{Table}|} \geq \log \frac{p-|\text{Table}|}{Rq_{\text{on}}|\text{Table}|} \geq \log \frac{p-d(3q_{\text{on}}+N)}{Rdq_{\text{on}}(3q_{\text{on}}+N)}.$$

We further observe that $\log \frac{p-d(3q_{\text{on}}+N)}{Rdq_{\text{on}}(3q_{\text{on}}+N)} \geq 1$ for $d = \lfloor p/((2Rq_{\text{on}}+1)(3q_{\text{on}}+N)) \rfloor$. Thus, with this value of d , we have that this net profit becomes at least 1 bit for each of the d executions of \mathcal{A}_{on} . Hence, the total bitlength of the encoding is at most

$$\begin{aligned} S + \log \binom{|\mathbb{G}|}{p} + \sum_{l=0}^p \log(p-l) - d &= \log \frac{|\mathbb{G}|!}{(|\mathbb{G}|-p)!} + S - d \\ &\leq \log \frac{|\mathbb{G}|!}{(|\mathbb{G}|-p)!} + S - \frac{p}{(2Rq_{\text{on}}+1)(3q_{\text{on}}+N)} + 1 \\ &\leq \log \frac{|\mathbb{G}|!}{(|\mathbb{G}|-p)!} + S - \frac{p}{6Rq_{\text{on}}(q_{\text{on}}+N)} + 1. \end{aligned}$$

We want that this randomized encoding routine succeeds with probability at least $1/2$. In the encoding routine, we run \mathcal{A}_{on} up to R times and see if there is any successful execution. Since each execution fails with probability at most ε , if we choose $R = (1 + N \log p)/\varepsilon$, then such R executions all fail with the probability at most

$$(1-\varepsilon)^R \leq e^{-\varepsilon R} \leq 2^{-\varepsilon R} \leq 2^{-1-N \log p} \leq 2^{-1-\log \binom{p}{N}} \leq \frac{1}{2} \cdot \binom{p}{N}^{-1},$$

since $p^N \geq \binom{p}{N}$. Union bounding over all $\binom{p}{N}$ different inputs on \mathcal{A}_{on} , the probability that the encoding routine fails is at most $1/2$. Hence, with this value of R , the encoding length is at most

$$\log \frac{|\mathbb{G}|!}{(|\mathbb{G}|-p)!} + S + 1 - \frac{\varepsilon p}{6q_{\text{on}}(q_{\text{on}}+N)(N \log p + 1)},$$

which completes the proof. \square

Reminder of Theorem 8. Let $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ be a key-prefixed Schnorr signature scheme and $p > 2^{2k}$ be a prime number. Let $N \in \mathbb{N}$ be a parameter and $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that $\mathcal{A}_{\text{sig}}^{\text{pre}}$ makes at most $q_{\text{H}}^{\text{pre}}$ queries to the random oracle $\text{H} : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$ and outputs an S -bit hint $\text{str}_{\tau, \text{H}}$, and $\mathcal{A}_{\text{sig}}^{\text{on}}$ makes at most $q_{\text{G}}^{\text{on}} := q_{\text{G}}^{\text{on}}(k)$ queries

to the generic group oracles and at most $q_{\mathbb{H}}^{\text{on}}$ queries to the random oracle. Then $\Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, \mathbb{H}}}, \Pi}^{\tau, N}(k) = 1 \right] \leq \varepsilon$, with

$$\varepsilon = \tilde{\mathcal{O}} \left(\frac{SN(q_{\mathbb{G}}^{\text{on}} + N)(q_{\mathbb{G}}^{\text{on}} + 2N)}{p} \right) + \frac{Nq_{\mathbb{H}}^{\text{pre}}q_{\mathbb{S}}^{\text{on}}}{p^2} + \frac{q_{\mathbb{S}}^{\text{on}}(q_{\mathbb{S}}^{\text{on}} + q_{\mathbb{H}}^{\text{on}})}{p} + \frac{4(q_{\mathbb{H}}^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p},$$

where $q_{\mathbb{S}}^{\text{on}}$ denotes the number of queries to the signing oracle and the randomness is taken over the selection of τ and the random coins of $\mathcal{A}_{\text{sig}}^{\text{on}}$ (the hint $\text{str}_{\tau, \mathbb{H}} = \mathcal{A}_{\text{sig}}^{\text{pre}, \text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger). In particular, if $q_{\mathbb{G}}^{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then

$$\varepsilon = \frac{12SN(q_{\mathbb{G}}^{\text{on}})^2 \log p}{p} + \frac{Nq_{\mathbb{H}}^{\text{pre}}q_{\mathbb{S}}^{\text{on}}}{p^2} + \frac{q_{\mathbb{S}}^{\text{on}}(q_{\mathbb{S}}^{\text{on}} + q_{\mathbb{H}}^{\text{on}})}{p} + \frac{4(q_{\mathbb{H}}^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}.$$

Proof of Theorem 8: Given a generic adversary with preprocessing $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ attacking Schnorr signature scheme, we construct the following efficient generic algorithm with preprocessing $(\mathcal{A}_{\text{bridge}}^{\text{pre}}, \mathcal{A}_{\text{bridge}}^{\text{on}})$ which tries to succeed in the 1-out-of- N generic BRIDGE ^{N} -finding game $\text{BridgeChal}_{\mathcal{A}_{\text{bridge}, \text{str}_{\tau, \mathbb{H}}}^{\tau, N}}(k, \vec{x})$:

Algorithm $(\mathcal{A}_{\text{bridge}}^{\text{pre}}, \mathcal{A}_{\text{bridge}}^{\text{on}})$:

The algorithm is given $p, \mathbf{g} = \tau(1), \text{pk}_i = \tau(x_i), 1 \leq i \leq N$ as input.

1. $\mathcal{A}_{\text{bridge}}^{\text{pre}}$ simply runs $\mathcal{A}_{\text{sig}}^{\text{pre}}$ to generate an S -bit hint $\text{str}_{\tau, \mathbb{H}}$.
2. $\mathcal{A}_{\text{bridge}}^{\text{on}}$ initializes the set $\mathbf{H}_{\text{resp}} = \{\}$ which stores the random oracle input/output pairs observed during online processing.
3. $\mathcal{A}_{\text{bridge}}^{\text{on}}$ takes the hint $\text{str}_{\tau, \mathbb{H}}$ and initializes the list $\mathcal{L} = \{(\tau(1), \vec{0}, 1), (\text{pk}_i, \hat{u}_i, 0) \text{ for } 1 \leq i \leq N\}$, and runs $\mathcal{A}_{\text{sig}}^{\text{on}}$ with the hint $\text{str}_{\tau, \mathbb{H}}$ and a number of access to the generic group oracles $\text{GO} = (\text{Mult}(\cdot, \cdot), \text{Inv}(\cdot), \text{DLog}(\cdot), \text{Sign}_i(\cdot))$ for $1 \leq i \leq N$, and the random oracle $\mathbb{H}(\cdot)$. We maintain the invariant that every output of a generic group query during the online phase appears in the list \mathcal{L} . We consider the following cases:
 - (a) Whenever $\mathcal{A}_{\text{sig}}^{\text{on}}$ submits a query w to the random oracle \mathbb{H} :
 - If there is a pair $(w, R) \in \mathbf{H}_{\text{resp}}$ for some string $R \in \{0, 1\}^{k_1}$ then return R .
 - Otherwise, select $R \leftarrow_{\$} \{0, 1\}^{k_1}$ and add (w, R) to \mathbf{H}_{resp} .
 - If w has the form $w = (\text{pk}_j \| \mathbf{a} \| m_i)$ where the value \mathbf{a} has not been observed previously (i.e., is not in the list \mathcal{L} then we query $b = \text{DLog}(\mathbf{a})$ and add $(\mathbf{a}, \vec{0}, b)$ to \mathcal{L} .
 - (b) Whenever $\mathcal{A}_{\text{sig}}^{\text{on}}$ submits a query \mathbf{a} to the generic group oracle $\text{Inv}(\cdot)$:
 - If \mathbf{a} is not in \mathcal{L} then we immediately query $b = \text{DLog}(\mathbf{a})$ and add $(\mathbf{a}, \vec{0}, b)$ to \mathcal{L} .
 - Otherwise, $(\mathbf{a}, \vec{a}, b) \in \mathcal{L}$ for some \vec{a} and b . Then we query $\text{Inv}(\mathbf{a}) = \tau(-\vec{a} \cdot \vec{x} - b)$, output the result and add $(\tau(-\vec{a} \cdot \vec{x} - b), -\vec{a}, -b)$ to \mathcal{L} .
 - (c) Whenever $\mathcal{A}_{\text{sig}}^{\text{on}}$ submits a query \mathbf{a}, \mathbf{b} to the generic group oracle $\text{Mult}(\cdot, \cdot)$:
 - If the element \mathbf{a} (resp. \mathbf{b}) is not in \mathcal{L} then query $b_0 = \text{DLog}(\mathbf{a})$ (resp. $b_1 = \text{DLog}(\mathbf{b})$) and add the element $(\mathbf{a}, \vec{0}, b_0)$ (resp. $(\mathbf{b}, \vec{0}, b_1)$) to \mathcal{L} .

- Otherwise both elements $(\mathbf{a}, \vec{a}_0, b_0), (\mathbf{b}, \vec{a}_1, b_1) \in \mathcal{L}$. Then we return $\text{Mult}(\mathbf{a}, \mathbf{b}) = \tau((\vec{a}_0 + \vec{a}_1) \cdot \vec{x} + b_0 + b_1)$ and add $(\tau((\vec{a}_0 + \vec{a}_1) \cdot \vec{x} + b_0 + b_1), \vec{a}_0 + \vec{a}_1, b_0 + b_1) \in \mathcal{L}$.
- (d) Whenever $\mathcal{A}_{\text{sig}}^{\text{on}}$ submits a query m_i to the signing oracle $\text{Sign}(x_j, \cdot)$:
 - The attacker tries to forge a signature without knowledge of the secret key x_j , relying on the ability to program the random oracle as follows:
 - i. Pick s_i, e_i randomly and compute $I_i = \tau(s_i - x_j e_i) = \text{Mult}(\mathfrak{s}_i, \text{Pow}(\text{pk}_j, -e_i))$ where $\mathfrak{s}_i = \text{Pow}(\mathfrak{g}, s_i)$.
 - ii. If $\text{H}(\text{pk}_j \| I_i \| m_i)$ has been previously queried, then return \perp .
 - iii. Otherwise, program $\text{H}(\text{pk}_j \| I_i \| m_i) := e_i$ and return $\sigma_i = (s_i, e_i)$.
 - We remark that a side effect of querying the Sign_j oracle is the addition of the tuples $(\tau(s_i), \vec{0}, s_i), (\tau(x_j e_i), e_i \hat{u}_i, 0)$ and $(\tau(s_i - x_j e_i), -e_i \hat{u}_i, s_i)$ to \mathcal{L} , since these values are computed using the generic group oracles Inv, Mult .
- (e) If at any point we have some string η such that $(\eta, \vec{a}, b) \in \mathcal{L}$ and $(\eta, \vec{c}, d) \in \mathcal{L}$ for $(\vec{a}, b) \neq (\vec{c}, d)$ then we can immediately have a BRIDGE^N instance $(\tau((\vec{a} - \vec{c}) \cdot \vec{x}), \vec{a} - \vec{c}, 0) \in \mathcal{L}$ and $(\tau(d - b), \vec{0}, d - b) \in \mathcal{L}$ since $\tau((\vec{a} - \vec{c}) \cdot \vec{x}) = \tau(d - b)$. Thus, without loss of generality, we can assume that each string η occurs at most once in the list \mathcal{L} .
- 4. After $\mathcal{A}_{\text{sig}}^{\text{on}}$ outputs $\sigma_{i^*} = (s_{i^*}, e_{i^*})$ and m_{i^*} , identify the index $i^* \in [N]$ such that $\text{Vf}(\text{pk}_{i^*}, m_{i^*}, \sigma_{i^*}) = 1$.
- 5. Compute $\tau(-e_{i^*} x_{i^*}) = \text{Inv}(\text{Pow}(\tau(x_{i^*}), e_{i^*}))$. This will ensure that the elements $(\tau(-e_{i^*} x_{i^*}), -e_{i^*} \hat{u}_{i^*}, 0)$ and $(\tau(e_{i^*} x_{i^*}), e_{i^*} \hat{u}_{i^*}, 0)$ are both added to \mathcal{L} .
- 6. Compute $\mathfrak{s}_{i^*} = \text{Pow}(\mathfrak{g}, s_{i^*})$ to ensure that $(\mathfrak{s}_{i^*}, \vec{0}, s_{i^*}) \in \mathcal{L}$.
- 7. Finally, compute $I_{i^*} = \text{Mult}(\mathfrak{s}_{i^*}, \tau(-e_{i^*} x_{i^*})) = \tau(s_{i^*} - x_{i^*} e_{i^*})$ which ensures that $(I_{i^*}, -e_{i^*} \hat{u}_{i^*}, s_{i^*}) \in \mathcal{L}$.

Analysis. We first remark that if the signature is valid then we must have $e_{i^*} = \text{H}(I_{i^*} \| m_{i^*})$ and $\text{DLog}(I_{i^*}) = s_{i^*} - x_{i^*} e_{i^*} = \vec{a} \cdot \vec{x} + b$.

We first analyze the probability that the attacker will outputs a valid signature forgery after when we use random oracle programming to simulate the honest signing oracles without the secret keys. We introduce two bad events FailtoSign and DetectProgramming such that the probability of a valid signature forgery is at least

$$\Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, \text{H}}}^{\tau, N}}(k) = 1 \right] - \Pr[\text{DetectProgramming}] - \Pr[\text{FailtoSign}],$$

where $\Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, \text{H}}}^{\tau, N}}(k) = 1]$ denotes the probability that the attacker is successful when playing with the real signing oracles.

FailtoSign is the event where our reduction outputs \perp in Step 3.(d) due signing oracle failure, i.e., during some query i , we generate I_i and find that the random oracle query $\text{H}(\text{pk}_j \| I_i \| m_i)$ was made previously during the *online* phase. As before, we can argue that

$$\Pr[\text{FailtoSign}] \leq \frac{q_{\text{S}}^{\text{on}}(q_{\text{S}}^{\text{on}} + q_{\text{H}}^{\text{on}})}{p},$$

i.e., see [Claim 1](#).

Intuitively, `DetectProgramming` denotes the event that the random oracle query $H(\text{pk}_j \| I_i \| m_i)$ was made in the *offline phase*. Now we have the following claim to upper bound the probability of the event `DetectProgramming`:

Claim 4. $\Pr[\text{DetectProgramming}] \leq \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2}$.

We push the proof of [Claim 4](#) right after the proof of our theorem. If the event `DetectProgramming` occurs, then it is possible that the signature the hint $\text{str}_{\tau, H}$ is somehow correlated with $H(\text{pk}_j \| I_i \| m_i)$, which could potentially allow the signature attacker to halt early because it guesses that we are programming the random oracle. Assuming that the event `DetectProgramming` does not occur, the attacker will not be able to distinguish between a programmed response and the signatures generated via the honest signing oracle.

We now upper bound the probability that the attacker outputs a valid signature forgery without the bridge event BRIDGE^N occurring. Unlike our previous analysis, we rely on a compression argument. `FailtoFind`(I_{i^*}) denotes the event that the signature is valid, but we find that I_{i^*} was not previously recorded in our list \mathcal{L} before we computed $\text{Mult}(\mathfrak{s}_{i^*}, \tau(-e_{i^*}x_{i^*}))$. Observe that if the attacker had previously queried $H(\text{pk}_{i^*} \| I_{i^*} \| m_{i^*})$ during the online phase, then I_{i^*} would have been added to \mathcal{L} already. Thus, in this case, a signature forgery allows us to predict the value of the random oracle $e_{i^*} = H(\text{pk}_{i^*} \| I_{i^*} \| m_{i^*})$. Intuitively, if `FailtoFind`(I_{i^*}) is too large then we can derive a contradiction by compressing the random oracle.

Similarly, we define `BadQuery` as the event that at some point for an element of the form $(I, -e\hat{u}_i, s) \in \mathcal{L}$, we make the random oracle query $H(\text{pk}_i \| I \| m)$ and receive the response $H(\text{pk}_i \| I \| m) = e$. Such a lucky query would allow the attacker to forge a signature without causing BRIDGE^N to occur. Intuitively, if such an event occurs we can also compress the random oracle, i.e., by using $\log q_H^{\text{on}} \leq k_1$ bits to encode the index of the lucky query.

Claim 5. $\Pr[\text{FailtoFind}(I_{i^*}) \cup \text{BadQuery}] \leq \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}$.

We also push the proof of [Claim 5](#) after the proof of our theorem. Now we have shown that

$$\begin{aligned} & \Pr \left[\text{BridgeChal}_{\mathcal{A}_{\text{bridge}, \text{str}_{\tau, H}}^{\text{on}}}^{\tau, N}(k) = 1 \right] \\ & \geq \Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, H}}^{\text{on}}}^{\tau, N}(k) = 1 \right] - \Pr[\text{DetectProgramming}] - \Pr[\text{FailtoSign}] \\ & \quad - \Pr[\text{FailtoFind}(I_{i^*})] - \Pr[\text{BadQuery}] \\ & \geq \Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, H}}^{\text{on}}}^{\tau, N}(k) = 1 \right] - \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2} - \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} - \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} - \frac{N^2(S + k_1)}{p}. \end{aligned}$$

Finally, we can apply [Theorem 7](#) to conclude that

$$\Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, H}}^{\text{on}}}^{\tau, N}(k) = 1 \right]$$

$$\leq \tilde{\mathcal{O}} \left(\frac{SN(q_G^{\text{on}} + N)(q_G^{\text{on}} + 2N)}{p} \right) + \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} + \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}.$$

Similar to [Theorem 7](#), if $q_G^{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then

$$\begin{aligned} & \Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig, str}_{\tau, \text{H}}}^{\tau, N}, \Pi}^{\tau, N}}(k) = 1 \right] \\ & \leq \frac{12SN(q_G^{\text{on}})^2 \log p}{p} + \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} + \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}, \end{aligned}$$

which completes the proof. \square

Reminder of Claim 4. $\Pr[\text{DetectProgramming}] \leq \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2}.$

Proof of Claim 4: Let $I_i = \text{Pow}(\mathbf{g}, r_i)$ be the value generated during the i^{th} query to the signing oracle. Let $B_{i,j}$ denote the bad event that the preprocessing attacker previously submitted a query of the form $\text{H}(\mathbf{pk}_j \| I_i \| \cdot)$. Since \mathbf{pk}_j and I_i are selected randomly, we have $\Pr[B_{i,j}] \leq q_H^{\text{pre}}/p^2$. We now define the event $\text{DetectProgramming} = \bigcup_{j \leq N, i \leq q_S} B_{i,j}$. Applying union bounds, we have

$$\Pr[\text{DetectProgramming}] = \Pr[\bigcup_{i,j} B_{i,j}] \leq \sum_{i,j} \Pr[B_{i,j}] \leq \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2}. \quad \square$$

Reminder of Claim 5. $\Pr[\text{FailtoFind}(I_{i^*}) \cup \text{BadQuery}] \leq \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}.$

Proof of Claim 5: Fixing τ and H , and for some parameter t (which we will determine later), we define $\varepsilon_{F, \tau, \text{H}}$ to be

$$\varepsilon_{F, \tau, \text{H}} := \min_{Y: |Y| \leq Nt} \Pr[\text{FailtoFind}(I_{i^*}) \cup \text{BadQuery}],$$

where the probability is taken over the random selection of $x_1, \dots, x_N \in \mathbb{Z}_p \setminus Y$ and the random coins of $\mathcal{A}_{\text{sig, str}_{\tau, \text{H}}}^{\text{on}}$. We can argue that removing $Nt \approx NS$ points does not impact this probability, i.e., $\varepsilon_{F, \tau, \text{H}} \approx \varepsilon'_{F, \tau, \text{H}}$ where $\varepsilon'_{F, \tau, \text{H}}$ is the same probability when $x_1, \dots, x_N \in \mathbb{Z}_p$ are chosen without any restrictions. In particular, $\varepsilon'_{F, \tau, \text{H}} \leq \varepsilon_{F, \tau, \text{H}} + N^2t/p$ where the term N^2t/p upper bounds the probability that sample some point $x_i \in Y$.

We now sample t uniformly random strings R_1, \dots, R_t , e.g., by querying the random oracle at $\mathcal{O}(t)$ fixed points. We repeat the following experiment t times to extract t input/output predictions, where each iteration $i \in [t]$ uses R_i as its random coin when sampling \vec{x}^i and q^i :

- (1) For $i \in [t]$, sample $\vec{x}^i = (x_1^i, \dots, x_N^i)$ subject to the restriction that $x_j^i \neq x_j^{i'}$ for any pair $(i, j) \neq (i', j')$.
- (2) Pick $q^i \in [0, q_H^{\text{on}}]$ uniformly at random.

- (3) If $q^i = 0$, then wait for the (attempted) forged signature $\sigma_i = (s_i, e_i)$ to be output, compute $I_i = \tau(s_i - e_i x_i^j)$, and output the prediction $\mathbf{H}(\tau(x_i^j) \| I_i \| m_i) = e_i$.
- (4) If instead $q^i > 0$, then we simulate $\mathcal{A}_{\text{sig, str}_{\tau, \mathbf{H}}}^{\text{on}}$ until the query q^i to the random oracle. At this point, the query has the form $\mathbf{H}(\text{pk}_j \| I \| m)$ for some I such that $(I, e \hat{u}_j, 0) \in \mathcal{L}^i$. In this case, we can extract e from \mathcal{L}^i and output the prediction $\mathbf{H}(\text{pk}_j \| I \| m) = e$.

During each iteration i , we output a correct input/output pair with probability at least $\varepsilon_{F, \tau, \mathbf{H}} / (q_{\mathbf{H}}^{\text{on}} + 1)$. Thus, the procedure above correctly outputs t input/output pairs with probability at least $(\varepsilon_{F, \tau, \mathbf{H}} / (q_{\mathbf{H}}^{\text{on}} + 1))^t$.

Let **SucceedExp** denotes an event that the experiment above correctly outputs t correct input/output pairs (without querying the random oracle at these t points) when τ and \mathbf{H} are picked randomly.

Now let $\varepsilon_F := \mathbb{E}_{\tau, \mathbf{H}}[\varepsilon_{F, \tau, \mathbf{H}}]$ and suppose that $\varepsilon_F > 4(q_{\mathbf{H}}^{\text{on}} + 1)2^{-k_1}$ for contradiction, then we can use Markov inequality to argue that

$$\begin{aligned} \Pr_{\tau, \mathbf{H}} \left[\varepsilon_{F, \tau, \mathbf{H}} > \frac{2(q_{\mathbf{H}}^{\text{on}} + 1)}{2^{k_1}} \right] &= 1 - \Pr_{\tau, \mathbf{H}} \left[\varepsilon_{F, \tau, \mathbf{H}} \leq \frac{2(q_{\mathbf{H}}^{\text{on}} + 1)}{2^{k_1}} \right] \\ &= 1 - \Pr_{\tau, \mathbf{H}} \left[1 - \varepsilon_{F, \tau, \mathbf{H}} \geq 1 - \frac{2(q_{\mathbf{H}}^{\text{on}} + 1)}{2^{k_1}} \right] \\ &\geq 1 - \frac{1 - \varepsilon_F}{1 - 2(q_{\mathbf{H}}^{\text{on}} + 1)2^{-k_1}} \\ &> \frac{\varepsilon_F/2}{1 - \varepsilon_F/2} \geq \frac{\varepsilon_F}{2}. \end{aligned}$$

Thus, if we pick τ and \mathbf{H} randomly and run the above procedure, we will succeed with probability at least

$$\Pr[\text{SucceedExp}] \geq \frac{\varepsilon_F}{2} \left(\frac{\varepsilon_{F, \tau, \mathbf{H}}}{q_{\mathbf{H}}^{\text{on}} + 1} \right)^t \geq \frac{\varepsilon_F}{2} 2^{(1-k_1)t} > 2(q_{\mathbf{H}}^{\text{on}} + 1)2^{-k_1} 2^{(1-k_1)t}.$$

However, on the other hand, [Lemma 4](#) below tells us that

$$\Pr[\text{SucceedExp}] \leq 2^{-k_1 t + S}.$$

Picking $t = S + k_1$, we derive a contradiction, since, with this value of t we have

$$2(q_{\mathbf{H}}^{\text{on}} + 1)2^{-k_1 S - k_1^2 + S} < \Pr[\text{SucceedExp}] \leq 2^{-k_1 S - k_1^2 + S}.$$

The contradiction comes from the assumption that $\varepsilon_F > 4(q_{\mathbf{H}}^{\text{on}} + 1)2^{-k_1}$. Thus, we have

$$\begin{aligned} \varepsilon'_F &:= \Pr[\text{FailtoFind}(I_{i^*}) \cup \text{BadQuery}] \\ &= \mathbb{E}_{\tau, \mathbf{H}}[\varepsilon'_{F, \tau, \mathbf{H}}] \\ &\leq \mathbb{E}_{\tau, \mathbf{H}} \left[\varepsilon_{F, \tau, \mathbf{H}} + \frac{N^2(S + k_1)}{p} \right] \end{aligned}$$

$$= \varepsilon_F + \frac{N^2(S + k_1)}{p} \leq \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}. \quad \square$$

Lemma 4 has appeared in prior work in various forms. To the best of our knowledge, the original usage is from [?] though our statement is closer to a form from [?]. Since we rephrase the lemma slightly, we include a proof of **Lemma 4** below for completeness.

Definition 4 (k -bit prediction game). Let B be a uniformly random bit string and let \mathcal{A} be an algorithm that receives a hint $h = f(B) \in H$ which may depend arbitrarily on B and can additionally query B at specific indices before outputting indices i_1, \dots, i_k and bits b_1, \dots, b_k . We say that \mathcal{A} wins the k -bit prediction game if for all $j \leq k$ we have $B[i_j] = b_j$ and \mathcal{A} did not previously query for $B[i_j]$.

Lemma 4 ([?,?]). Any attacker \mathcal{A} wins the k -bit prediction game with probability at most $|H|/2^k$.

Proof. Let $\text{Success}_{\mathcal{A}}$ denote the attacker's success probability when $h = f(B)$ depends on B . Suppose for contradiction that $\Pr[\text{Success}_{\mathcal{A}}] > |H|2^{-k}$, then consider the algorithm \mathcal{A}' which takes no hint, samples $h \in H$ uniformly, and simulates \mathcal{A} with hint h . We have

$$\Pr[\text{Success}_{\mathcal{A}'}] \geq \Pr[h = f(B)] \Pr[\text{Success}_{\mathcal{A}}] > \frac{1}{|H|} |H|2^{-k} = 2^{-k}.$$

To obtain a contradiction, we observe that \mathcal{A}' succeeds with probability at most 2^{-k} . In particular, \mathcal{A}' starts with no hint and outputs i_1, \dots, i_k and bits b_1, \dots, b_k such that $B[i_j]$ has not been queried for all $j \leq k$. In this case, we can view each $B[i_j]$ as a uniformly random bit sampled *after* \mathcal{A}' outputs. Thus, $\Pr[\forall j \leq k, b_j = B[i_j]] = 2^{-k}$, which implies that $\Pr[\text{success}_{\mathcal{A}'}] \leq 2^{-k}$. \square

C.4 Missing Proofs from Section 6

Reminder of Theorem 9. The Chaum-Pedersen-FDH signature scheme is $\left(N, q_H, q_G, q_S, \mathcal{O}\left(\frac{q+N}{2^k}\right)\right)$ -MU-UF-CMA secure under the generic group model of prime order $p \approx 2^{2k}$ and the programmable random oracle model, where q denotes the total number of queries made by an adversary.

Proof Sketch of Theorem 9: Let Π be the Chaum-Pedersen-FDH signature scheme. We follow a similar reduction as in **Theorem 6** using the signing oracle without knowledge of the secret key x_j in **Figure 3** (top). Note that the corresponding public key is $\text{pk}_j = \tau(x_j)$. Whenever the attacker queries the random oracle of the form $H(h\|\eta\|\mathbf{a}\|\mathbf{b})$ we can ensure that η, \mathbf{a} and \mathbf{b} all appear in our list \mathcal{L} by using the oracle $\text{DLog}(\cdot)$ on any fresh input. Letting $e = H(h\|\eta\|\mathbf{a}\|\mathbf{b})$,

we then compute $\text{Pow}(\text{pk}_j, e)$ and $\text{Inv}(\text{Pow}(\text{pk}_j, e))$ to ensure that both the tuples $(\text{Pow}(\text{pk}_j, e), e\hat{u}_j, 0)$ and $(\text{Inv}(\text{Pow}(\text{pk}_j, e)), -e\hat{u}_j, 0)$ appear in \mathcal{L} .

We first consider the event FailtoSign that our reduction outputs \perp because of a signing oracle failure. Conditioning on this event not occurring the programmed signing oracle is equivalent to the real one. Thus, in our reduction the attacker successfully produces a forged signature with probability at least

$$\Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}}, \Pi}^{\tau, N}(k) = 1] - \Pr[\text{FailtoSign}].$$

When \mathcal{A}_{sig} outputs a (potential) forgery $\sigma_{i^*} = (\eta_{i^*}, \mathbf{a}_{i^*}, \mathbf{b}_{i^*}, s_{i^*})$ and the corresponding message m_{i^*} for some $i^* \in [N]$, we perform the following computations in hope of causing the bridge event BRIDGE^N (if it has not already):

- (1) Query the oracle $\text{DLog}(\cdot)$ on any fresh inputs $\eta_{i^*}, \mathbf{a}_{i^*}, \mathbf{b}_{i^*}$ (if applicable),
- (2) Compute $h_{i^*} = \text{H}'(\text{pk}_{i^*} \| m_{i^*}) = \text{Pow}(\mathbf{g}, \text{H}(\text{pk}_{i^*} \| m_{i^*}))$,
- (3) Compute $\mathbf{s}_{i^*} = \text{Pow}(\mathbf{g}, s_{i^*})$,
- (4) Compute the hash output $e_{i^*} = \text{H}(h_{i^*} \| \eta_{i^*} \| \mathbf{a}_{i^*} \| \mathbf{b}_{i^*})$, and
- (5) Compute $A_{i^*} = \text{Mult}(\mathbf{s}_{i^*}, \text{Inv}(\text{Pow}(\text{pk}_{i^*}, e_{i^*}))) = \tau(s_{i^*} - x_{i^*}e_{i^*})$.

This ensures that the following tuples appear in \mathcal{L} :

- $(\mathbf{s}_{i^*}, \vec{0}_j, s_{i^*})$,
- $(h_{i^*}, \vec{0}, \text{H}(\text{pk}_{i^*} \| m_{i^*}))$,
- $(\text{Pow}(\text{pk}_{i^*}, e_{i^*}), \hat{u}_{i^*}e_{i^*}, 0)$,
- $(\text{Pow}(\text{pk}_{i^*}, -e_{i^*}), -\hat{u}_{i^*}e_{i^*}, 0)$, and
- $(A_{i^*}, -e_{i^*}\hat{u}_{i^*}, s_{i^*})$.

We can now define the event $\text{FailtoFind}(A_{i^*})$ to be the event that A_{i^*} was not previously recorded in our list \mathcal{L} before we computed A_{i^*} in the last step. Similarly, we let BadQuery denote the event that the signature is valid but for the only prior tuple $(A_{i^*}, \vec{a}, b) \in \mathcal{L}$ we have $\vec{a} = -e_{i^*}\hat{u}_{i^*}$. If the signature is valid and neither event $\text{FailtoFind}(A_{i^*})$ nor BadQuery occurs then the bridge event BRIDGE^N must have occurred and we immediately win the game since $(A_{i^*}, \vec{a}, b) \in \mathcal{L}$, $(A_{i^*}, -e_{i^*}\hat{u}_{i^*}, s_{i^*}) \in \mathcal{L}$ and $\vec{a} \neq -e_{i^*}\hat{u}_{i^*}$.

Our analysis of the failure events FailtoSign , $\text{FailtoFind}(A_{i^*})$, and BadQuery is similar to the analysis in [Theorem 6](#) with a few minor difference. First, we consider the hash output $\text{H}(h_{i^*} \| \eta_{i^*} \| A_{i^*} \| \mathbf{b}_{i^*})$ instead of $\text{H}(I_{i^*} \| m_{i^*})$ since we are analyzing Chaum-Pedersen-FDH instead of short Schnorr Signatures. Additionally, since we are analyzing FDH when we analyze BadQuery we have to consider the possibility of a random oracle collisions. The attacker makes at most q_H random oracle queries and each query to a signing oracle generates at most 2 additional random oracle queries. Thus, the probability the attacker outputs a collision can be upper bounded as $\tilde{q}^2/2^{2k}$ since the domain of our random oracle where $\tilde{q} = q_H + q_G + 2q_S$. Similar to [Theorem 6](#) we can show that $\Pr[\text{FailtoSign}] \leq q_S(q_H + q_S)/p$, $\Pr[\text{FailtoFind}(A_{i^*})] \leq \frac{q_H + q_S}{p - |\mathcal{L}|} + \frac{1}{2^{2k}}$, and $\Pr[\text{BadQuery}] \leq q_H/2^{2k} + \tilde{q}^2/2^{2k}$, with a minor difference that the hash output is $2k$ bits instead of k bits. Since we have $|\mathcal{L}| \leq N + 3q_G + 1$ we conclude that

$$\Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}}, \Pi}^{\tau, N}(k) = 1]$$

$$\begin{aligned}
 &\leq \Pr[\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k) = 1] + \Pr[\text{FailtoSign}] + \Pr[\text{FailtoFind}(A_{i^*})] + \Pr[\text{BadQuery}] \\
 &\leq \frac{q_{\mathbb{G}}N + 3q_{\mathbb{G}}(q_{\mathbb{G}} + 1)/2}{p - (N + 3q_{\mathbb{G}} + 1)^2 - N} + \frac{q_{\mathbb{S}}(q_{\mathbb{H}} + q_{\mathbb{S}})}{p} + \frac{q_{\mathbb{H}} + q_{\mathbb{S}}}{p - (N + 3q_{\mathbb{G}} + 1)} + \frac{q_{\mathbb{H}} + \tilde{q}^2 + 1}{2^{2k}} \\
 &= \mathcal{O}\left(\frac{q + N}{2^k}\right),
 \end{aligned}$$

where the second inequality follows by applying [Theorem 5](#) to upper bound $\Pr[\text{BridgeChal}_{\mathcal{A}}^{\tau, N}(k) = 1]$. \square

Reminder of [Theorem 10](#). Let $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ be a key-prefixed Chaum-Pedersen-FDH signature scheme and $p > 2^{2k}$ be a prime number. Let $N \in \mathbb{N}$ be a parameter and $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that $\mathcal{A}_{\text{sig}}^{\text{pre}}$ makes at most $q_{\mathbb{H}}^{\text{pre}} < 2^{3k}$ queries to the random oracle $\mathbb{H} : \{0, 1\}^* \rightarrow \{0, 1\}^{2k}$ and outputs an S -bit hint $\text{str}_{\tau, \mathbb{H}}$, and $\mathcal{A}_{\text{sig}}^{\text{on}}$ makes at most $q_{\mathbb{G}}^{\text{on}} := q_{\mathbb{G}}^{\text{on}}(k)$ queries to the generic group oracles and at most $q_{\mathbb{H}}^{\text{on}}$ queries to the random oracle. Then $\Pr\left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, \mathbb{H}}}^{\tau, N}, \Pi}(k) = 1\right] \leq \varepsilon$, with

$$\varepsilon = \tilde{\mathcal{O}}\left(\frac{SN(q_{\mathbb{G}}^{\text{on}} + N)(q_{\mathbb{G}}^{\text{on}} + 2N)}{p}\right) + \frac{Nq_{\mathbb{H}}^{\text{pre}}q_{\mathbb{S}}^{\text{on}}}{p^2} + \frac{q_{\mathbb{S}}^{\text{on}}(q_{\mathbb{S}}^{\text{on}} + q_{\mathbb{H}}^{\text{on}})}{p} + \frac{4(q_{\mathbb{H}}^{\text{on}} + \tilde{q}_{\text{on}}^2 + 1)}{2^{2k}} + \frac{3N^2(S + 2k)}{2p},$$

where $q_{\mathbb{S}}^{\text{on}}$ denotes the number of queries to the signing oracle, $\tilde{q}_{\text{on}} = q_{\mathbb{H}}^{\text{on}} + 2q_{\mathbb{S}}^{\text{on}}$, and the randomness is taken over the selection of τ and the random coins of $\mathcal{A}_{\text{sig}}^{\text{on}}$ (the hint $\text{str}_{\tau, \mathbb{H}} = \mathcal{A}_{\text{sig}}^{\text{pre}, \text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger).

Proof Sketch of [Theorem 10](#): We follow a similar security reduction as in [Theorem 8](#), except for the differences as follows:

- When the attacker tries to forge a signature on the message m_i without knowledge of the secret key x_j , s/he follow [Figure 3](#) (top). Note that the corresponding public key is $\text{pk}_j = \tau(x_j)$ for the secret key x_j .
- Whenever the attacker submits a random oracle query of the form $\mathbb{H}(\text{pk}_j \| h \| \boldsymbol{\eta} \| \mathbf{a} \| \mathbf{b})$ we use the oracle $\text{DLog}(\cdot)$ to ensure that $\boldsymbol{\eta}$, \mathbf{a} and \mathbf{b} appear in our list \mathcal{L} if they do not already. Letting $e = \mathbb{H}(\text{pk}_j \| h \| \boldsymbol{\eta} \| \mathbf{a} \| \mathbf{b})$, we then compute $\text{Pow}(\text{pk}_j, e)$ and $\text{Inv}(\text{Pow}(\text{pk}_j, e))$ to ensure that both the tuples $(\text{Pow}(\text{pk}_j, e), e\hat{u}_j, 0)$ and $(\text{Inv}(\text{Pow}(\text{pk}_j, e)), -e\hat{u}_j, 0)$ appear in \mathcal{L} .
- After \mathcal{A}_{sig} outputs $\sigma_{i^*} = (\boldsymbol{\eta}_{i^*}, \mathbf{a}_{i^*}, \mathbf{b}_{i^*}, s_{i^*})$ and m_{i^*} for some $i^* \in [N]$ we perform the following computations in hope of causing the bridge event BRIDGE^N (if it has not already):
 - (1) Query the oracle $\text{DLog}(\cdot)$ on any fresh inputs $\boldsymbol{\eta}_{i^*}, \mathbf{a}_{i^*}, \mathbf{b}_{i^*}$ (if applicable),
 - (2) Compute $h_{i^*} = \mathbb{H}'(\text{pk}_{i^*} \| m_{i^*}) = \text{Pow}(\mathbf{g}, \mathbb{H}(\text{pk}_{i^*} \| m_{i^*}))$,
 - (3) Compute $\mathbf{s}_{i^*} = \text{Pow}(\mathbf{g}, s_{i^*})$,
 - (4) Compute the hash output $e_{i^*} = \mathbb{H}(\text{pk}_{i^*} \| h_{i^*} \| \boldsymbol{\eta}_{i^*} \| \mathbf{a}_{i^*} \| \mathbf{b}_{i^*})$, and
 - (5) Compute $A_{i^*} = \text{Mult}(\mathbf{s}_{i^*}, \text{Inv}(\text{Pow}(\text{pk}_{i^*}, e_{i^*}))) = \tau(s_{i^*} - x_{i^*}e_{i^*})$.

This ensures that the following tuples appear in \mathcal{L} :

$$\circ (\mathbf{s}_{i^*}, \mathbf{0}, s_{i^*}),$$

- $(h_{i^*}, \vec{0}, H(\text{pk}_{i^*} \| m_{i^*}))$,
- $(\text{Pow}(\text{pk}_{i^*}, e_{i^*}), \hat{u}_{i^*} e_{i^*}, 0)$,
- $(\text{Pow}(\text{pk}_{i^*}, -e_{i^*}), -\hat{u}_{i^*} e_{i^*}, 0)$, and
- $(A_{i^*}, -e_{i^*} \hat{u}_{i^*}, s_{i^*})$.

To analyze the probability that the bridge event occurs we first analyze the following possibilities for the reduction to terminate early with the output \perp . Everything goes the same with the analysis in [Theorem 8](#) except for the following:

- (1) Consider the i th query to the signing oracle and suppose we ask for a signature for m_i under public key pk_j . Let $B_{i,j}$ denote the bad event that the pre-processing attacker previously submitted a query of the form $H(\text{pk}_j \| H'(\text{pk}_j \| m_i) \| \eta_i \| \mathbf{a}_i \| \mathbf{b}_i)$. Since we pick pk_j randomly, and as well as \mathbf{a}_i (observe that η_i , \mathbf{a}_i , and \mathbf{b}_i might be correlated), we have $\Pr[B_{i,j}] \leq q_H^{\text{pre}}/p^2$. We now define the event $\text{DetectProgramming} = \bigcup_{j \leq N, i \leq q_S} B_{i,j}$. Applying union bounds, we have

$$\Pr[\text{DetectProgramming}] = \Pr[\bigcup_{i,j} B_{i,j}] \leq \sum_{i,j} \Pr[B_{i,j}] \leq \frac{N q_H^{\text{pre}} q_S^{\text{on}}}{p^2}.$$

Assuming that the event DetectProgramming does not occur, the attacker will not be able to distinguish between a programmed response and the original response.

- (2) One way to output failure is during the signing oracle if $H(\text{pk}_j \| H'(\text{pk}_j \| m_i) \| \eta_i \| \mathbf{a}_i \| \mathbf{b}_i)$ has been previously queried during the *online* phase. We define this event as FailtoSign . As before we can argue that $\Pr[\text{FailtoSign}] \leq \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p}$ i.e., see [Theorem 9](#).
- (3) We also consider a bad event $\text{Collision}_{H'}$ which denotes the event that the attacker finds a collision on $H'(\text{pk}_{i^*} \| \cdot)$ so that the attacker can forge a signature without the bridge event BRIDGE^N occurring. [Claim 6](#) upper bounds the probability of this event using a similar compression argument as in [Claim 5](#). The proof of [Claim 6](#) is given below.

$$\Pr[\text{FailtoFind} \cup \text{BadQuery}] \leq \frac{4(q_H^{\text{on}} + 1)}{2^{2k}} + \frac{N^2(S + 2k)}{p}.$$

Claim 6. $\Pr[\text{Collision}_{H'}] \leq \frac{4\tilde{q}_{\text{on}}^2}{2^{2k}} + \frac{N^2(S + 2k)}{2p}$, where $\tilde{q}_{\text{on}} = q_H^{\text{on}} + 2q_S^{\text{on}}$.

Applying [Theorem 7](#) and from the similar argument in [Theorem 8](#), we have

$$\begin{aligned} & \Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau}, \text{H}}^{\tau, N}}(k) = 1 \right] \\ & \leq \Pr \left[\text{BridgeChal}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau}, \text{H}}^{\tau, N}}(k) = 1 \right] + \Pr[\text{DetectProgramming}] + \Pr[\text{FailtoSign}] \\ & \quad + \Pr[\text{FailtoFind}] + \Pr[\text{BadQuery}] + \Pr[\text{Collision}_{H'}] \\ & \leq \tilde{\mathcal{O}} \left(\frac{SN(q_S^{\text{on}} + N)(q_S^{\text{on}} + 2N)}{p} \right) + \frac{N q_H^{\text{pre}} q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} \end{aligned}$$

$$+ \frac{4(q_H^{\text{on}} + \tilde{q}_{\text{on}}^2 + 1)}{2^{2k}} + \frac{3N^2(S + 2k)}{2p},$$

which completes the proof. \square

Remark 7. Similar to [Theorem 8](#), if $q_G^{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then we can argue that

$$\begin{aligned} & \Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig, str}_{\tau, \text{H}}}^{\tau, N}}(k) = 1 \right] \\ & \leq \frac{12SN(q_G^{\text{on}})^2 \log p}{p} + \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^3} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p^2} + \frac{4(q_H^{\text{on}} + \tilde{q}_{\text{on}}^2 + 1)}{2^{2k}} + \frac{3N^2(S + 2k)}{2p}, \end{aligned}$$

for a key-prefixed Chaum-Pedersen signature scheme Π . \triangleleft

Reminder of Claim 6. $\Pr[\text{Collision}_{\text{H}'}] \leq \frac{4\tilde{q}_{\text{on}}^2}{2^{2k}} + \frac{N^2(S+2k)}{2p}$, where $\tilde{q}_{\text{on}} = q_H^{\text{on}} + 2q_S^{\text{on}}$.

Proof of Claim 6: Fixing τ and H , and for some parameter t (which we will determine later), we define $\varepsilon_{F, \tau, \text{H}}$ to be $\varepsilon_{F, \tau, \text{H}} := \min_{Y: |Y| \leq Nt} \Pr[\text{Collision}_{\text{H}'}]$, where the probability is taken over the random selection of $x_1, \dots, x_N \in \mathbb{Z}_p \setminus Y$ and the random coins of $\mathcal{A}_{\text{sig, str}_{\tau, \text{H}}}^{\text{on}}$. Similar to [Claim 5](#), we have that $\varepsilon'_{F, \tau, \text{H}} \leq \varepsilon_{F, \tau, \text{H}} + N^2t/p$ where $\varepsilon'_{F, \tau, \text{H}}$ is the same probability when $x_1, \dots, x_N \in \mathbb{Z}_p$ are chosen without any restrictions. We also sample t uniformly random strings R_1, \dots, R_t by querying the random oracle at $\mathcal{O}(t)$ fixed points and repeat the following experiment t times to extract t input/output predictions by using R_i as its random coin for each $i \in [t]$.

Let $z_1, \dots, z_{\tilde{q}_{\text{on}}}$ be the queries to the random oracles either directly or indirectly through the signing oracle, and consider the following experiment:

- (1) For $i \in [t]$, sample $\bar{x}^i = (x_1^i, \dots, x_N^i)$ subject to the restriction that $x_j^i \neq x_j^{i'}$ for any pair $(i, j) \neq (i', j')$.
- (2) Pick a pair $(q_1^i, q_2^i) \in [\tilde{q}_{\text{on}}] \times [\tilde{q}_{\text{on}}]$ uniformly at random subject to the restriction that $q_1^i < q_2^i$.
- (3) Simulate $\mathcal{A}_{\text{sig, str}_{\tau, \text{H}}}^{\text{on}}$ until the query q_1^i to the random oracle and observe the output $W_{q_1^i} := \text{H}(z_{q_1^i})$. Record the input/output pair $(z_{q_1^i}, W_{q_1^i})$ of this query.
- (4) For the query q_2^i , we predict that the output of query $z_{q_2^i}$ is equal to $W_{q_1^i}$ *without actually* querying the random oracle, i.e., output the input/output pair prediction $(z_{q_2^i}, W_{q_1^i})$.

During each iteration i , we output a correct input/output pair with probability at least $2\varepsilon_{F, \tau, \text{H}}/\tilde{q}_{\text{on}}^2$. Thus, the procedure above correctly outputs t input/output pairs with probability at least $(2\varepsilon_{F, \tau, \text{H}}/\tilde{q}_{\text{on}}^2)^t$. Similar to [Claim 5](#), we define SucceedExp as an event that the experiment above correctly outputs t input/output pairs (without querying the random oracle at these t points) when τ and H are picked randomly. Now let $\varepsilon_F := \mathbb{E}_{\tau, \text{H}}[\varepsilon_{F, \tau, \text{H}}]$ and to derive a contradiction, suppose that $\varepsilon_F > 4\tilde{q}_{\text{on}}^2/2^{2k}$. Then Markov inequality tells us that

$$\Pr_{\tau, \text{H}} \left[\varepsilon_{F, \tau, \text{H}} > \frac{2\tilde{q}_{\text{on}}^2}{2^{2k}} \right] > \frac{\varepsilon_F}{2}.$$

Thus, if we pick τ and H randomly and run the above procedure, we will succeed with probability at least

$$\Pr[\text{SucceedExp}] \geq \frac{\varepsilon_F}{2} \left(\frac{2\varepsilon_{F,\tau,H}}{\tilde{q}_{\text{on}}^2} \right)^t > 2\tilde{q}_{\text{on}}^2 2^{-2k} 2^{(2-2k)t}.$$

Similarly, since we have $\Pr[\text{SucceedExp}] \leq 2^{-2kt+S}$ from Lemma 4, picking $t = (S + 2k)/2$ we have a contradiction as we have

$$2\tilde{q}_{\text{on}}^2 2^{-kS-2k^2+S} < \Pr[\text{SucceedExp}] \leq 2^{-kS-2k^2+S}.$$

Hence, we finally have

$$\begin{aligned} \varepsilon'_F &:= \Pr[\text{Collision}_{H'}] \\ &= \mathbb{E}_{\tau,H}[\varepsilon'_{F,\tau,H}] \\ &\leq \mathbb{E}_{\tau,H} \left[\varepsilon_{F,\tau,H} + \frac{N^2(S+2k)}{2p} \right] \\ &= \varepsilon_F + \frac{N^2(S+2k)}{2p} \leq \frac{4\tilde{q}_{\text{on}}^2}{2^{2k}} + \frac{N^2(S+2k)}{2p}. \quad \square \end{aligned}$$

Reminder of Theorem 11. *The (short) Katz-Wang signature scheme is $(N, q_H, q_G, q_S, \mathcal{O}(\frac{q+N}{2^k}))$ -MU-UF-CMA secure under the generic group model of prime order $p \approx 2^{2k}$ and the programmable random oracle model, where q denotes the total number of queries made by an adversary.*

Proof Sketch of Theorem 11: Let Π be the Katz-Wang signature scheme. We follow essentially the same security reduction as in Theorem 6 using the signing oracle in Figure 3 (bottom). One minor difference here is that after \mathcal{A}_{sig} outputs $\sigma_{i^*} = (s_{i^*}, e_{i^*})$ and m_{i^*} for some $i^* \in [N]$, we first compute $\mathbf{a}_{j,i^*} = \text{Mult}(\text{Pow}(\tau(p_j), s_{i^*}), \text{Pow}(\text{Inv}(\mathbf{h}_j), e_{i^*}))$ for $j = 1, 2$, and check to see if we previously had any tuple of the form $(\mathbf{a}_{j,i^*}, \vec{a}_j, b_j) \in \mathcal{L}$ for any of $j = 1, 2$. If no such tuple exists, then we return \perp , and otherwise, we let \vec{a}, b be given such that $\mathbf{a}_{j,i^*} = \tau(\vec{a}_j \cdot \vec{x} + b_j)$ for some j (or both). If $\vec{a}_j + e_{i^*} \hat{u}_{i^*} = \vec{0}$, then we return \perp , and otherwise, we have a BRIDGE^N instance $(\tau(s_{i^*} - b_j), \vec{a}_j + e_{i^*} \hat{u}_{i^*}, 0) \in \mathcal{L}$ and $(\tau(s_{i^*} - b_j), \vec{0}, s_{i^*} - b_j) \in \mathcal{L}$. Then we have the following possibilities to output \perp for failure correspondingly:

- (1) One way to output failure is during the signing oracle, if $H(\mathbf{pk}_j \| \mathbf{a}_{1,i} \| \mathbf{a}_{2,i} \| m_i)$ has been previously queried during the online phase. We define this event as **FailtoSign**. Since \mathbf{pk}_j is fixed before attacker runs, we cannot view \mathbf{pk}_j as random. Note that every time the attacker queries the signing oracle, we would generate a new query to the random oracle H . Thus, we would have at most $q_H + q_S$ input/output pairs recorded for the random oracle. Since there are p possible choices for $\mathbf{a}_{1,i}$, the probability that $(\mathbf{pk}_j \| \mathbf{a}_{1,i} \| \mathbf{a}_{2,i} \| m_i)$ is one of the inputs is at most $(q_H + q_S)/p$. Applying union bound over q_S queries to the signing oracle, we conclude that $\Pr[\text{FailtoSign}] \leq q_S(q_H + q_S)/p$.

- (2) The second way to output failure is if *both* the value \mathbf{a}_{1,i^*} and \mathbf{a}_{2,i^*} does not previously appear in the list \mathcal{L} . We remark that if either value appears in the list, then we can still generate the BRIDGE^N event. We define this event by $\text{FailtoFind}(\mathbf{a}_{1,i^*}, \mathbf{a}_{2,i^*})$. Since we already know that for a single instance, $\Pr[\text{FailtoFind}(\mathbf{a}_{j,i^*})] \leq \frac{q_{\text{H}}+q_{\text{S}}}{p-|\mathcal{L}|} + 2^{-k}$ holds for each $j = 1, 2$ (i.e., see [Claim 2](#)), we have that $\Pr[\text{FailtoFind}(\mathbf{a}_{1,i^*}, \mathbf{a}_{2,i^*})] \leq \frac{q_{\text{H}}+q_{\text{S}}}{p-|\mathcal{L}|} + \frac{1}{2^k}$, because of the observation that $\Pr[A \cap B] \leq \Pr[A]$.
- (3) Finally, we could output failure if $\vec{a}_j = -e_{i^*}\hat{u}_{i^*}$ for both $j = 1$ and $j = 2$. We call this event BadQuery . From [Claim 3](#), we know that the probability to have a bad query for a single instance \mathbf{a}_{j,i^*} is upper bounded by $q_{\text{H}}/2^k$. Hence, from the same observation from (2), the probability to have bad queries for both instances is also upper bounded by $q_{\text{H}}/2^k$. Thus, $\Pr[\text{BadQuery}] \leq q_{\text{H}}/2^k$.

Since we have $|\mathcal{L}| \leq N + 3q_{\text{G}} + 1$, from [Theorem 5](#) we conclude that

$$\begin{aligned} & \Pr[\text{SigForge}_{\mathcal{A}_{\text{sig},\Pi}^{\tau,N}}(k) = 1] \\ & \leq \frac{q_{\text{G}}N + 3q_{\text{G}}(q_{\text{G}} + 1)/2}{p - (N + 3q_{\text{G}} + 1)^2 - N} + \frac{Nq_{\text{S}}(q_{\text{H}} + q_{\text{S}})}{p^2} + \frac{q_{\text{H}} + q_{\text{S}}}{p - (N + 3q_{\text{G}} + 1)} + \frac{q_{\text{H}} + 1}{2^k} \\ & = \mathcal{O}\left(\frac{q + N}{2^k}\right). \quad \square \end{aligned}$$

Reminder of [Theorem 12](#). Let $\Pi = (\text{Kg}, \text{Sign}, \text{Vfy})$ be a Katz-Wang signature scheme and $p > 2^{2k}$ be a prime number. Let $N \in \mathbb{N}$ be a parameter and $(\mathcal{A}_{\text{sig}}^{\text{pre}}, \mathcal{A}_{\text{sig}}^{\text{on}})$ be a pair of generic algorithms with an encoding map $\tau : \mathbb{Z}_p \rightarrow \mathbb{G}$ such that $\mathcal{A}_{\text{sig}}^{\text{pre}}$ makes at most $q_{\text{H}}^{\text{pre}} < 2^{3k}$ queries to the random oracle at most $q_{\text{H}}^{\text{pre}} < 2^{3k}$ queries to the random oracle $\text{H} : \{0, 1\}^* \rightarrow \{0, 1\}^{k_1}$ and outputs an S -bit hint $\text{str}_{\tau, \text{H}}$, and $\mathcal{A}_{\text{sig}}^{\text{on}}$ makes at most $q_{\text{G}}^{\text{on}} := q_{\text{G}}^{\text{on}}(k)$ queries to the generic group oracles and at most q_{H}^{on} queries to the random oracle. Then $\Pr[\text{SigForge}_{\mathcal{A}_{\text{sig}, \text{str}_{\tau, \text{H}}}^{\tau, N}}(k) = 1] \leq \varepsilon$, with

$$\varepsilon = \tilde{\mathcal{O}}\left(\frac{SN(q_{\text{G}}^{\text{on}} + N)(q_{\text{G}}^{\text{on}} + 2N)}{p}\right) + \frac{Nq_{\text{H}}^{\text{pre}}q_{\text{S}}^{\text{on}}}{p^2} + \frac{q_{\text{S}}^{\text{on}}(q_{\text{S}}^{\text{on}} + q_{\text{H}}^{\text{on}})}{p} + \frac{4(q_{\text{H}}^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p},$$

where q_{S}^{on} denotes the number of queries to the signing oracle and the randomness is taken over the selection of τ and the random coins of $\mathcal{A}_{\text{sig}}^{\text{on}}$ (the hint $\text{str}_{\tau, \text{H}} = \mathcal{A}_{\text{sig}}^{\text{pre}, \text{GO}}(\mathbf{g})$ is selected independently of the random coins used by the challenger).

Proof Sketch of [Theorem 12](#): We follow essentially the same security reduction as in [Theorem 8](#), except for the differences as follows:

- When the attacker tries to forge a signature on the message m_i without knowledge of the secret key x_j , s/he follow [Figure 3](#) (bottom).
- After \mathcal{A}_{sig} outputs $\sigma_{i^*} = (s_{i^*}, e_{i^*})$ and m_{i^*} for some $i^* \in [N]$, we first compute $\mathbf{a}_{j,i^*} = \text{Mult}(\text{Pow}(\tau(p_j), s_{i^*}), \text{Pow}(\text{Inv}(\mathbf{h}_j), e_{i^*}))$ for $j = 1, 2$, and check

to see if we previously had any tuple of the form $(\mathbf{a}_{j,i^*}, \vec{a}_j, b_j) \in \mathcal{L}$ for any of $j = 1, 2$. If no such tuple exists, then we return \perp , and otherwise, we let \vec{a}, b be given such that $\mathbf{a}_{j,i^*} = \tau(\vec{a}_j \cdot \vec{x} + b_j)$ for some j (or both). If $\vec{a}_j + e_{i^*} \hat{u}_{i^*} = \vec{0}$, then we return \perp , and otherwise, we have a **BRIDGE**^{*N*} instance $(\tau(s_{i^*} - b_j), \vec{a}_j + e_{i^*} \hat{u}_{i^*}, 0) \in \mathcal{L}$ and $(\tau(s_{i^*} - b_j), \vec{0}, s_{i^*} - b_j) \in \mathcal{L}$.

Then we have the following possibilities to output \perp for failure correspondingly. Everything goes the same with the analysis in [Theorem 8](#) except for the following:

- (1) Let $B_{i,j}$ denote the bad event that the preprocessing attacker previously submitted a query of the form $H(\text{pk}_j \| \mathbf{a}_{1,i} \| \mathbf{a}_{2,i} \| m_i)$. Since we pick pk_j randomly, and as well as $\mathbf{a}_{1,i}$, we have $\Pr[B_{i,j}] \leq q_H^{\text{pre}}/p^2$. We now define the event $\text{DetectProgramming} = \bigcup_{j \leq N, i \leq q_S} B_{i,j}$. Applying union bounds, we have

$$\Pr[\text{DetectProgramming}] = \Pr[\bigcup_{i,j} B_{i,j}] \leq \sum_{i,j} \Pr[B_{i,j}] \leq \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2}.$$

Assuming that the event **DetectProgramming** does not occur, the attacker will not be able to distinguish between a programmed response and the original response.

- (2) One way to output failure is during the signing oracle if $H(\text{pk}_j \| \mathbf{a}_{1,i} \| \mathbf{a}_{2,i} \| m_i)$ has been previously queried during the online phase. We define this event as **FailtoSign**. As before, we can argue that $\Pr[\text{FailtoSign}] \leq \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p}$, i.e., see [Theorem 11](#).
- (3) Another way to output \perp is if the event **DetectProgramming** did not occur and both the value \mathbf{a}_{1,i^*} and \mathbf{a}_{2,i^*} does not previously appear in the list \mathcal{L} . We call this event **FailtoFind** $(\mathbf{a}_{1,i^*}, \mathbf{a}_{2,i^*})$. The final case to output failure is if both $\vec{a}_j = -e_{i^*} \hat{u}_{i^*}$ holds for $j = 1$ and $j = 2$, and we call this event **BadQuery**. Similar to the argument from [Theorem 11](#), each probability is bounded by the event with the single instance. Hence, we can directly adapt the same argument from [Theorem 8](#) and say that

$$\Pr[\text{FailtoFind}(\mathbf{a}_{1,i^*}, \mathbf{a}_{2,i^*}) \cup \text{BadQuery}] \leq \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}.$$

Applying [Theorem 7](#) and from the similar argument in [Theorem 8](#), we have

$$\begin{aligned} & \Pr \left[\text{SigForge}_{\text{sig, str}_{\tau, H}}^{\tau, N}(\Pi)(k) = 1 \right] \\ & \leq \tilde{\mathcal{O}} \left(\frac{SN(q_G^{\text{on}} + N)(q_G^{\text{on}} + 2N)}{p} \right) + \frac{Nq_H^{\text{pre}}q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} \\ & \quad + \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}, \end{aligned}$$

which completes the proof. \square

Remark 8. Similar to [Theorem 8](#), if $q_G^{\text{on}} \geq 10N(1 + 2 \log p)$ and $S \geq 10 \log(8p)$, then we can argue that

$$\begin{aligned} & \Pr \left[\text{SigForge}_{\mathcal{A}_{\text{sig, str}_{\tau, H}^{\text{on}}}, \Pi}^{\tau, N}(k) = 1 \right] \\ & \leq \frac{12SN(q_G^{\text{on}})^2 \log p}{p} + \frac{Nq_H^{\text{pre}} q_S^{\text{on}}}{p^2} + \frac{q_S^{\text{on}}(q_S^{\text{on}} + q_H^{\text{on}})}{p} + \frac{4(q_H^{\text{on}} + 1)}{2^{k_1}} + \frac{N^2(S + k_1)}{p}, \end{aligned}$$

for the Katz-Wang signature scheme *II*. ◁

D Prior Analysis on Schnorr Signatures

In this section, we will discuss in detail why the security bounds from Neven et al. [\[NSW09\]](#) do not yield k bits of security for *short* Schnorr signatures, as we mentioned in the introduction.

Theorem 13 ([\[NSW09, Theorem 1\]](#)). *If the discrete logarithm problem in a group \mathbb{G} is $(t_{\text{dlog}}, \epsilon_{\text{dlog}})$ -hard, then the Schnorr signature scheme is $(t_{\text{uf-cma}}, q_S, q_H, \epsilon_{\text{uf-cma}})$ -secure for*

$$\epsilon_{\text{uf-cma}} = \sqrt{(q_H + q_S + 1) \cdot \epsilon_{\text{dlog}}} + \frac{q_H + q_S + 1}{2^n} + \frac{q_S(q_H + q_S + 1)}{p}$$

and $t_{\text{uf-cma}} = t_{\text{dlog}}/2 - q_S t_{\text{exp}} + \mathcal{O}(q_H + q_S + 1)$, where t_{exp} is the cost of an exponentiation in the group \mathbb{G} , $p = |\mathbb{G}|$ denotes the size of the group and n denotes the length of a hash output.

As discussed by Neven et al. [\[NSW09\]](#) [Theorem 13](#) shows that we can use shorter hashes, i.e., $n = k$ to achieve k -bit security. However, we must also use a larger group size $p \approx 2^{3k}$ to achieve k -bit security due to the $\sqrt{(q_H + q_S + 1) \cdot \epsilon_{\text{dlog}}}$ term in the security bound. In particular, for any group we have $\epsilon_{\text{dlog}} \geq t_{\text{dlog}}^2/p$ as the following generic attack demonstrates: Given $h = g^x$, we can compute g^{jt} for $j = 1, \dots, t_{\text{dlog}}$ and g^{x+i} for $i = 1, \dots, t_{\text{dlog}}$. If $x < t_{\text{dlog}}^2$, then we will find some i, j such that $g^{x+i} = g^{jt_{\text{dlog}}}$, which allows us to solve for $x = jt_{\text{dlog}} - i$. The attack takes time $\mathcal{O}(t_{\text{dlog}})$ and succeeds with probability t_{dlog}^2/p . If we set $p \approx 2^{2k}$ then $\epsilon_{\text{dlog}} \geq t^2/p$ and [Theorem 13](#) does not guarantee k -bit security, e.g., plugging in $t_{\text{dlog}} = 2^{2k/3}$ and $q_H + q_S + 1 = 2^{2k/3}$ the security bound becomes meaningless as we have

$$\begin{aligned} \epsilon_{\text{uf-cma}} & \geq \sqrt{(q_H + q_S + 1) \cdot \epsilon_{\text{dlog}}} \\ & = \sqrt{2^{2k/3} \cdot 2^{4k/3}/p} \\ & \approx 1. \end{aligned}$$

We can address the issues by increasing the size p of our group \mathbb{G} , but if we take $p \approx 2^{3k}$ the total length of a Schnorr signature is *at least* $4k$ bits even

if we truncate the hash output to just $n = k$ -bits. Neven et al. [NSW09] also present a bound with better dependence on the group size, but this second bound requires larger hash outputs. The second bound assumes that our group is a generic group, but does not assume that the hash function is a random oracle. Instead it relies on assumptions that the hash function satisfies random-prefix preimage resistance (rpp) and random-prefix second-preimage resistance (rpsp) – see [NSW09] for formal definitions.

Theorem 14 ([NSW09, Theorem 2]). *Let $H : \{0, 1\}^* \rightarrow \{0, 1\}^n \hookrightarrow \mathbb{Z}_p$ be some hash function, and let G be some fixed group modelled as a generic group over the set of bit strings \mathbb{G} . Let $f : \mathbb{G} \rightarrow \{0, 1\}^d$ be an almost-invertible conversion function. If both the rpp[D] and the rpsp[D] are hard for H , with respect to the domain $D = \{0, 1\}^d$, then the Schnorr signature scheme $\text{Sch}[H]$ is secure in the generic group model.*

In particular, for an adversary \mathcal{A} against the Schnorr signature scheme, they proved that

$$\text{Adv}_{\text{Sch}[H]}^{\text{uf-cma}}(\mathcal{A}) \leq \frac{q_G}{\delta} \cdot \epsilon_{\text{rpp}} + \frac{q_S + 2}{\delta} \cdot \epsilon_{\text{rpsp}} + \mathcal{O}\left(\frac{(q_S + q_G)^2}{p}\right),$$

where q_G (resp. q_S) denote the upper bounds on the number of queries that \mathcal{A} makes to its generic group oracle (resp. signing oracle), and ϵ_{rpp} (resp. ϵ_{rpsp}) denote the probability to win the rpp (resp. rpsp) problem.

To achieve k bits of security using the upper bound in Theorem 14, we can use a smaller group size, i.e., $p \approx 2^{2k}$, but we would need longer hash outputs, i.e., $n = 2k$. If we assume that it takes one time unit for evaluating the hash function H , then generic attacks imply that $\epsilon_{\text{rpp}} = \Omega(t/2^n)$ and $\epsilon_{\text{rpsp}} = \Omega(t/2^n)$ [NSW09]. Thus, to achieve k -bit security we must set $n \geq 2k$ — otherwise when $q_G = \Omega(t)$ (resp. $q_S = \Omega(t)$) the term $q_G \cdot \epsilon_{\text{rpp}} = \Omega\left(\frac{t^2}{2^n}\right)$ (resp. $q_S \cdot \epsilon_{\text{rpp}} = \Omega\left(\frac{t^2}{2^n}\right)$) is too large to achieve k -bit security. Thus, to obtain k -bit security from Theorem 14 the length of the Schnorr signature will be at least $n + \log p \geq 4k$ bits.