

The many faces of Schnorr

Victor Shoup
DFINITY

June 30, 2023

Abstract

Recently, a number of highly optimized threshold signing protocols for Schnorr signatures have been proposed. A key feature of these protocols is that they produce so-called “presignatures” in a “offline” phase (using a relatively heavyweight, high-latency subprotocol), which are then consumed in an “online” phase to generate signatures (using a relatively lightweight, low-latency subprotocol). The idea is to build up a large cache of presignatures in periods of low demand, so as to be able to quickly respond to bursts of signing requests in periods of high demand. Unfortunately, it is well known that using such presignatures naively leads to subexponential attacks. Thus, any protocols based on presignatures must mitigate against these attacks.

One such notable protocol is FROST, which provides security even with an unlimited number of presignatures; moreover, assuming unused presignatures are available, signing requests can be processed concurrently with minimal latency. Unfortunately, FROST is not a robust protocol, at least in the asynchronous communication model (arguably the most realistic model for such a protocol). Indeed, a single corrupt party can prevent any signatures from being produced. Recently, a protocol called ROAST was developed to remedy this situation. Unfortunately, ROAST is significantly less efficient than FROST (each signing request runs many instances of FROST concurrently).

A more recent protocol is SPRINT, which provides robustness without synchrony assumptions, and actually provides better throughput than FROST. Unfortunately, SPRINT is only secure in very restricted modes of operation. Specifically, to avoid a subexponential attack, only a limited number of presignatures may be produced in advance of signing requests, which somewhat defeats the purpose of presignatures.

Our main new result is to show how to securely combine the techniques used in FROST and SPRINT, allowing one to build a threshold Schnorr signing protocol that (i) is secure and robust without synchrony assumptions (like SPRINT), (ii) provides security even with an unlimited number of presignatures, and (assuming unused presignatures are available) signing requests can be processed concurrently with minimal latency (like FROST), (iii) achieves high throughput (like SPRINT), and (iv) achieves optimal resilience.

Besides achieving this particular technical result, one of our main goals in this paper is to provide a *unifying framework* in order to better understand the techniques used in various protocols. To that end, we attempt to isolate and abstract the main ideas of each protocol, stripping away superfluous details, so that these ideas can be more readily combined and implemented in different ways. More specifically, to the extent possible, we try to avoid talking about distributed protocols at all, and rather, we examine the security of the ordinary, non-threshold Schnorr scheme in “enhanced” attack modes that correspond to attacks on various types of threshold signing protocols.

Another one of our goals to carry out a security analysis of these enhanced attack modes in the Generic Group Model (GGM), sometimes in conjunction with the Random Oracle Model (ROM). Despite the limitations of these models, we feel that giving security proofs in the GGM or GGM+ROM provides useful insight into the concrete security of the various enhanced attack modes we consider.

1 Introduction

1.1 Background

Recently, a number of highly optimized threshold signing protocols for Schnorr signatures have been proposed. Threshold signing protocols are useful in that they can provide both *security* and *robustness*, even if some of the parties on the signing committee are corrupt — *security*, in the sense that even if a bounded number parties are corrupt, they cannot forge a signature, and *robustness*, in the sense that even if a bounded number parties are corrupt, they cannot stop the honest parties from producing signatures.

Recall that for the Schnorr signature scheme, the public key is of the form $\mathcal{D} = d\mathcal{G}$, where $d \in \mathbb{Z}_q$ is the secret key and \mathcal{G} is a generator for a group E of prime order q (which we write here using additive notation to reflect the fact that E is nowadays typically an elliptic curve). A signature on a message m is a pair $(\mathcal{R}, z) \in E \times \mathbb{Z}_q$, where $z\mathcal{G} = \mathcal{R} + h\mathcal{D}$ and $h \in \mathbb{Z}_q$ is a hash of \mathcal{R} and m (and typically \mathcal{D} as well). To generate such a signature in the non-threshold setting, the signer generates $r \in \mathbb{Z}_q$ at random, computes $\mathcal{R} \leftarrow r\mathcal{G}$ and $z \leftarrow r + hd$, and outputs the signature (\mathcal{R}, z) .

In the threshold setting, we have n parties on a signing committee, some of which may be corrupt, and a certain threshold number of parties is needed to sign a message (and assuming this threshold is high enough, some honest party must actually participate in signing the message). The usual technique used is Shamir secret sharing, so that each of the n parties obtains \mathcal{D} and its share of d . To generate a public-key/secret-key pair, some kind of distributed key generation (DKG) protocol must be executed, which can be rather expensive.

To sign an individual message, in principle, the same DKG protocol could be used to generate the “ephemeral” public-key/secret-key pair (\mathcal{R}, r) , where each party obtains \mathcal{R} and its share of r . Once this is done, each party can locally compute its share of the signature (since this is a linear operation), and then these shares can be revealed and combined to form a signature.

The problem with this approach is that an expensive DKG protocol must be run for each signing operation. To reduce this cost, a few optimizations have been considered.

One obvious optimization follows from the observation that the “ephemeral” public-key/secret-key pair (\mathcal{R}, r) is completely independent of the message to be signed. Therefore, we could potentially use an “offline/online” strategy, in which we generate such ephemeral key pairs in an offline fashion, building up a cache of them in advance of actual signing requests. In this context, such an ephemeral public key is called a **presignature**. The idea is to build up a large cache of presignatures in periods of low demand, so as to be able to quickly respond to bursts of signing requests in periods of high demand. Note, however, that while computing presignatures in this way can improve *latency*, it does not improve *throughput*.

Unfortunately, using presignatures naively in this way breaks the security of Schnorr signatures. Indeed, the usual proof of security of ordinary, non-threshold Schnorr signatures relies in an essential way on the fact that the randomly generated group element \mathcal{R} is not revealed before the request to sign m is given. Moreover, this is not just an artifact of the proof: there are actual subexponential attacks on signing protocols that use presignatures in this way [DEF⁺18] (which we review below).

To mitigate against these presignature attacks, the FROST protocol [KG20, CKM21] was introduced. FROST provides security even with an unlimited number of presignatures; moreover, assuming unused presignatures are available, signing requests can be processed concurrently with minimal latency. However, FROST is not a robust protocol. Indeed, a single corrupt party can prevent any signatures from being produced. Nevertheless, FROST does enjoy a property called *identifiable abort*, which allows misbehaving parties that prevent protocol termination to be identified and removed from the signing committee. The use of *identifiable aborts* in the context of threshold signatures is also found in the work of [GG20]. However, the notion of *identifiable aborts* only makes sense in a *synchronous communication setting*. Indeed, in an *asynchronous communication setting*, it is impossible to tell the difference between a party that is misbehaving by staying silent and a party that is just slow or temporarily disconnected from the rest of the parties. Thus, at least in an asynchronous communication setting, FROST does not provide robustness. This makes FROST unusable in distributed systems for which both security and robustness are required without synchrony assumptions. Indeed, for a protocol with parties distributed around the globe, synchrony assumptions seem quite unrealistic.

This limitation of FROST was highlighted in [RRJ⁺22], who propose a new protocol called ROAST. To obtain robustness without synchrony assumptions, the ROAST protocol uses FROST (or any protocol with similar security properties) as a subprotocol, running it concurrently $O(n)$ times per signing request. Thus, while ROAST achieves robustness without synchrony assumptions, this comes at a significant performance cost.

More recently, the SPRINT protocol [BHK⁺23] was proposed, which aims to achieve security and robustness without synchrony assumptions, and to do so while actually providing *better throughput* than FROST by using improved presignature generation protocols based on batch randomness extraction techniques (an idea that goes back to [HN06]). While SPRINT does achieve this goal, it is only secure in very restricted modes of operation. Specifically, only a limited number of presignatures may be generated in advance of signing requests, which somewhat defeats the purpose of presignatures. Indeed, the security theorem in [BHK⁺23] only applies to a chosen message attack in which a single, fixed-size batch of presignatures is generated, which are subsequently used to sign a corresponding batch of messages. As we discuss below in Section 4.1, if many such batches of presignatures are generated in advance, the same subexponential attacks mentioned above can be used on SPRINT.

1.2 Our contributions

On a purely technical level, our main new result is to show how the batch randomness extraction technique used in SPRINT can be securely combined with the main technical idea of FROST for making presignatures safe, thus allowing one to build a threshold Schnorr signing protocol that

- is secure and robust without synchrony assumptions (like SPRINT),
- provides security even with an unlimited number of presignatures, and (assuming unused presignatures are available) signing requests can be processed concurrently with minimal latency (like FROST),

- achieves high throughput (like SPRINT), and
- achieves optimal resilience (i.e., tolerates up to $f < n/3$ corrupt parties).

Note that one variant of SPRINT also considers so-called “packed” secret sharing, which can give even higher throughput at the cost of suboptimal resilience.¹ We do not consider this type of protocol here, although our techniques and analysis may well apply.

Besides achieving this particular technical result, one of our main goals in this paper is to provide a *unifying framework* in order to better understand the techniques used in various papers. Indeed, the analyses in the papers [KG20, CKM21, BHK⁺23] are quite targeted to very specific protocols (although [CKM21] makes some attempt to be a bit more modular), and it is not clear how ideas from one protocol can be used in a different context. Here, we attempt to isolate and abstract the main ideas of each protocol, stripping away superfluous details, so that these ideas can be more readily combined and implemented in different ways. Indeed, our approach is much more like that of [GS21], in that we try to avoid talking about distributed protocols at all, and rather, we examine the security of the ordinary, non-threshold Schnorr scheme in “enhanced” attack modes that correspond to attacks on various types of threshold signing protocols (which may use presignatures, for example) — the details of these threshold protocols do not matter that much, so long as they are designed in a reasonably modular way so as to satisfy certain natural security properties. (That said, to analyze various batch randomness extraction techniques, we do have to explicitly model the fact that there are multiple parties in the signing protocol who contribute randomness, but still, we abstract away most other nonessential details.)

Another one of our goals to carry out a security analysis of these enhanced attack modes in the Generic Group Model (GGM). Such an analysis in the GGM has already been done by [NSW09] for the basic attack mode on Schnorr, but not for any of the enhanced attack modes we consider here. The analysis in [NSW09] proves the security of Schnorr for the basic attack mode in the GGM under specific preimage assumptions on the underlying hash function. In fact, we reprove the results in [NSW09]. Our reasons for this our twofold. First, we want to establish a general framework for proving results on various Schnorr attack modes. This framework is very similar to that introduced in [GS21], in which attack in the GGM is reduced to a purely “symbolic” attack that allows for a much more modular and intuitive security analysis. Second, we actually prove a bit more than what is proved in [NSW09], observing that if we use both the GGM and Random Oracle Model (ROM), where the hash function is modeled as a random oracle, we get a very tight security bound: any adversary that makes at most N oracle queries (to either the signing, group, or random oracles) forges a signature with probability $O(N^2/q + N/M)$. Here, M is the size of the output space of the hash function.

We feel that giving security proofs in the GGM or GGM+ROM provides useful insight into the practical security of the various enhanced attack modes we consider. For example, the attack modes that correspond to FROST and SPRINT have been analyzed in the literature in the ROM, with reductions to the one-more discrete logarithm problem (for

¹“Packed” secret sharing is a technique introduced in [FY92], in which many secrets are packed into a single Shamir secret sharing, and is not to be confused with “batched” secret sharing, in which many independent Shamir secret sharings are generated concurrently, as in [DN07], for example. Indeed, protocols that implement the strategies we outline here may very well use “batched” secret sharing.

FROST) or the discrete logarithm problem (for SPRINT). However, these reductions all go via the so-called “forking lemma” [PS96], which yields very “loose” security reductions. Even though security proofs in the GGM or GGM+ROM have limitations in the generality of attacks they consider, they also have value by giving a better understanding of concrete security against the types of attacks that are arguably most likely to be carried out in practice.

We give security proofs in the GGM+ROM for a number of enhanced attack modes, including those corresponding to FROST, SPRINT, and our new technique that combines the best of both FROST and SPRINT. In all cases, we find that the adversary’s forging probability is still $O(N^2/q + N/M)$, just as for the basic attack mode. For several enhanced attack modes, we also give security proofs in the GGM under specific preimage assumptions on the hash function. Note that our analysis of our new technique combining FROST and SPRINT (which is covered in Section 4.3) is only done in the GGM+ROM. We speculate that a security proof in the ROM via a reduction to the one-more discrete logarithm problem should be possible — we leave that as an open problem.

In addition to all of the above, we also model **additive key derivation**. Here, when the adversary makes a signing query, he additionally specifies an additive tweak $e \in \mathbb{Z}_q$ to derive the effective public key as $\mathcal{D}' := \mathcal{D} + e\mathcal{G}$. This corresponds to using a scheme like BIP32 [Wui20] to derive subkeys from a master key. This type of key derivation is especially important in a threshold setting, as there is a significant cost to maintaining a secret key — for example, it will likely need to be reshared with regular frequency, both to achieve proactive security and to support membership changes to the signing committee. With additive key derivation, a signing committee can just maintain a single master key, and derive subkeys as necessary on behalf of individual external users (or “smart contracts” in a blockchain setting). Moreover, because of the simple additive nature of the key derivation, it is generally trivial to deal with these derived keys in a distributed computation. Not surprisingly, including the (effective) public key in the hash used to derive h is necessary and sufficient to obtain security proofs for all of the attack modes we consider.

2 Preliminaries

2.1 Schnorr Signatures

From now on, we consider the Schnorr signature scheme over an elliptic curve. Let E be an elliptic curve defined over \mathbb{Z}_p and generated by a point \mathcal{G} of prime order q , and let E^* be the set of points (x, y) on the curve excluding the point at infinity \mathcal{O} .

The secret key for ECDSA is a random $d \in \mathbb{Z}_q$, the public key is $\mathcal{D} = d\mathcal{G} \in E$. The scheme makes use of a hash function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q$. The signing and verification algorithms are shown in Figure 1. Here, we assume a serialization function

$$\langle \cdot \rangle : E \rightarrow \{0, 1\}^*$$

that is prefix-free and is 1-1 (as well as easy to compute and to invert).

NOTE: The scheme presented in Figure 1 does not quite fully capture either BIP340 (the bitcoin version of Schnorr) or EdDSA — each have their own quirks. However, it seems

Sign message m :	Verify signature $(\mathcal{R}, z) \in E \times \mathbb{Z}_q$ on m :
$r \xleftarrow{\$} \mathbb{Z}_q, \mathcal{R} \leftarrow r\mathcal{G} \in E$ $h \leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m) \in \mathbb{Z}_q$ $z \leftarrow r + hd$ return the signature (\mathcal{R}, z)	$h \leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m) \in \mathbb{Z}_q$ check that $z\mathcal{G} = \mathcal{R} + h\mathcal{D}$

Figure 1: Schnorr signing and verification algorithms

reasonable to speculate that all of the results proved here can easily be adapted to those particular schemes.

2.2 Enhanced attack modes

In the basic attack game for signatures, the adversary makes a series of signing queries and then must forge a signature on some message that was not submitted as a signing query. This attack game needs to be modified in order to model attacks that can be carried out in the threshold setting. There are three variations to consider:

Presignatures. Here, the adversary instructs the challenger to generate presignatures $\mathcal{R}_1, \mathcal{R}_2, \dots$, which are random elements of E that are given to the adversary. In a signing query, the adversary specifies the index k of an unused presignature and a message m_k ; the challenger then signs m_k using \mathcal{R}_k .

This models the situation in the threshold setting where we do the expensive presignature computation in advance using a secure DKG protocol. Any secure DKG protocol may be used. For example, [GS22] provides fairly efficient DKG protocols that are secure and robust, with optimal resilience, in the asynchronous setting.

Biased presignatures. Here, when the adversary makes a signing query, in addition to specifying k and m_k , the adversary specifies a “bias” $(u_k, u'_k) \in \mathbb{Z}_q^* \times \mathbb{Z}_q$; the challenger then signs m_k using $\mathcal{R}'_k := u_k \mathcal{R}_k + u'_k \mathcal{G}$.

This models the situation in the threshold setting where we use a Verifiable Secret Sharing (VSS) protocol with Feldman commitments to generate presignatures, rather than Pedersen commitments. In such a situation, the adversary can bias the result. This type of biasing was discussed in [GJKR07] in the synchronous communication setting, and in Section A.3.6 [GS22] in the asynchronous communication setting. Typically, in the synchronous setting, we can take $\mu_k = 1$, while in the asynchronous setting, we can take μ_k to be a number between 1 and the number of parties on the signing committee.

Additive key derivation. Here, when the adversary makes a signing query, he additionally specifies an additive tweak $e_k \in \mathbb{Z}_q$ to derive the effective public key as $\mathcal{D}'_k := \mathcal{D} + e_k \mathcal{G}$. With this modification, the notion of a forgery must also be appropriately modified, so that the forgery includes a tweak $e^* \in \mathbb{Z}_q$ in addition to a message

m^* , and the forgery counts so long as $(m^*, e^*) \neq (m_k, e_k)$ for any (m_k, e_k) submitted as part of a signing query.

This corresponds to using a scheme like BIP32 [Wui20] to derive subkeys from a master key.

Additive key derivation can be considered either by itself, or in combination of one of the two variants above.

2.3 Proof techniques and known attacks

In the usual analysis of Schnorr, we model H as a random oracle. The main idea of the security proof is to reduce an attack on the signature scheme to an attack on the interactive identification scheme. In the latter attack, the adversary, playing the role of prover, may initiate many conversations with the challenger, who is playing the role of verifier. The adversary wins the attack game if he can make any of these verifiers accept.² To carry out this reduction, we program the random oracle, which allows us to (a) simulate signing queries and (b) translate the random challenges in the identification attack game into random oracle outputs in the signature attack game.

To simulate signing queries, when we get a message m to sign, we generate $z, h \in \mathbb{Z}_q$ at random, compute $\mathcal{R} \leftarrow z\mathcal{G} - h\mathcal{D}$, and program the random oracle representing H so that $H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m) := h$. This simulation fails only if $H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m)$ was already defined, which happens only with negligible probability since \mathcal{R} is chosen after m is specified.

With presignatures, the above proof falls apart, precisely because \mathcal{R} is chosen and given to the adversary before the adversary specifies m . Indeed, as is well known [DEF⁺18], there are attacks. Suppose the adversary is given presignatures $\mathcal{R}_1, \dots, \mathcal{R}_K$. The adversary sets

$$\mathcal{R}^* := \sum_{k \in [K]} \mathcal{R}_k,$$

and attempts to find messages m^*, m_1, \dots, m_K such that

$$H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}^* \rangle \parallel m^*) = \sum_{k \in [K]} H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}_k \rangle \parallel m_k).$$

This is an instance of the $(K + 1)$ -sum problem, a generalization of the Birthday Problem studied by Wagner [Wag02]. Indeed, the adversary can generate $(K + 1)$ lists of random numbers, where the first list is obtained by computing $H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}^* \rangle \parallel m^*)$ for various messages m^* , the second by computing $H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}_1 \rangle \parallel m_1)$ for various messages m_1 , and so on. This can generally be done much faster than the time $O(\sqrt{q})$ needed to break the discrete logarithm problem in E . Once this is done, the adversary can obtain signatures (\mathcal{R}_k, z_k) on m_k for $k \in [K]$. From this, the adversary can compute $z^* \leftarrow \sum_{k \in [K]} z_k$ so that (\mathcal{R}^*, z^*) is a valid signature on m^* .

²One can then reduce the security of the interactive identification scheme to the hardness of the discrete logarithm using the “forking lemma”.

2.4 Re-randomized presignatures

One mitigation to this security weakness is to use **re-randomized presignatures**, just as in [GS21]. The idea is that a random tweak $\delta_k \in \mathbb{Z}_k$ is chosen by the challenger after the signing request is made, so that the original presignature \mathcal{R}_k is replaced by the effective presignature $\mathcal{R}'_k := \mathcal{R}_k + \delta_k \mathcal{G}$ (the value δ_k is given to the adversary to model that once chosen by the system it is publicly known). The same mitigation can be applied to biased presignatures: the effective presignature is then $\mathcal{R}'_k := u_k \mathcal{R}_k + (u'_k + \delta_k) \mathcal{G}$.

To implement this technique in a threshold setting, some type of “Random Beacon” must be used. A Random Beacon is a mechanism for obtaining public random values that remain hidden and unpredictable until a time determined by the protocol. For example, a Random Beacon can be efficiently implemented using a threshold BLS signature scheme [BLS01, Bol03]. Since the re-randomization is linear, in terms of working with linear secret sharing, the impact is negligible. Depending on the details of the system, obtaining the value δ_k from the Random Beacon may result in some additional latency — but not necessarily so. For example, on a distributed system such as the Internet Computer [DFI22], signing requests must go through a consensus mechanism, which itself may be implemented so that it uses a threshold BLS signature to achieve finalization; that very same threshold BLS signature can be used to derive δ_k .

Let us reconsider the proof of security with this mitigation. We will consider the re-randomized biased presignature setting (which includes the re-randomized presignature setting as a special case where $u_k = 1$ and $u'_k = 0$). We will also combine this with additive key derivation. Again, the main part of the proof is to simulate signing queries by programming the random oracle representing H . The simulator generates the presignature \mathcal{R}_k as

$$\mathcal{R}_k \leftarrow \zeta_k \mathcal{G} - \eta_k \mathcal{D},$$

where $\zeta_k, \eta_k \in \mathbb{Z}_q$ are chosen at random. At a later time, the adversary makes a corresponding signing query, where he specifies a message m_k an additive key tweak $e_k \in \mathbb{Z}_q$, and an presignature tweak $(u_k, u'_k) \in \mathbb{Z}_q^* \times \mathbb{Z}_q$. So the effective public key is $\mathcal{D}'_k := \mathcal{D} + e_k \mathcal{G}$, the effective presignature (used in the actual signature) is $\mathcal{R}'_k := u_k \mathcal{R} + (u'_k + \delta_k) \mathcal{G}$ and the resulting signature is (\mathcal{R}'_k, z_k) , where

$$z_k \mathcal{G} = \mathcal{R}'_k + h_k \mathcal{D}'_k = (u_k \mathcal{R}_k + (u'_k + \delta_k) \mathcal{G}) + h_k (\mathcal{D}_k + e_k \mathcal{G}),$$

which is equivalent to

$$\underbrace{u_k^{-1}(z_k - u'_k - \delta_k - e_k h_k)}_{=\zeta_k} \mathcal{G} = \mathcal{R}_k + \underbrace{u_k^{-1} h_k}_{=\eta_k} \mathcal{D}.$$

So the simulator can simply compute

$$h_k \leftarrow u_k \eta_k$$

and

$$z_k \leftarrow u_k \zeta_k + u'_k + \delta_k + e_k h_k,$$

and then program the random oracle so that $H(\langle \mathcal{D}'_k \rangle \parallel \langle \mathcal{R}'_k \rangle \parallel m_k) := h_k$. Because δ_k is chosen only after the adversary makes the signing request, the input is unlikely to have been used before and the programming of the oracle will fail only with negligible probability.

We give an alternative proof of security of re-randomized presignatures in the generic group model, below in Section 3.2.

2.5 Re-randomizing presignatures via hashing

The FROST [KG20] and FROST2 protocols [CKM21] use a hash function to derive the re-randomization tweak and uses a second random group element as a part of the presignature. We abstract away the details of that protocol in a way that is still useful in the context of a threshold Schnorr scheme built using robust MPC primitives, such as in [GS22]. To this end, a presignature consists of a pair of random group elements $(\mathcal{R}_k, \mathcal{S}_k)$. To sign a message m_k , the effective presignature (used in the actual signature) is

$$\mathcal{R}'_k := \mathcal{R}_k + \delta_k \mathcal{S}_k,$$

where

$$\delta_k := \Delta(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}_k \rangle \parallel \langle \mathcal{S}_k \rangle \parallel \langle k \rangle \parallel m_k).$$

Here, Δ is a hash function whose output space is \mathbb{Z}_q . Note that \mathcal{R}_k and \mathcal{S}_k could be biased presignatures.

The main advantage of this approach to re-randomizing presignatures is that in the threshold setting, we do not need a Random Beacon, as in Section 2.4.

The FROST2 protocol was analyzed in [CKM21] in the random oracle model, giving a reduction to one-more discrete log (OMDL). Below in Section 3.3, we give an analysis of the above abstract variant in the generic group model (where we also model the hash functions as random oracles). We believe this is useful because (a) the reduction to OMDL is extremely loose and our analysis here gives what is probably a more realistic bound on the effectiveness of any generic attacks, and (b) working in the generic group model allows us to examine further variants more quickly and easily.

NOTE: If we instead derive $\mathcal{R}'_k := \mathcal{R}_k + \delta_k \mathcal{G}$, where

$$\delta_k := \Delta(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}_k \rangle \parallel \langle k \rangle \parallel m_k).$$

one can carry out essentially the same attack as in Section 2.3. Indeed, suppose the adversary is given presignatures $\mathcal{R}_1, \dots, \mathcal{R}_K$. For $k \in [K]$, define

$$\delta_k(m) := \Delta(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}_k \rangle \parallel \langle k \rangle \parallel m)$$

and

$$h_k(m) := H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}_k + \delta_k(m) \mathcal{G} \rangle \parallel m).$$

The adversary sets

$$\mathcal{R}^* := \sum_{k \in [K]} \mathcal{R}_k,$$

and tries to find messages m^*, m_1, \dots, m_K such that

$$H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}^* \rangle \parallel m^*) = \sum_{k \in [K]} h_k(m_k).$$

This can again be done by solving an instance of the $(K + 1)$ -sum problem. Once this is done, the adversary asks for signatures $(\mathcal{R}_k + \delta_k(m_k)\mathcal{G}, z_k)$ on m_k for $k \in [K]$, computes

$$z^* \leftarrow \sum_{k \in [K]} (z_k - \delta_k(m_k)),$$

and outputs the forgery (\mathcal{R}^*, z^*) on m^* .

3 Generic Group Model analysis

In the above analysis, we give a reduction to breaking the interactive Schnorr identification scheme. That security property can be reduced to the DL problem via a “forking lemma” argument. This gives a very “loose” reduction. An alternative approach is to carry out an analysis in the Generic Group Model (GGM). For the basic Schnorr attack game, this has already been done in [NSW09]. However, we want to extend this to various extended attack games.

3.1 Analysis of basic attack

Our approach and proof technique will be as in [GS21].

3.1.1 The EC-GGM

We review the EC-GGM (Elliptic Curve Generic Group Model), introduced in [GS21]. We assume an elliptic curve E is defined by an equation $y^2 = F(x)$ over \mathbb{Z}_p and that the curve contains q points including the point at infinity \mathcal{O} . Here, p and q are odd primes. Let E^* be the set of non-zero points (excluding the point at infinity) on the curve, i.e., $(x, y) \in \mathbb{Z}_p \times \mathbb{Z}_p$ that satisfy $y^2 = F(x)$. From now on, we shall not be making any use of the usual group law for E , but simply treat E as a set; however, for a point $\mathcal{P} = (x, y) \in E^*$, we write $-\mathcal{P}$ to denote the point $(x, -y) \in E^*$.

An **encoding function** for E is a function

$$\pi : \mathbb{Z}_q \mapsto E$$

that is

- injective,
- **identity preserving**, meaning that $\pi(0) = \mathcal{O}$, and
- **inverse preserving**, meaning that for all $i \in \mathbb{Z}_q$, $\pi(-i) = -\pi(i)$.

In the EC-GGM, parties know E and interact with a **group oracle** \mathcal{O}_{grp} that works as follows:

- \mathcal{O}_{grp} on initialization chooses an encoding function π at random from the set of all encoding functions
- \mathcal{O}_{grp} responds to two types of queries:
 - (map, i) , where $i \in \mathbb{Z}_q$:
 - * return $\pi(i)$ // models computing $i\mathcal{G}$
 - $(\text{add}, \mathcal{P}_1, \mathcal{P}_2, c_1, c_2)$, where $\mathcal{P}_1, \mathcal{P}_2 \in E$ and $c_1, c_2 \in \mathbb{Z}_q$:
 - * return $\pi(c_1\pi^{-1}(\mathcal{P}_1) + c_2\pi^{-1}(\mathcal{P}_2))$ // models computing $c_1\mathcal{P}_1 + c_2\mathcal{P}_2$

NOTES:

1. The intuition is that the random choice of encoding function hides relations between group elements.
2. However, to make things more realistic, the encodings themselves have the same format as in a concrete elliptic curve, even though we do not at all use the group law of an elliptic curve.
3. Also to make things more realistic, the trivial relationship between a point and its inverse (that they share the same x -coordinate) is preserved.
4. Our model only captures the situation of elliptic curves over \mathbb{Z}_p of prime order and cofactor 1. This is sufficient for many settings, and it covers all of the “**secp**” curves in [Cer10].
5. We have enhanced slightly the EC-GCM model from [GS21]: in that paper, the **add** query only supports coefficients $c_1 = c_2 = 1$. This “enhanced **add** query” only strengthens the model and brings it more in line with other formulations of the GGM (such as [Zha22]).

3.1.2 Modeling the attack on Schnorr in the EC-GCM

In the EC-GGM model, the generator \mathcal{G} is encoded as $\pi(1)$ and the public key \mathcal{D} is encoded as $\pi(d)$ for randomly chosen $d \in \mathbb{Z}_q$. These encodings of \mathcal{G} and \mathcal{D} are given to the adversary at the start of the signing attack game.

The adversary then makes a sequence of queries to both the group and signing oracles. The signing oracle on a message m itself works as usual, generating $r \in \mathbb{Z}_q$ at random, but it uses the group oracle to compute the encoding of $\mathcal{R} = r\mathcal{G}$. After that, the signing oracle computes $h \leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m) \in \mathbb{Z}_q$ and $z \leftarrow r + hd$, and then gives the signature (\mathcal{R}, z) to the adversary.

At the end of the signing attack game, the adversary outputs a forgery (\mathcal{R}^*, z^*) on a message m^* . The signature is then verified using the verification algorithm, computing $h^* \leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}^* \rangle \parallel m^*) \in \mathbb{Z}_q$ and checking that $z^*\mathcal{G} = \mathcal{R}^* + h^*\mathcal{D}$ using the group oracle. WLOG, we may assume that the adversary has already performed this check and made

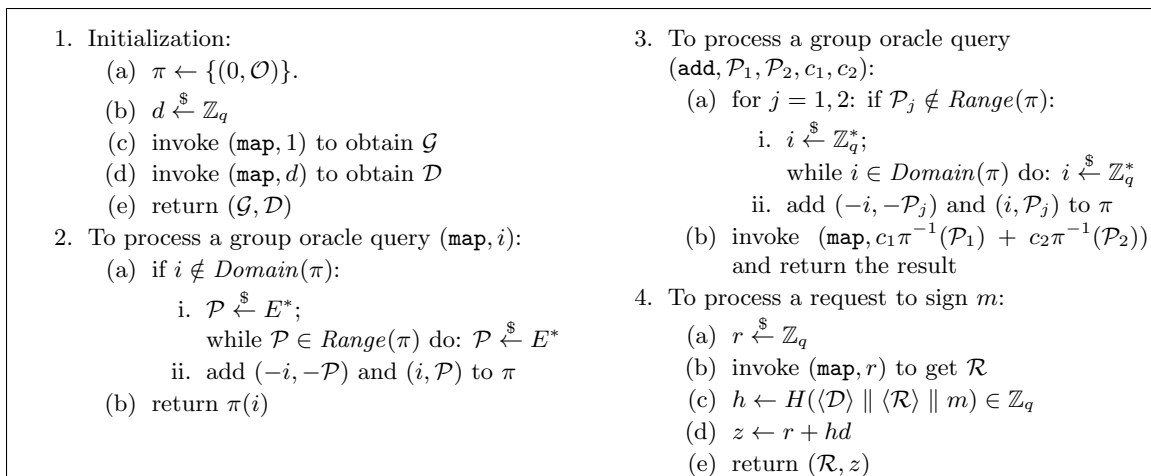


Figure 2: Lazy-Sim

the corresponding calls to the group oracle. The adversary wins the signing attack game if (\mathcal{R}^*, z^*) is a valid signature on m^* and m^* was not submitted as an input to the signing oracle.

We let N_{sig} be a bound on the number of signing queries made by the adversary, and N_{grp} be a bound on the number of group oracle queries made by the adversary. For simplicity, we assume that N_{grp} includes the group oracle queries made in the initialization step and in the verification step of the adversary’s forgery attempt. We let N be a bound on the number of group oracle and signing queries made by during the attack. Later in the paper, we will consider scenarios where the adversary also makes queries to a random oracle, and in these scenarios, N will also bound the number of random oracle queries as well.

A lazy simulation of the signature attack game. Instead of choosing the encoding function π at random at the beginning of the attack game, we can lazily construct π a bit at a time. That is, we represent π as a set of pairs (i, \mathcal{P}) which grows over time — such a pair (i, \mathcal{P}) represents the relation $\pi(i) = \mathcal{P}$. Here, we give the entire logic for both the group and signing oracles in the forgery attack game. Figure 2 gives the details of **Lazy-Sim**. This is exactly the same as the lazy simulator in Fig. 2 in [GS21], except for the logic for processing signing requests, which has been changed to Schnorr signatures instead of ECDSA signatures (and the “enhanced **add** queries”)

This lazy simulation is perfectly faithful. Specifically, the advantage of any adversary in the signature attack game using this lazy simulation of the group oracle is identical to that using the group oracle as originally defined.

A symbolic simulation of the signature attack game. We now define a **symbolic** simulation of the attack game. The essential difference in this game is that $\text{Domain}(\pi)$ will now consist of polynomials of the form $a + bD$, where $a, b \in \mathbb{Z}_q$ and D is a variable (or indeterminant). Here, D symbolically represents the value of d . Note that π will otherwise still satisfy all of the requirements of an encoding function. Figure 3 gives the details of **Symbolic-Sym**. This is exactly the same as the lazy simulator in Fig. 3 in [GS21], except

<ol style="list-style-type: none"> 1. Initialization: <ol style="list-style-type: none"> (a) $\pi \leftarrow \{(0, \mathcal{O})\}$. (b) invoke $(\mathbf{map}, 1)$ to obtain \mathcal{G} (c) invoke $(\mathbf{map}, \mathcal{D})$ to obtain \mathcal{D} (d) return $(\mathcal{G}, \mathcal{D})$ 2. To process a group oracle query (\mathbf{map}, i): <ol style="list-style-type: none"> (a) if $i \notin \text{Domain}(\pi)$: <ol style="list-style-type: none"> i. $\mathcal{P} \xleftarrow{\\$} E^*$; if $\mathcal{P} \in \text{Range}(\pi)$ then abort ii. add $(-i, -\mathcal{P})$ and (i, \mathcal{P}) to π (b) return $\pi(i)$ 	<ol style="list-style-type: none"> 3. To process a group oracle query $(\mathbf{add}, \mathcal{P}_1, \mathcal{P}_2)$: <ol style="list-style-type: none"> (a) for $j = 1, 2$: if $\mathcal{P}_j \notin \text{Range}(\pi)$: <ol style="list-style-type: none"> i. $i \xleftarrow{\\$} \mathbb{Z}_q^*$; if $i \in \text{Domain}(\pi)$ then abort ii. add $(-i, -\mathcal{P}_j)$ and (i, \mathcal{P}_j) to π (b) invoke $(\mathbf{map}, c_1\pi^{-1}(\mathcal{P}_1) + c_2\pi^{-1}(\mathcal{P}_2))$ and return the result 4. To process a request to sign m: <ol style="list-style-type: none"> (a) $\mathcal{R} \xleftarrow{\\$} E, z \xleftarrow{\\$} \mathbb{Z}_q$ (b) $h \leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m) \in \mathbb{Z}_q$ (c) $r \leftarrow z - h\mathcal{D}$ (d) if $r \in \text{Domain}(\pi)$ or $\mathcal{R} \in \text{Range}(\pi)$ then abort (e) add $(-r, -\mathcal{R})$ and (r, \mathcal{R}) to π (f) return (\mathcal{R}, z)
--	--

Figure 3: Symbolic-Sim

for the logic for processing signing requests (and the “enhanced **add** queries”).

Essentially, the signing oracle in the symbolic simulation (i) chooses $\mathcal{R} \in E$ and $z \in \mathbb{Z}_q$ at random, (ii) sets $r \leftarrow z - h\mathcal{D}$, where $h = H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m)$, (iii) “programs” π so that $\pi(r) = \mathcal{R}$ and $\pi(-r) = -\mathcal{R}$, and (iv) returns the signature (\mathcal{R}, z) .

The following lemma is fairly straightforward, and can be proved along the same lines as Lemma 1 in [GS21].

Lemma 1. *The difference between the adversary’s forging advantage in the Lazy-Sim and Symbolic-Sim games is $O(N^2/q)$.*

Indeed, throughout this paper, we hew closely to the general strategies developed in [GS21], one of which is to carry out the generic group analysis in a modular fashion, moving first from the real attack to a symbolic simulation of the attack, and then to finish off the analysis in this symbolic simulation. The move from real attack to symbolic simulation is usually straightforward and fairly mechanical, and allows us to then focus fashion on the “meat” of the proof in a more intuitive fashion.

3.1.3 Proving security in the EC-GGM

By virtue of Lemma 1, it suffices to prove the security of Schnorr in the Symbolic-Sim game. We do this by reducing the security to specific preimage resistance security properties of H .

Assume the adversary’s forgery is the signature (\mathcal{R}^*, z^*) on the message m^* . Suppose $\pi^{-1}(\mathcal{R}^*) = a + b\mathcal{D}$. By the verification equation, we must also have $\pi^{-1}(\mathcal{R}^*) = z^* - h^*\mathcal{D}$. Let $h^* := H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}^* \rangle \parallel m^*)$. Then we must have $a = z^*$ and $b = -h^*$.

Type I forgery: $\mathcal{R}^* = \pm\mathcal{R}$ for some \mathcal{R} output by the signing oracle.

Let $\mathcal{R}^* = \epsilon\mathcal{R}$, with $\epsilon \in \{\pm 1\}$. Suppose m was the input to signing oracle that produced the signature (\mathcal{R}, z) , and let $h := H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R} \rangle \parallel m)$. Then we must have

$$z^* - h^*\mathcal{D} = \epsilon(z - h\mathcal{D}).$$

In particular, $h^* = \epsilon h$.

In this case, the adversary must essentially win a UOWHF-like attack on H , which we call **Preimage Attack I** — details below.

Type II forgery: not type I and $h^* \neq 0$.

Since $b = -h^* \neq 0$, the group element \mathcal{R}^* was generated at random as the result of a group oracle query made by the adversary. (Note that the assumption $h^* \neq 0$ is used here to rule out the possibility of \mathcal{R}^* being cooked up directly by the adversary, which is allowed in the EC-GGM model.)

So in this case, the adversary must essentially win a certain type of preimage attack game on H , which we call **Preimage Attack II** — details below.

Type III forgery: not type I and $h^* = 0$.

In this case, the adversary must find a preimage of zero under H , which we call **Preimage Attack III**. (Note this case does not arise in the analysis of [NSW09] because they do not allow access to π^{-1} as is done in the EC-GCM.)

3.1.4 Preimage attack games

In the above analysis, we sketched a reduction to various preimage attacks on H . Here, we state these preimage attacks in more detail. (Preimage Attacks I and II are stated in greater generality than what we need here to cover other situations we will encounter later.)

Preimage Attack I on H .

- For $k = 1, 2, \dots$, the adversary makes a *challenge query*, giving (m_k, \mathcal{D}'_k) to challenger, who responds with random \mathcal{R}_k .

Let $h_k^* = H(\langle \mathcal{D}'_k \rangle \parallel \langle \mathcal{R}_k \rangle \parallel m_k)$.

- To win, the adversary outputs k , $(m^*, \mathcal{D}^*) \neq (m_k, \mathcal{D}'_k)$, and $\epsilon \in \{\pm 1\}$ such that

$$H(\langle \mathcal{D}^* \rangle \parallel \langle \epsilon \mathcal{R}_k \rangle \parallel m^*) = \epsilon h_k^*.$$

Preimage Attack II on H .

- For $i = 1, 2, \dots$, the adversary makes a *challenge query*, giving h_i^* to challenger, who responds with random \mathcal{R}_i^* .
- To win, the adversary outputs i , (m^*, \mathcal{D}^*) , and $\epsilon \in \{\pm 1\}$ such that

$$H(\langle \mathcal{D}^* \rangle \parallel \langle \epsilon \mathcal{R}_i^* \rangle \parallel m^*) = \epsilon h_i^*.$$

Preimage Attack III on H . To win, the adversary outputs a bit string x such that $H(x) = 0$.

3.1.5 Concrete security bounds

We can derive concrete security bounds for the analysis in Section 3.1.3. If an adversary \mathcal{A} has an advantage \aleph in forging a signature, then

$$\aleph = O(N^2/q + \aleph_{\text{I}} + \aleph_{\text{II}} + \aleph_{\text{III}}). \quad (1)$$

Here, \aleph_X is the advantage of an adversary \mathcal{A}_X in winning Preimage Attack X , for $X \in \{\text{I}, \text{II}, \text{III}\}$. Each \mathcal{A}_X has roughly the same running time as \mathcal{A} . Moreover, \mathcal{A}_{I} makes at most N_{sig} challenge queries and \mathcal{A}_{II} makes at most N_{grp} challenge queries. This follows from Lemma 1, together with the analysis in Section 3.1.3.

Suppose we model H as a random oracle (as well as working in the EC-GGM). In this paper, we generally assume that the output space of H is \mathbb{Z}_q . However, it is useful to consider a smaller output space as well, as this can be used to generate shorter signatures (using the standard technique where a signature consists of (h, z) , rather than (\mathcal{R}, z)).

So suppose H has an output space of size M . Also, assume that N also bounds the number of queries made by \mathcal{A} to the random oracle representing H . Note that this also bounds the number of random oracle queries made by each \mathcal{A}_X . Then (1) implies

$$\aleph = O(N^2/q + N/M). \quad (2)$$

Indeed, consider an adversary that carries out Preimage Attack I or II and makes at most N_{h} random oracle queries and N_{ch} challenge queries. Then such an adversary wins this attack with probability at most $O(N_{\text{ch}}^2/q + N_{\text{h}}/M)$. The term $O(N_{\text{ch}}^2/q)$ bounds the probability that there are collisions among the any of the N_{ch} random group elements generated by the challenger. Similarly, an adversary that carries out Preimage Attack III and makes at most N_{h} random oracle queries wins this attack with probability $O(N_{\text{h}}/M)$. The bound (2) immediately follows.

The above analysis is similar to that in [NSW09], except that they consider preimage attacks with only a single challenge, and then make a “guessing” argument to complete the reduction to the hardness of winning such a single-challenge preimage attack. This leads to somewhat artificially pessimistic security bounds.

3.2 Analysis of attack with re-randomized presignatures

We assume unbiased presignatures and no key derivation, but with presignatures re-randomized (as in Section 2.4). Our approach for designing the symbolic simulation in this setting is similar to that [GS21], specifically Fig. 7 of that paper, in which each presignature \mathcal{R}_k corresponds to a variable \mathbf{R}_k , meaning that $\pi(\mathbf{R}_k) = \mathcal{R}_k$. When a signing query on a message m_k is made that uses the presignature \mathcal{R}_k , in the symbolic simulation, the signing oracle computes

$$\delta_k \stackrel{\$}{\leftarrow} \mathbb{Z}_q, \quad \mathcal{R}'_k \stackrel{\$}{\leftarrow} E, \quad z_k \stackrel{\$}{\leftarrow} \mathbb{Z}_q, \quad h_k \leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}'_k \rangle \parallel m_k).$$

Before returning the signature (\mathcal{R}'_k, z_k) and the tweak δ_k , the signing oracle programs π so that $\pi(\mathbf{R}_k + \delta_k) = \mathcal{R}'_k$, and then substitutes

$$\mathbf{R}_k \mapsto z_k - \delta_k - h_k \mathbf{D}$$

throughout $\text{Domain}(\pi)$. The symbolic simulation will “fail” if any of these substitutions cause $\text{Domain}(\pi)$ to “collapse” (i.e., if two distinct elements of $\text{Domain}(\pi)$ before the substitution become equal afterwards).

This same “substitution strategy” for dealing with presignatures in the GGM was used extensively in [GS21], and works equally well here. The analog of Lemma 1 above for this symbolic simulator can easily be proven along the same lines as Lemma 2 in [GS21].

Type I forgery: $\mathcal{R}^* = \pm\mathcal{R}$ for some \mathcal{R} output by signing oracle.

This is handled exactly the same as Type I in the basic attack in Section 3.1.3.

Type II forgery: not type I and $h^* \neq 0$.

By some simple case analysis, we can assume that \mathcal{R}^* was randomly generated in processing a group oracle query made by the adversary. Suppose that initially

$$\pi^{-1}(\mathcal{R}^*) = a + bD + \sum_k c_k R_k,$$

where the c_k ’s are all nonzero. If the sum on k is empty, this can be handled the same as Type II in the basic attack in Section 3.1.3. Otherwise, in order for the forgery to be valid, the R_k variables need to be eliminated by substitution, to end up with

$$\pi^{-1}(\mathcal{R}^*) = z^* - h^*D.$$

Suppose that all but one has been eliminated, say R_ℓ , so that at that time,

$$\pi^{-1}(\mathcal{R}^*) = a' + b'D + c_\ell R_\ell.$$

The last substitution is $R_\ell \mapsto z_\ell - \delta_\ell - h_\ell D$, yielding

$$\pi^{-1}(\mathcal{R}^*) = \underbrace{\{a' + c_\ell(z_\ell - \delta_\ell)\}}_{=z^*} + \underbrace{\{b' - c_\ell h_\ell\}}_{=-h^*} D.$$

So in this case, the adversary can win a certain type of preimage attack game on H , which we call **Preimage Attack II'** — details below.

Type III forgery: not type I and $h^* = 0$.

This is handled exactly the same as Type III in the basic attack in Section 3.1.3.

3.2.1 Another preimage attack

We describe in more detail the preimage attack used in the above security analysis. This attack is stated in greater generality than what we need here to cover other situations we will encounter later.

Preimage Attack II' on H .

- The challenger gives a collection $\{\mathcal{R}_i^*\}_{i=1}^{N_{\text{ch}}}$ of random challenges to the adversary (each \mathcal{R}_i^* is a random element of E).
- For $k = 1, 2, \dots$, the adversary submits a *completion query* to the challenger consisting of an index set $\mathcal{I}_k \subseteq \{1, \dots, N_{\text{ch}}\}$ that is disjoint from $\mathcal{I}_1 \cup \dots \cup \mathcal{I}_{k-1}$, along with \mathcal{D}'_k , m_k , and $\{(b_i, c_i)\}_{i \in \mathcal{I}_k}$, where each $(b_i, c_i) \in \mathbb{Z}_q \times \mathbb{Z}_q^*$.
 - The challenger generates \mathcal{R}'_k at random and returns this to the adversary.
 - Let $h_k = H(\langle \mathcal{D}'_k \rangle \parallel \langle \mathcal{R}'_k \rangle \parallel m_k)$ and $h_i^* = b_i - c_i \cdot h_k$ for $i \in \mathcal{I}_k$.
- To win, the adversary outputs $i \in \mathcal{I}$, (m^*, \mathcal{D}^*) , and $\epsilon \in \{\pm 1\}$ such that

$$H(\langle \mathcal{D}^* \rangle \parallel \langle \epsilon \mathcal{R}_i^* \rangle \parallel m^*) = \epsilon h_i^*.$$

In the above attack game, the various \mathcal{R}_i^* values correspond to outputs from the group oracle in the symbolic simulation of the signing attack, while the various \mathcal{R}'_k values correspond to the outputs of the signing oracle. The k th completion query in the above attack game corresponds to the k th signing query in the symbolic simulation of the signing attack, and the set of indices \mathcal{I}_k represents those group elements that were output by the group oracle whose last remaining presignature variable is being eliminated by substitution from this signing request.

3.2.2 Concrete security bounds

If an adversary \mathcal{A} has an advantage \aleph in forging a signature, then

$$\aleph = O(N^2/q + \aleph_{\text{I}} + \aleph_{\text{II}} + \aleph_{\text{II}'} + \aleph_{\text{III}}). \quad (3)$$

Here, \aleph_X is the advantage of an adversary \mathcal{A}_X in winning Preimage Attack X , for $X \in \{\text{I}, \text{II}, \text{II}', \text{III}\}$. Each \mathcal{A}_X has roughly the same running time as \mathcal{A} . Moreover, \mathcal{A}_{I} makes at most N_{sig} challenge queries, \mathcal{A}_{II} makes at most N_{grp} challenge queries, and $\mathcal{A}_{\text{II}'}$ receives at most N_{grp} challenges and makes at most N_{sig} completion queries.

Now suppose we model H as a random oracle with an output space of size M . Also, assume that N also bounds the number of queries made by \mathcal{A} to the random oracle representing H . Note that this also bounds the number of random oracle queries made by each \mathcal{A}_X . Then (3) implies

$$\aleph = O(N^2/q + N/M). \quad (4)$$

To see this, suppose that in Preimage Attack II', the adversary receives N_{ch} random challenges, and makes at most N_{cmp} completion queries and at most N_{h} random oracle queries. Assume no collisions among the random challenges occur. This means that for a given random oracle query of the form

$$H(\langle \mathcal{D}^* \rangle \parallel \langle \epsilon \mathcal{R}_i^* \rangle \parallel m^*), \quad (5)$$

there is a unique index i and values b_i, c_i , the h_k such that

$$H(\langle \mathcal{D}^* \rangle \parallel \langle \epsilon \mathcal{R}_i^* \rangle \parallel m^*) = b_i - c_i \cdot h_k$$

must hold in order for the random oracle query (5) to lead to a win. Here, k is the index of the completion query which included i in \mathcal{I}_k . Moreover, assuming that $\mathcal{D}'_k \neq \pm \mathcal{R}_i^*$ and the adversary did not happen to query $H(\langle \mathcal{D}'_k \rangle \parallel \langle \mathcal{R}'_k \rangle \parallel m_k)$ before the k th completion query was made, the value $h_k := H(\langle \mathcal{D}'_k \rangle \parallel \langle \mathcal{R}'_k \rangle \parallel m_k)$ is random and independent of $H(\langle \mathcal{D}^* \rangle \parallel \langle \epsilon \mathcal{R}_i^* \rangle \parallel m^*)$, b_i , and c_i , and so the random oracle query (5) leads to a win with probability at most $1/M$. From this, we see that the adversary wins the attack with probability at most

$$O((N_{\text{ch}} + N_{\text{cmp}} + N_{\text{h}})^2/q + N_{\text{h}}/M).$$

The bound (4) now follows.

3.2.3 Variations

If we use biased presignatures, then effectively \mathcal{R}_k gets replaced by $u_k \mathcal{R}_k + u'_k \mathcal{G}$ just before signing a message, where $u_k \neq 0$ and u'_k are explicitly given by the adversary. So in the symbolic simulation, the signing oracle programs π so that $\pi(u_k \mathbf{R}_k + u'_k + \delta_k) = \mathcal{R}'_k$ and substitutes

$$\mathbf{R}_k \mapsto u_k^{-1}(z_k - u'_k - \delta_k - h_k \mathbf{D}).$$

The general argument does not really change at all. If we use additive key derivation, deriving $\mathcal{D}'_k := \mathcal{D} + e_k \mathcal{G}$, then this substitution becomes

$$\mathbf{R}_k \mapsto u_k^{-1}(z_k - h_k e_k - u'_k - \delta_k - h_k \mathbf{D}).$$

The same concrete security bounds in Section 3.2.2 also hold here.

3.3 Re-randomizing presignatures via hashing

We now analyze the security of Schnorr signatures in the GGM with presignatures that are re-randomized via hashing, as discussed in Section 2.5. Here, we will model Δ as a random oracle with output space \mathbb{Z}_q . We will also model as H as random oracle.

We will assume unbiased presignatures for now and later examine biased presignatures as well as additive key derivation.

So to model this situation, in the symbolic simulation, we introduce variables \mathbf{R}_k and \mathbf{S}_k , where $\pi(\mathbf{R}_k) = \mathcal{R}_k$ and $\pi(\mathbf{S}_k) = \mathcal{S}_k$. When a signing query on a message m_k is made that uses the presignature $(\mathcal{R}_k, \mathcal{S}_k)$, a preliminary computation

$$\begin{aligned} \delta_k &\leftarrow \Delta(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}_k \rangle \parallel \langle \mathcal{S}_k \rangle \parallel \langle k \rangle \parallel m_k), \\ \mathcal{R}'_k &\leftarrow \mathcal{R}_k + \delta_k \mathcal{S}_k, \\ h_k &\leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \mathcal{R}'_k \rangle \parallel m_k) \end{aligned}$$

is made. WLOG, we can assume that the adversary has computed these values himself. Next, the signing oracle generates z_k at random. Before returning the signature (\mathcal{R}'_k, z_k) , the signing oracle also substitutes

$$\mathbf{R}_k \mapsto z_k - \delta_k \mathbf{S}_k - h_k \mathbf{D} \tag{6}$$

throughout $\text{Domain}(\pi)$.

We leave it to the reader to verify that the analog of Lemma 1 above holds as well for this symbolic simulator.

As usual, suppose the forgery is a signature (\mathcal{R}^*, z^*) on a message m^* . Note that the signing oracle does not generate any new group elements, so we do not need to categorize forgeries as we did before. In particular, we may assume that \mathcal{R}^* was randomly generated by the a group oracle query made by the adversary. Suppose that initially

$$\pi^{-1}(\mathcal{R}^*) = a + bD + \sum_k (c_k R_k + d_k S_k), \quad (7)$$

where each (c_k, d_k) is nonzero (as a pair). Note that the constants a, b , and c_k, d_k for all indices k are fixed before \mathcal{R}^* is randomly generated. In order for this forgery to be valid, the R_k variables need to be eliminated by substitution, to end up with

$$\pi^{-1}(\mathcal{R}^*) = z^* - h^*D.$$

In fact, after substitution, we have

$$\pi^{-1}(\mathcal{R}^*) = \underbrace{\{a + \sum_k c_k z_k\}}_{=z^*} + \underbrace{\{b - \sum_k c_k h_k\}}_{=-h^*} D + \sum_k \underbrace{\{d_k - c_k \delta_k\}}_{=0} S_k. \quad (8)$$

For the forgery to be valid, we must have $d_k - c_k \delta_k = 0$ for each index k . If $c_k = 0$, then $d_k = 0$ as well, as since we are assuming that $(c_k, d_k) \neq (0, 0)$, this implies $c_k \neq 0$. In particular,

$$\delta_k = \frac{d_k}{c_k}$$

for each index k . This means that at the time we generate \mathcal{R}^* , we can inspect the queries to the random oracle Δ to find for each index k an input

$$(\langle D \rangle \parallel \langle \mathcal{R}_k \rangle \parallel \langle S_k \rangle \parallel \langle k \rangle \parallel m_k)$$

to Δ that yields the output d_k/c_k . If the forgery is to be valid, then except with negligible probability, the adversary must have already made such a query and it will be unique. Thus, at the time we generate \mathcal{R}^* at random, the inputs to H that determine the h_k 's have already been determined. Thus, in order for the adversary to find m^* , he must essentially win Preimage Attack II on H , as in Section 3.1, but where now H is modeled as a random oracle.

Concrete security bounds. To make the above analysis concrete, we have to calculate the probability that the above inspection process fails. For it to fail, it means that either (a) the adversary finds a collision in Δ , or (b) for some \mathcal{R}^* output by the group oracle, for each k in (7) for which the adversary has not already made a relevant query to Δ whose output hits d_k/c_k , the adversary must make such a query at a later time whose output (by pure luck) hits d_k/c_k . The probability that (a) or (b) occurs is at most $O(N^2/q)$ — more precisely, (a) occurs with probability $O(N_h^2)$ and (b) occurs with probability $O(N_{\text{grp}}N_h)$. From this, it follows that if H is modeled as a random oracle with an output space of size M , the adversary's forging advantage is $O(N^2/q + N/M)$.

3.3.1 Variations

If we use biased presignatures, then effectively \mathcal{R}_k gets replaced by $\tilde{\mathcal{R}}_k := u_k \mathcal{R}_k + u'_k \mathcal{G}$ and \mathcal{S}_k gets replaced by $\tilde{\mathcal{S}}_k := v_k \mathcal{S}_k + v'_k \mathcal{G}$, where $u_k \neq 0$, $u'_k, v_k \neq 0$, and v'_k are explicitly given by the adversary. The input to Δ used to derive δ_k is then

$$(\langle \mathcal{D} \rangle \parallel \langle \tilde{\mathcal{R}}_k \rangle \parallel \langle \tilde{\mathcal{S}}_k \rangle \parallel \langle k \rangle \parallel m_k) \quad (9)$$

The substitution (6) then becomes

$$\mathbf{R}_k \mapsto u_k^{-1}(z_k - u'_k - \delta_k v'_k - \delta_k v_k \mathbf{S}_k - h_k \mathcal{D}) \quad (10)$$

and (8) becomes

$$\begin{aligned} \pi^{-1}(\mathcal{R}^*) = & \underbrace{\left\{ a + \sum_k c_k u_k^{-1} (z_k - u'_k - \delta_k v_k) \right\}}_{=z^*} + \underbrace{\left\{ b - \sum_k c_k u_k^{-1} h_k \right\}}_{=-h^*} \mathcal{D} + \\ & \sum_k \underbrace{\left\{ d_k - c_k u_k^{-1} v_k \delta_k \right\}}_{=0} \mathbf{S}_k. \end{aligned} \quad (11)$$

The main argument does not change too much. One thing we may have to adjust is that now we are inspecting the queries to Δ , looking for inputs that output

$$\delta_k = \frac{u_k}{v_k} \frac{d_k}{c_k}.$$

However, we have to look for such inputs *before* the signing request is made that determines u_k and v_k . Nevertheless, since the biased presignatures $\tilde{\mathcal{R}}$ and $\tilde{\mathcal{S}}$ are input to Δ , we can actually use the GGM to determine u_k and v_k . That is, we are looking for inputs to Δ of the form (9) such that

- $\pi^{-1}(\tilde{\mathcal{R}}) = (u_k \mathbf{R}_k + \dots)$ and $\pi^{-1}(\tilde{\mathcal{S}}) = (v_k \mathbf{S}_k + \dots)$, and
- the output is

$$\frac{u_k}{v_k} \frac{d_k}{c_k}.$$

Indeed, we can assume WLOG that for all queries to Δ , the corresponding group element encodings $\tilde{\mathcal{R}}$ and $\tilde{\mathcal{S}}$ are already in $\text{Range}(\pi)$. Moreover, $\pi^{-1}(\tilde{\mathcal{R}})$ and $\pi^{-1}(\tilde{\mathcal{S}})$ need to be of this form if they are to be of the required form $u_k \mathbf{R}_k + u'_k$ and $v_k \mathbf{S}_k + v'_k$ at the time the actual signing request is made (none of the substitutions performed between now and then will affect the coefficients of \mathbf{R}_k or \mathbf{S}_k).

If we use additive key derivation, deriving $\mathcal{D}'_k := \mathcal{D} + e_k \mathcal{G}$, then we also need to include \mathcal{D}'_k as input to Δ , in place of \mathcal{D} . The substitution (10) then becomes

$$\mathbf{R}_k \mapsto u_k^{-1}(z_k - u'_k - h_k e_k - \delta_k v'_k - \delta_k v_k \mathbf{S}_k - h_k \mathcal{D}) \quad (12)$$

and the constant term in (11) is adjusted accordingly. Again, the main argument does not change too much.

Concrete security bounds. For each of these variations, the same security bounds as above: if H is modeled as a random oracle with an output space of size M , the adversary’s forging advantage is $O(N^2/q + N/M)$.

4 Batch randomness extraction

In a typical distributed algorithm for creating a sharing of a random secret $r \in \mathbb{Z}_q$ together with the public value $\mathcal{R} = r\mathcal{G} \in E$, each party will run a VSS protocol as the “dealer”, and the results of these are then combined (usually just added together) to create \mathcal{R} . However, there are well known “batching” techniques that allow us to extract many such \mathcal{R} ’s for the price of one.

We assume n parties P_1, \dots, P_n , up to f of which may be corrupt. We assume *static* corruptions and $n > 3f$ (which is optimal in the asynchronous communication setting).

We assume an ideal *transcript-building functionality* in which, in the first phase, each party inputs $r^{(j)} \in \mathbb{Z}_q$ while $\mathcal{R}^{(j)} = r^{(j)}\mathcal{G} \in E$ is given to the adversary. This phase corresponds to a protocol in which each party P_j acts as a dealer in a VSS protocol to share $r^{(j)}$ and publish $\mathcal{R}^{(j)}$ (for example, one may use a VSS protocol based on Feldman commitments). In the second phase, the adversary chooses a subset set \mathcal{J} of $\{1, \dots, n\}$ of size $n - f$, and the ideal functionality then outputs to each party \mathcal{J} and $\{\mathcal{R}^{(j)}\}_{j \in \mathcal{J}}$. This phase corresponds to running a consensus protocol to agree on the set \mathcal{J} .

In our application of this transcript-building functionality, honest parties P_j will always input random $r^{(j)} \in \mathbb{Z}_q$, while corrupt parties may input arbitrary values. To make use of such an ideal functionality in a signing protocol, we also need to augment it with additional capabilities, corresponding to running a DKG to generate random $d \in \mathbb{Z}_q$ and publish $\mathcal{D} := d\mathcal{G}$, and to opening specified linear combinations of d and $\{r^{(j)}\}_{j \in \mathcal{J}}$.

The basic idea of this batching technique is to use a Vandermonde matrix

$$V \in \mathbb{Z}_q^{(n-2f) \times (n-f)},$$

where the j th column is

$$V_j = (1, \alpha_j, \dots, \alpha_j^{n-2f-1})^\top$$

and all the α_j ’s are distinct. Given the group elements $\mathcal{R}^{(j_1)}, \dots, \mathcal{R}^{(j_{n-f})}$ as output by the transcript-building functionality, of which at least $n - 2f$ are randomly generated by honest parties and the rest are adversarially chosen, we compute

$$\begin{pmatrix} \bar{\mathcal{R}}^{(1)} \\ \vdots \\ \bar{\mathcal{R}}^{(n-2f)} \end{pmatrix} = V \begin{pmatrix} \mathcal{R}^{(j_1)} \\ \vdots \\ \mathcal{R}^{(j_{n-f})} \end{pmatrix}. \quad (13)$$

The idea is then that in the Schnorr signature attack game, we can use the derived presignatures $\bar{\mathcal{R}}^{(1)}, \dots, \bar{\mathcal{R}}^{(n-2f)}$ for $n - 2f$ signing queries. Since $n - 2f > n/3$, this essentially gives us a batch of $\Omega(n)$ presignatures for the price of one.

Note that instead of a Vandermonde matrix, any $(n - 2f) \times (n - f)$ matrix may be used that has the property that any subset of $n - 2f$ columns is linearly independent. Such a matrix is called *super-invertible* in [HN06]. While there are various constructions, a

Vandermonde matrix is convenient as it may lead to more efficient implementations in some settings. For example, we can choose the α_j 's to be the images of small integers $(1, 2, \dots$ or $\pm 1, \pm 2, \dots)$, in which case we can compute the matrix-vector product (13) using $O(n^2)$ “short” scalar multiplications. For moderately sized n (with $n = O(\log q)$), this is probably the most efficient algorithm. For larger n , asymptotically fast algorithms may be employed, which exploit the fact that V is a structured matrix.

This general technique of batch randomness extraction first appeared in [HN06] in a somewhat different setting. It was first proposed in the context of threshold Schnorr signatures in [BHK⁺23].

In what follows, we assume there are many batches, each indexed by k , so we denote the initial presignatures (as output by one instance of the transcript-building functionality) as $\mathcal{R}_k^{(j)}$ for $j \in \mathcal{J}_k$, and the derived presignatures as $\bar{\mathcal{R}}_k^{(i)}$ for $i \in \mathcal{I} := \{1, \dots, n - 2f\}$. So a particular signing query specifies a pair of indices (k, i) so that a message $m_{k,i}$ is signed using the derived presignature $\bar{\mathcal{R}}_k^{(i)}$.

4.1 Re-randomized presignatures

Just as in Section 2.4, when signing a message $m_{k,i}$ we can re-randomize the derived presignature $\bar{\mathcal{R}}_k^{(i)}$, replacing it with

$$\hat{\mathcal{R}}_k^{(i)} := \bar{\mathcal{R}}_k^{(i)} + \delta_{k,i}\mathcal{G},$$

where $\delta_{k,i}$ is generated (by a Random Beacon) only after the signing request has been made.

As in Section 2.4, we give an efficient reduction from this scheme to the security of the interactive Schnorr identification scheme, modeling H as a random oracle. As usual, the key is to show how to simulate the signing queries. Here, we are working in a hybrid model with the transcript-building ideal functionality. Let \mathcal{H} denote the set of indices of the honest parties and \mathcal{C} denote the set of indices of the corrupt parties. In the k th instance of this functionality, for each $j \in \mathcal{H}$, our simulator will generate $\zeta_k^{(j)}, \eta_k^{(j)} \in \mathbb{Z}_q$ at random, and then compute

$$\mathcal{R}_k^{(j)} \leftarrow \zeta_k^{(j)}\mathcal{G} - \eta_k^{(j)}\mathcal{D}.$$

When this ideal functionality finishes, for each $i \in \mathcal{I}$, we have

$$\bar{\mathcal{R}}_k^{(i)} = \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \mathcal{R}_k^{(j)} + \mu_{k,i} \mathcal{G}, \tag{14}$$

where $\mathcal{H}_k \subseteq \mathcal{H}$ is a set of size at least $n - 2f$, and each $\lambda_{k,i}^{(j)}, \mu_{k,i} \in \mathbb{Z}_q$. The values \mathcal{H}_k and $\lambda_{k,i}^{(j)}, \mu_{k,i}$ are all known to the simulator (note that the values $\mu_{k,i}$ can be computed by the simulator because the adversary must input the values $r_k^{(j)}$ for $j \in \mathcal{J} \setminus \mathcal{H}$ directly to the simulator via the transcript-building interface). Moreover, the matrix

$$A_k = (\lambda_{k,i}^{(j)})_{i \in \mathcal{I}, j \in \mathcal{H}_k}$$

which has $n - 2f$ rows and $|\mathcal{H}_k| \geq n - 2f$ columns, has full rank.

Therefore, we have

$$\begin{aligned}\bar{\mathcal{R}}_k^{(i)} &= \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \mathcal{R}_k^{(j)} + \mu_{k,i} \mathcal{G} \\ &= \{\mu_{k,i} + \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \zeta_k^{(j)}\} \mathcal{G} - \{\sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \eta_k^{(j)}\} \mathcal{D}.\end{aligned}$$

This implies that the re-randomized presignature is

$$\hat{\mathcal{R}}_k^{(i)} = \bar{\mathcal{R}}_k^{(i)} + \delta_{k,i} \mathcal{G} = \{\delta_{k,i} + \mu_{k,i} + \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \zeta_k^{(j)}\} \mathcal{G} - \{\sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \eta_k^{(j)}\} \mathcal{D}$$

So to sign a message $m_{k,i}$, the simulator will choose $\delta_{k,i}$ at random, and output the signature $(\hat{\mathcal{R}}_k^{(i)}, z_{k,i})$ along with the value $\delta_{k,i}$, where

$$z_{k,i} := \delta_{k,i} + \mu_{k,i} + \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \zeta_k^{(j)},$$

and, additionally, programs the random oracle so that

$$H(\langle \mathcal{D} \rangle \parallel \langle \hat{\mathcal{R}}_k^{(i)} \rangle \parallel m_{k,i}) := h_{k,i},$$

where

$$h_{k,i} := \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \eta_k^{(j)}.$$

The fact that the matrix A_k has full rank and that the $\eta_k^{(j)}$'s are random and independent (of each other as well as everything in the adversary's view, including the values $\lambda_{k,i}^{(j)}$) means that the $h_{k,i}$'s are also random and independent, so the output of the random oracle has the right distribution.

The above analysis carries over in a straightforward way to handle additive key derivation. If the effective key for a given signing request is $\mathcal{D}'_{k,i} = \mathcal{D} + e_{k,i} \mathcal{G}$, then in the above simulation, we have to subtract $e_{k,i} h_{k,i}$ from the value $z_{k,i}$ we originally computed, and program the random oracle at the point corresponding to $\mathcal{D}'_{k,i}$, rather than \mathcal{D} .

It is a curious fact that the above proof relies crucially on the assumption that the output space of H is all of \mathbb{Z}_q . However, this appears to just be an artifact of the proof, as we can prove security in the GGM without this restriction (see below in Section 4.2).

Relation to SPRINT. Our analysis here highlights and presents in a more simplified and modular form ideas that are already in present in the SPRINT protocol from [BHK⁺23]. Note that in SPRINT, rather than the $\delta_{k,i}$'s being the output of a Random Beacon, they are actually the output of a hash function modeled as a random oracle. In fact, in SPRINT, the inputs to this random oracle includes a batch of messages $\{m_{k,i}\}_i$ to be signed using the corresponding batch of derived presignatures $\{\bar{\mathcal{R}}_{k,i}\}_i$, so that all of $\delta_{k,i}$'s for this entire batch are generated at once. However, the analysis really calls for a Random Beacon, rather than a random oracle. Indeed, the security theorem proved in [BHK⁺23] actually only analyzes an attack with just a single batch of signing requests. It works by guessing which random

oracle query represents the Random Beacon, and this (among other things) results in a quite inefficient security reduction. To be useful, one must model an attack in which many batches of signing requests are processed. While [BHK⁺23] is mute on this point, it would appear that their theorem could be extended to prove the security in an attack in which batches of signing requests are processed sequentially. However, this means that only a *single* batch of unused presignatures can be outstanding at a time: if there are many such batches of unused presignatures, the same attack as described in Section 2.5 can be carried out (using one presignature per batch).

This seems to somewhat defeat the purpose of presignatures — the goal is to build up a large cache of presignatures in periods of low demand, so as to be able to quickly process bursts of signing requests in periods of high demand. However, with SPRINT, after a single batch of presignatures is produced, it must be consumed by processing a corresponding batch of signing requests before the next batch of presignatures can be produced. Regardless of the size of these batches, latency and/or throughput will be adversely affected by this restriction. For example, when an individual signing request comes in, we will have to make it wait until the batch of signing requests is full, or we can process it, discarding any unused presignatures in the batch and initiating production of the next batch of presignatures. In the latter case, by discarding unused presignatures, the overall throughput of the system is reduced; moreover, the next signing request that comes in will have to wait for the production of that next batch of presignatures to complete.

4.2 Generic group model analysis

In Section 4.1, we analyzed the scheme with re-randomized presignatures in the random oracle model, giving a reduction to the security of Schnorr’s identification scheme. Here, we give an analysis in the generic group model.

Analogous to (14), we have

$$\bar{\mathbf{R}}_k^{(i)} = \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \mathbf{R}_k^{(j)} + \mu_{k,i}, \quad (15)$$

for $i \in \mathcal{I} = \{1, \dots, n-2f\}$, where the $\mathbf{R}_k^{(j)}$ and $\bar{\mathbf{R}}_k^{(i)}$ are variables which symbolically represent the discrete logarithms of the group elements $\mathcal{R}_k^{(j)}$ and $\bar{\mathcal{R}}_k^{(i)}$.

It is convenient to extend (15) to all $i \in \mathcal{I}_k := \{1, \dots, |\mathcal{H}_k|\} \supseteq \mathcal{I}$. To do this, we can simply extend the j th column of the Vandermonde matrix V to include higher powers of α_j . This defines a bijective \mathbb{Z}_q -linear map between the \mathbb{Z}_q -vector spaces $\mathbb{Z}_q + \sum_{j \in \mathcal{H}_k} \mathbb{Z}_q \mathbf{R}_k^{(j)}$ and $\mathbb{Z}_q + \sum_{i \in \mathcal{I}_k} \mathbb{Z}_q \bar{\mathbf{R}}_k^{(i)}$ (which acts as the identity on \mathbb{Z}_q).

We can even extend this further, defining $\bar{\mathbf{R}}_k^{(i)}$ for all $i \in \mathcal{I}' := \{1, \dots, |\mathcal{H}|\}$, by simply defining $\bar{\mathbf{R}}_k^{(i)}$ for each $i \in \mathcal{I}' \setminus \mathcal{I}_k$ to be one of the variables $\mathbf{R}_k^{(j)}$ for $j \in \mathcal{H} \setminus \mathcal{H}_k$. This defines a bijective \mathbb{Z}_q -linear map between the \mathbb{Z}_q -vector spaces $\mathbb{Z}_q + \sum_{j \in \mathcal{H}} \mathbb{Z}_q \mathbf{R}_k^{(j)}$ and $\mathbb{Z}_q + \sum_{i \in \mathcal{I}'} \mathbb{Z}_q \bar{\mathbf{R}}_k^{(i)}$.

Note that for a given k , this bijective map is only defined after the corresponding transcript-building functionality terminates. In the symbolic simulation, when we this instance of the transcript-building functionality terminates, we use this bijective map to substitute, throughout $\text{Domain}(\pi)$, each variable $\mathbf{R}_k^{(j)}$, for $j \in \mathcal{H}$, by its corresponding value in $\mathbb{Z}_q + \sum_{i \in \mathcal{I}'} \mathbb{Z}_q \bar{\mathbf{R}}_k^{(i)}$ under this map.

Now consider what happens at a later time in the symbolic simulation (after we have already substituted the variables $\mathbf{R}_k^{(j)}$ with the variables $\bar{\mathbf{R}}_k^{(i)}$) when we sign a message $m_{k,i}$ using the derived presignature $\bar{\mathcal{R}}_k^{(i)}$, which is re-randomized as $\hat{\mathcal{R}}_k^{(i)} := \bar{\mathcal{R}}_k^{(i)} + \delta_{k,i}\mathcal{G}$. Here, $\delta_{k,i}$ is generated by a Random Beacon only after the signing request has been made. Here, the symbolic simulation chooses $z_{k,i}, \delta_{k,i} \in \mathbb{Z}_q$ and $\hat{\mathcal{R}}_k^{(i)} \in E$ at random, programs π so that $\pi(\bar{\mathbf{R}}_k^{(i)} + \delta_{k,i}) = \hat{\mathcal{R}}_k^{(i)}$, and makes the substitution

$$\bar{\mathbf{R}}_k^{(i)} \mapsto z_{k,i} - \delta_{k,i} - h_{k,i}\mathbf{D}.$$

throughout $\text{Domain}(\pi)$.

After defining the symbolic simulation in this way, the rest of the argument is essentially the same as in Section 3.2. The argument is also easily adapted to handle additive key derivation. We get essentially the same reduction to the various preimage problems as in Section 3.2. The same concrete security bounds in Section 3.2.2 also hold here.

4.3 Re-randomizing presignatures via hashing

In this section, we present a protocol that combines the technique of re-randomization via hashing (as in Section 3.3) with batch randomness extraction. We analyze the protocol in the GGM plus ROM.

We first extend the transcript-building functionality to a *transcript-pair-building functionality*. In this functionality, each party P_j inputs $(r^{(j)}, s^{(j)}) \in \mathbb{Z}_q \times \mathbb{Z}_q$, with $(\mathcal{R}^{(j)}, \mathcal{S}^{(j)}) = (r^{(j)}\mathcal{G}, s^{(j)}\mathcal{G}) \in E \times E$ being given to the adversary. As before, the adversary chooses a subset set \mathcal{J} of $\{1, \dots, n\}$ of size $n - f$, and the ideal functionality then outputs to each party \mathcal{J} and $\{(\mathcal{R}^{(j)}, \mathcal{S}^{(j)})\}_{j \in \mathcal{J}}$. As before, honest parties P_j will always input random $(r^{(j)}, s^{(j)}) \in \mathbb{Z}_q \times \mathbb{Z}_q$, while corrupt parties may input arbitrary values.

We use the same Vandermonde matrix V to compute

$$\begin{pmatrix} \bar{\mathcal{R}}^{(1)} \\ \vdots \\ \bar{\mathcal{R}}^{(n-2f)} \end{pmatrix} = V \begin{pmatrix} \mathcal{R}^{(j_1)} \\ \vdots \\ \mathcal{R}^{(j_{n-f})} \end{pmatrix}$$

as above, as well as

$$\begin{pmatrix} \bar{\mathcal{S}}^{(1)} \\ \vdots \\ \bar{\mathcal{S}}^{(n-2f)} \end{pmatrix} = V \begin{pmatrix} \mathcal{S}^{(j_1)} \\ \vdots \\ \mathcal{S}^{(j_{n-f})} \end{pmatrix}.$$

After the transcript-pair-building functionality terminates, the honest parties will access a Random Beacon to obtain a random value ρ_k from some large set. Note that while we use a Random Beacon here, it is only needed in the offline preprocessing phase.

To sign a message $m_{k,i}$, for any k and $i \in \mathcal{I}$, the derived presignature is

$$\hat{\mathcal{R}}_k^{(i)} := \bar{\mathcal{R}}_k^{(i)} + \delta_{k,i}\bar{\mathcal{S}}_k^{(i)},$$

where

$$\delta_{k,i} := \Delta(\langle \mathcal{D} \rangle \parallel \langle k \rangle \parallel \langle i \rangle \parallel \langle \rho_k \rangle \parallel m_{k,i}).$$

As in Section 3.3, both Δ and H are modeled as a random oracles.

NOTES:

1. If we use additive key derivation, then the signing request includes $e_{k,i} \in \mathbb{Z}_q$, and we replace the public key \mathcal{D} by $\mathcal{D}'_{k,i} := \mathcal{D} + e_{k,i}\mathcal{G}$ in the calculation of $\delta_{k,i}$, and we replace the secret key d by $d + e_{i,k}$ in the signing algorithm. While we will not analyze this variation in detail, the analysis we present applies equally well to this variation.
2. We can easily apply an additional layer of batching, so that we generate many batches at once. This means we can use “batched asynchronous VSS” protocols to more efficiently generate many dealings at once (for example, as in Section 8.9 of [GS22]), leading to even more efficient protocols. In addition, we can easily implement the Random Beacon using such an asynchronous VSS protocol at essentially no extra cost.

Analogous to (15), we have

$$\bar{\mathbf{S}}_k^{(i)} = \sum_{j \in \mathcal{H}_k} \lambda_{k,i}^{(j)} \mathbf{R}_k^{(j)} + \nu_{k,i}, \quad (16)$$

for $i \in \mathcal{I} = \{1, \dots, n - 2f\}$, where $\mathbf{S}_k^{(j)}$ and $\bar{\mathbf{S}}_k^{(i)}$ are variables which symbolically represent the discrete logarithms of the group elements $\mathcal{S}_k^{(j)}$ and $\bar{\mathcal{S}}_k^{(i)}$. Just as we did in relation to (15), we can extend this to all $i \in \mathcal{I}'$. This defines a bijective \mathbb{Z}_q -linear map between the \mathbb{Z}_q -vector spaces $\mathbb{Z}_q + \sum_{j \in \mathcal{H}} \mathbb{Z}_q \mathcal{S}_k^{(j)}$ and $\mathbb{Z}_q + \sum_{i \in \mathcal{I}'} \mathbb{Z}_q \bar{\mathcal{S}}_k^{(i)}$ (which acts as the identity on \mathbb{Z}_q).

Analogous to what we did in Section 4.2, in the symbolic simulation, when this instance of the transcript-building functionality terminates, we use this bijective map to substitute, throughout $\text{Domain}(\pi)$, each variable $\mathbf{R}_k^{(j)}$, for $j \in \mathcal{H}$, by its corresponding value in $\mathbb{Z}_q + \sum_{i \in \mathcal{I}'} \mathbb{Z}_q \bar{\mathbf{R}}_k^{(i)}$ under this map, and each variable $\mathbf{S}_k^{(j)}$, for $j \in \mathcal{H}$, by its corresponding value in $\mathbb{Z}_q + \sum_{i \in \mathcal{I}'} \mathbb{Z}_q \bar{\mathbf{S}}_k^{(i)}$.

Analogous to what we did in the Generic Group Model analysis in Section 3.3, when a signing query on a message $m_{k,i}$ is made that uses the derived presignature $(\bar{\mathcal{R}}_{k,i}, \bar{\mathcal{S}}_{k,i})$, a preliminary computation

$$\begin{aligned} \delta_{k,i} &\leftarrow \Delta(\langle \mathcal{D} \rangle \parallel \langle k \rangle \parallel \langle i \rangle \parallel \langle \rho_k \rangle \parallel m_{k,i}), \\ \hat{\mathcal{R}}_{k,i} &\leftarrow \bar{\mathcal{R}}_{k,i} + \delta_{k,i} \bar{\mathcal{S}}_{k,i}, \\ h_{k,i} &\leftarrow H(\langle \mathcal{D} \rangle \parallel \langle \hat{\mathcal{R}}_{k,i} \rangle \parallel m_{k,i}) \end{aligned}$$

is made. WLOG, we can assume that the adversary has computed these values himself. Next, the signing oracle generates $z_{k,i}$ at random. Before returning the signature $(\hat{\mathcal{R}}_{k,i}, z_k)$, the signing oracle also substitutes

$$\bar{\mathbf{R}}_k^{(i)} \mapsto z_{k,i} - \delta_{k,i} \bar{\mathbf{S}}_k^{(i)} - h_{k,i} \mathbf{D} \quad (17)$$

throughout $\text{Domain}(\pi)$.

Analogous to what we did in the Generic Group Model analysis in Section 3.3, suppose the forgery is a signature (\mathcal{R}^*, z^*) on a message m^* . We may assume that \mathcal{R}^* was randomly

generated by the a group oracle query made by the adversary. Suppose that initially

$$\begin{aligned} \pi^{-1}(\mathcal{R}^*) &= a + b\mathsf{D} + \\ &\quad \sum_{\ell,j} (c_{\ell,j}\mathsf{R}_\ell^{(j)} + d_{\ell,j}\mathsf{S}_\ell^{(j)}) + \\ &\quad \sum_{k,i} (\bar{c}_{k,i}\bar{\mathsf{R}}_k^{(i)} + \bar{d}_{k,i}\bar{\mathsf{S}}_k^{(i)}). \end{aligned}$$

Here,

- the sum on ℓ, j corresponds to presignatures from batches whose transcript-pair-building functionalities have not yet terminated (and whose corresponding Random Beacons ρ_ℓ have not yet been generated), while
- the sum on k, i corresponds to presignatures from batches whose transcript-pair-building functionalities have terminated (and whose corresponding Random Beacons ρ_k have not yet been generated).

Note that all of the constants $a, b, c_{\ell,j}, d_{\ell,j}, \bar{c}_{k,i}, \bar{d}_{k,i}$ are fixed before \mathcal{R}^* is generated at random. In order for the forgery to be valid, all of the variables except D must be eliminated by substitution to end up with

$$\pi^{-1}(\mathcal{R}^*) = z^* - h^*\mathsf{D}.$$

We argue below that the sum on ℓ, j must be empty, as otherwise the forgery will be valid with only negligible probability. Assuming this for now, we focus on the sum on k, i . After substitution, we have

$$\pi^{-1}(\mathcal{R}^*) = \underbrace{\{a + \sum_{k,i} \bar{c}_{k,i} z_{k,i}\}}_{=z^*} + \underbrace{\{b - \sum_{k,i} \bar{c}_{k,i} h_{k,i}\}}_{=-h^*} \mathsf{D} + \sum_{k,i} \underbrace{\{\bar{d}_{k,i} - \bar{c}_{k,i} \delta_{k,i}\}}_{=0} \bar{\mathsf{S}}_k^{(i)}. \quad (18)$$

For the forgery to be valid, we must have $\bar{d}_{k,i} - \bar{c}_{k,i} \delta_{k,i} = 0$ for each k, i .

The rest of the argument is analogous to what we did in the Generic Group Model analysis in Section 3.3. Namely, if the forgery is to be valid, then with overwhelming probability, the adversary must have already made queries to Δ that produce the outputs $\bar{d}_{k,i}/\bar{c}_{k,i}$. Thus, at the time we generate \mathcal{R}^* , the inputs to H that determine the $h_{k,i}$'s have already been determined. Therefore, in order for the adversary to find m^* , he must essentially win Preimage Attack II on H , as in Section 3.1, but where now H is modeled as a random oracle.

We now return to the claim that the sum on ℓ, j must be empty. Suppose it is not. Then for some ℓ the corresponding transcript-pair-building functionality has not terminated at the time \mathcal{R}^* is generated, which means that the corresponding random value ρ_ℓ has not yet been generated either. For the forgery to be valid, the corresponding transcript-pair-building functionality must terminate, which only then determines values $\bar{c}_{\ell,i}$ and $\bar{d}_{\ell,i}$ before ρ_ℓ is generated, and then we must also make the coefficient $\bar{d}_{\ell,i} - \bar{c}_{\ell,i} \delta_{\ell,i}$ on $\bar{\mathsf{S}}_{\ell,i}$ vanish for each i via substitution. Since ρ_ℓ is input to the hash Δ to determine each $\delta_{\ell,i}$, the probability that the adversary can find other inputs to Δ so that $\bar{d}_{\ell,i} - \bar{c}_{\ell,i} \delta_{\ell,i} = 0$ for each i will be negligible.

Concrete security bounds. One can show that if H is modeled as a random oracle with an output space of size M , the adversary’s forging advantage is $O(N^2/q + N/M)$.

References

- [BHK⁺23] F. Benhamouda, S. Halevi, H. Krawczyk, Y. Ma, and T. Rabin. Sprint: High-throughput robust distributed schnorr signatures. Cryptology ePrint Archive, Paper 2023/427, 2023. URL <https://eprint.iacr.org/2023/427>. <https://eprint.iacr.org/2023/427>.
- [BLS01] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the weil pairing. In C. Boyd, editor, *Advances in Cryptology - ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, December 9-13, 2001, Proceedings*, volume 2248 of *Lecture Notes in Computer Science*, pages 514–532. Springer, 2001.
- [Bol03] A. Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-diffie-hellman-group signature scheme. In Y. Desmedt, editor, *Public Key Cryptography - PKC 2003, 6th International Workshop on Theory and Practice in Public Key Cryptography, Miami, FL, USA, January 6-8, 2003, Proceedings*, volume 2567 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2003.
- [Cer10] Certicom Research. Sec 2: Recommended elliptic curve domain parameters, 2010. Version 2.0, <http://www.secg.org/sec2-v2.pdf>.
- [CKM21] E. Crites, C. Komlo, and M. Maller. How to prove schnorr assuming schnorr: Security of multi- and threshold signatures. Cryptology ePrint Archive, Paper 2021/1375, 2021. <https://eprint.iacr.org/2021/1375>.
- [DEF⁺18] M. Drijvers, K. Edalatnejad, B. Ford, E. Kiltz, J. Loss, G. Neven, and I. Stepanovs. On the security of two-round multi-signatures. Cryptology ePrint Archive, Paper 2018/417, 2018. <https://eprint.iacr.org/2018/417>.
- [DFI22] The DFINITY Team. The internet computer for geeks. Cryptology ePrint Archive, Report 2022/087, 2022. <https://ia.cr/2022/087>.
- [DN07] I. Damgård and J. B. Nielsen. Scalable and unconditionally secure multiparty computation. In A. Menezes, editor, *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, volume 4622 of *Lecture Notes in Computer Science*, pages 572–590. Springer, 2007.
- [FY92] M. K. Franklin and M. Yung. Communication complexity of secure computation (extended abstract). In S. R. Kosaraju, M. Fellows, A. Wigderson, and J. A. Ellis, editors, *Proceedings of the 24th Annual ACM Symposium on Theory of Computing, May 4-6, 1992, Victoria, British Columbia, Canada*, pages 699–710. ACM, 1992.

- [GG20] R. Gennaro and S. Goldfeder. One round threshold ECDSA with identifiable abort. Cryptology ePrint Archive, Report 2020/540, 2020. <https://ia.cr/2020/540>.
- [GJKR07] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Secure distributed key generation for discrete-log based cryptosystems. *J. Cryptol.*, 20(1):51–83, 2007.
- [GS21] J. Groth and V. Shoup. On the security of ECDSA with additive key derivation and presignatures. Cryptology ePrint Archive, Report 2021/1330, 2021. <https://ia.cr/2021/1330>.
- [GS22] J. Groth and V. Shoup. Design and analysis of a distributed ECDSA signing service. Cryptology ePrint Archive, Report 2022/506, 2022. <https://ia.cr/2022/506>.
- [HN06] M. Hirt and J. B. Nielsen. Robust multiparty computation with linear communication complexity. In C. Dwork, editor, *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings*, volume 4117 of *Lecture Notes in Computer Science*, pages 463–482. Springer, 2006.
- [KG20] C. Komlo and I. Goldberg. Frost: Flexible round-optimized schnorr threshold signatures. Cryptology ePrint Archive, Paper 2020/852, 2020. <https://eprint.iacr.org/2020/852>.
- [NSW09] G. Neven, N. P. Smart, and B. Warinschi. Hash function requirements for schnorr signatures. *J. Math. Cryptol.*, 3(1):69–87, 2009.
- [PS96] D. Pointcheval and J. Stern. Provably secure blind signature schemes. In K. Kim and T. Matsumoto, editors, *Advances in Cryptology - ASIACRYPT '96, International Conference on the Theory and Applications of Cryptology and Information Security, Kyongju, Korea, November 3-7, 1996, Proceedings*, volume 1163 of *Lecture Notes in Computer Science*, pages 252–265. Springer, 1996.
- [RRJ⁺22] T. Ruffing, V. Ronge, E. Jin, J. Schneider-Bensch, and D. Schröder. ROAST: Robust asynchronous schnorr threshold signatures. Cryptology ePrint Archive, Paper 2022/550, 2022. URL <https://eprint.iacr.org/2022/550>. <https://eprint.iacr.org/2022/550>.
- [Wag02] D. A. Wagner. A generalized birthday problem. In M. Yung, editor, *Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings*, volume 2442 of *Lecture Notes in Computer Science*, pages 288–303. Springer, 2002.
- [Wui20] P. Wuille. Bip32: Hierarchical deterministic wallets, 2020. <https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki>.
- [Zha22] M. Zhandry. To label, or not to label (in generic groups). Cryptology ePrint Archive, Paper 2022/226, 2022. <https://eprint.iacr.org/2022/226>.