

Fully Adaptive Schnorr Threshold Signatures

Elizabeth Crites¹, Chelsea Komlo², and Mary Maller³

¹ University of Edinburgh

² University of Waterloo, Zcash Foundation

³ Ethereum Foundation, PQShield

Abstract. We prove adaptive security of a simple three-round threshold Schnorr signature scheme, which we call **Sparkle**. The standard notion of security for threshold signatures considers a *static* adversary – one who must declare which parties are corrupt at the beginning of the protocol. The stronger *adaptive* adversary can at any time corrupt parties and learn their state. This notion is natural and practical, yet not proven to be met by most schemes in the literature.

In this paper, we demonstrate that **Sparkle** achieves several levels of security based on different corruption models and assumptions. To begin with, **Sparkle** is statically secure under minimal assumptions: the discrete logarithm assumption (DL) and the random oracle model (ROM). If an adaptive adversary corrupts fewer than $t/2$ out of a threshold of $t + 1$ signers, then **Sparkle** is adaptively secure under a weaker variant of the one-more discrete logarithm assumption (AOMDL) in the ROM. Finally, we prove that **Sparkle** achieves *full* adaptive security, with a corruption threshold of t , under AOMDL in the algebraic group model (AGM) with random oracles. Importantly, we show adaptive security without requiring secure erasures. Ours is the first proof achieving full adaptive security without exponential tightness loss for *any* threshold Schnorr signature scheme; moreover, the reduction is tight.

Table of Contents

1	Introduction	3
2	Related Work	5
3	Preliminaries	7
	3.1 General Notation	7
	3.2 Definitions and Assumptions	8
4	Threshold Signature Schemes	11
	4.1 Static Security	12
	4.2 Adaptive Security	14
5	Schnorr Threshold Signature Scheme Sparkle	14
6	Static Security Under Standard Assumptions	16
7	Adaptive Security Against up to $t/2$ Corruptions	17
8	Adaptive Security Against t Corruptions	18
9	Static and Adaptive Security Proofs	19
	9.1 Proof of Static Security	19
	9.2 Proof of Adaptive Security for up to $t/2$ Corruptions	22
	9.3 Proof of Adaptive Security for up to t Corruptions	26
A	Proof of Static Security in the AGM	34

1 Introduction

A threshold signature scheme allows a set of n possible signers to jointly produce a signature over a message and with respect to a single public key, so long as at least a threshold $t + 1$ of signers participate. Importantly, threshold signature schemes are secure even if t signers are under adversarial control. A recent line of work has explored multi-party signature schemes whose output is a standard (single-party) Schnorr signature [32, 34, 38]. Schnorr signatures admit an efficient and compact representation even in the multi-party setting, which makes them of particular interest for practical use [17].

Static vs. Adaptive Security. Most threshold signature schemes in the literature are proven *statically* secure. In the static setting, an adversary must declare which parties it wishes to corrupt in advance of any messages being sent. This model places an artificial restriction on the adversary’s capabilities: in reality, malicious actors may observe a system before targeting specific parties. Thus, adaptive security is a strictly stronger notion, and indeed there are schemes that are statically but not adaptively secure [14]. While there are generic methods for transforming a statically secure scheme into an adaptively secure one [16], such as guessing the corrupted parties and aborting if incorrect, these methods incur undesirable performance overhead and a tightness loss of $\binom{n}{t}$. This grows exponentially in the number of parties, and no adaptive guarantees can be made for larger n . Adaptive security without exponential loss is challenging to achieve. A number of other techniques for proving adaptive security have been proposed, but similarly require undesirable tradeoffs. Prior methods include erasure of secret state [16], which is not easily enforced in practice, or heavyweight tools, such as non-committing encryption [30].

In this work, we investigate the adaptive security of a simple three-round threshold Schnorr signature scheme, which we call **Sparkle**, under different corruption models and security assumptions. Achieving adaptive security is not just of theoretical interest: NIST recently published a call for threshold EdDSA/Schnorr schemes and included adaptive security as a main goal [11, 12], ideally supporting up to $n = 1024$ or more parties. Our techniques are likely of independent interest, as this paper introduces the first proof achieving full adaptive security without exponential tightness loss for *any* threshold Schnorr signature scheme; moreover, the reduction is tight.

Concurrent Security. A *concurrent* adversary may open an arbitrary number of signing sessions simultaneously and unforgeability should still hold. Concurrent security is also a difficult property to achieve, and indeed a host of multi-, threshold, and blind signature schemes were demonstrated to be broken by concurrent (ROS) attacks first observed by DEFKLNS19 [18] and exhibited in polynomial time in by BLLOR21 [8]. Our security reductions for **Sparkle** hold against a concurrent and adaptive adversary.

Schnorr Threshold Signature Sparkle. In this paper, we prove the adaptive security of a simple three-round threshold Schnorr signature scheme **Sparkle** that

Scheme	Static Assumptions	Sign					Combine	
		Performance			Bandwidth		Performance	
		rounds	exp	H	\mathbb{G}	\mathbb{F}	exp	H
FROST [32]	OMDL+ROM	2	$t + 2$	$t + 1$	2	1	t	t
FROST2 [5]	OMDL+ROM	2	3	2	2	1	1	1
Lindell22 [34]	DL+ROM	3	$11t + 1$	$18t + 10^7$	11	25	0	0
Sparkle	DL+ROM	3	1	$t + 2$	1	2	0	0

Fig. 1. Efficiency of Two- and Three-Round Schnorr Threshold Signature Schemes. All output a standard Schnorr signature. We only compare schemes that are secure against ROS attacks [8]. The number of network rounds between participants is given in the rounds column. exp stands for the number of group exponentiations. The total number of group and field elements sent by each signer is denoted by \mathbb{G} and \mathbb{F} , respectively. H denotes the total number of hashing operations performed. The cost of signature verification is identical for each scheme, and is simply the cost of verifying a single Schnorr signature. Estimates for Lindell22 are made with respect to a 128-bit security level for Fischlin [21], where $r = 8$ is the number of commitments for a Fischlin proof and the length of the zero vector is $b = 16$, such that $b \cdot r = 128$.

follows a commit-reveal paradigm. To begin with, we prove the static security of Sparkle from minimal assumptions: that the discrete logarithm assumption (DL) holds in the random oracle model (ROM). This is the same assumption and model for which Schnorr signatures themselves are proven secure. Our security reduction incurs no additional tightness loss compared to the security reduction for plain Schnorr signatures [40]. We compare the efficiency of Sparkle with existing two- and three-round threshold Schnorr signatures in Table 1.

Adaptive Security of Sparkle. We next consider an adaptive adversary who may corrupt up to $t/2$ out of a threshold of $t + 1$ parties over the course of the protocol. We prove that Sparkle is adaptively secure in the random oracle model under AOMDL, a weaker variant of the one-more discrete logarithm assumption formalized in [38]. Importantly, we do so without requiring the assumption of secure erasures. In the $t + 1$ -aomdl game, the adversary is given as input an AOMDL challenge that is a vector of group elements of length $t + 1$. The adversary is given access to a discrete logarithm oracle, which returns the discrete logarithm of a group element chosen by the adversary. To win the $t + 1$ -aomdl game, the adversary must output all $t + 1$ discrete logarithms of its challenge, having queried its DL oracle a maximum of t times. The AOMDL assumption is stronger than the discrete logarithm assumption because of the adversary’s ability to request for up to t discrete logarithm solutions before returning the $t + 1$ discrete logarithms of its challenge. On the other hand, the AOMDL assumption is strictly weaker than standard OMDL [7] because the adversary only receives the discrete logarithm of linear combinations of its challenge elements, which allows the oracle to run in polynomial time and makes it a falsifiable assumption [4, 28, 31].

In the case where the adversary can corrupt up to a full t parties, it is not clear how to prove the adaptive security of Sparkle under AOMDL+ROM alone. The reason is that in order to extract an AOMDL solution, the adversary is rewound




Fig. 2. Comparison of security models and assumptions for our Schnorr threshold signature scheme **Sparkle**. The signing threshold is $t + 1$. DL is the Discrete Logarithm Assumption, and AOMDL is the Algebraic One-More Discrete Logarithm Assumption. ROM is the Random Oracle Model, and AGM is the Algebraic Group Model.

once, and there is no guarantee that the adversary will corrupt the same set of parties after the fork as it did during the first iteration of the protocol. When the adversary can corrupt only $t/2$ parties, this causes no issues, as the total number of corruptions over both iterations does not exceed t . If the adversary could corrupt more parties, the reduction would query its DL oracle more than t times and would lose the $t + 1$ -aomdl game.

We thus look towards the algebraic group model (AGM) [22] for proving our strongest adaptivity result. The AGM assumes that whenever an adversary outputs a group element, it also outputs an algebraic representation specifying how the group element depends on previously seen values. In the AGM, we are able to prove full adaptive security of **Sparkle**, with corruption threshold t , under the AOMDL assumption and random oracles. Our security reduction is *straight-line*, i.e., does not rewind the adversary, and so avoids counting the number of corruptions over different forks of the adversary’s execution.

2 Related Work

Threshold Schnorr Signatures. Closest to the design of **Sparkle** is the MSDL scheme presented by Boneh, Drijvers, and Neven [10], the three-round MuSig scheme by Maxwell, Poelstra, Seurin, and Wuille [37], and the 2Schnorr scheme by Nicolosi, Krohn, Dodis, and Mazières [39]. However, MSDL and MuSig consider only the multisignature setting (n -out-of- n), and 2Schnorr considers only the 2-out-of-2 setting. Concurrent to this work, Makriyannis [36] defines a commit-reveal threshold Schnorr signature scheme similar to **Sparkle**, and proves security with respect to an idealized notion of threshold signatures by Canetti et al. [15]. However, the proof of adaptive security incurs exponential tightness loss relative to the number of parties.

Stinson and Strobl [43] present a threshold Schnorr scheme secure in the random oracle model (ROM) under the discrete logarithm assumption (DL).

However, their scheme requires performing a three-round distributed key generation protocol (DKG) [26] to generate the nonce for each signature, and hence has undesirable network overhead: at a minimum, it requires participants to perform four rounds in total. Further, the proof of security assumes only a static adversary.

Komlo and Goldberg [32] present a two-round Schnorr threshold signature FROST. Unlike prior threshold Schnorr schemes in the literature [25], FROST is secure against a concurrent adversary and is not susceptible to ROS attacks [8]. FROST2, introduced in [5], is an optimized version of FROST that reduces the number of exponentiations required for signing from $t + 1$ to one. (See Table 1 for a comparison of efficiency.) However, both FROST and FROST2 are proven secure assuming a static adversary and the stronger one-more discrete logarithm assumption (OMDL) [5]. While Sparkle adds an additional round of communication, it only requires a single exponentiation per signer and can be proven secure under standard assumptions, which are also criteria of interest in [11, 12].

Lindell presents a three-round Schnorr threshold signature [34] with the goal of defining a threshold scheme that is secure against ROS attacks and proven secure in the random oracle model (ROM) under the discrete logarithm assumption only. Security is modeled in the the universally composable framework (UC) [13] and therefore captures a concurrent adversary. However, the security proof assumes the adversary is static, and no claims are made regarding adaptive security. Sparkle similarly relies on only ROM and DL assumptions, but does not require the use of online-extractable zero-knowledge proofs [21], and hence is both significantly more efficient and a simpler design.

Adaptive Security of Threshold Signatures. While adaptive security for threshold schemes is a well-known topic in the literature, to the best of our knowledge, no proof of adaptive security without exponential tightness loss exists for a threshold Schnorr signature scheme that is secure against ROS attacks.

Generalized techniques for transforming statically secure threshold schemes into adaptively secure schemes have been defined in the literature [16, 30, 35]. However, these techniques introduce prohibitive performance overhead, such as requiring a robust distributed key generation mechanism (DKG) for nonce generation.

Almansa, Damgård, and Nielsen [2] present a threshold RSA scheme with proactive and adaptive security, but these results do not translate to the discrete logarithm setting. Libert, Joye, and Yung [33] present a variant of the threshold BLS [9] scheme that is adaptively secure. However, their variant is incompatible with single-party BLS verification, an often-critical goal for threshold schemes in practice. Bacho and Loss [3] demonstrate the adaptive security of threshold BLS directly in the algebraic group model from the $t + 1$ -omdl assumption. Their reduction is tight. Interestingly, they demonstrate that the $t + 1$ -omdl assumption is the minimum assumption under which threshold BLS can be proven adaptively secure.

Definitions. In this work, we employ a game-based approach to defining the static and adaptive security of a threshold signature, formalizing prior notions presented in the literature [33]. Alternative definitions of adaptive security in the UC setting have been proposed by Canetti et al. [15]. They prove their threshold ECDSA scheme adaptively secure assuming that ECDSA is secure, in addition to other non-interactive and falsifiable assumptions. However, their construction focuses on n -out-of- n multi-party signing, and their techniques critically do not translate to the t -out-of- n setting unless $\binom{n}{t}$ is small. Our fully adaptive t -out-of- n construction requires the AGM, and incorporating algebraic adversaries into the UC setting is known to be a hard problem [1].

On the other hand, Lindell [34] shows that their protocol UC-realizes the Schnorr functionality with aborts, in the presence of an adversary that non-adaptively corrupts t parties. Secure evaluation of the Schnorr functionality is stronger than unforgeability. As noted in [15], it is arguably overly strong in the sense that it necessitates certain design decisions, such as incorporating online-extractable zero-knowledge proofs. Indeed, [34] elected for Fischlin proofs (see Table 1). The only method to bias the nonces in both Sparkle and [34] is to abort; this is not the case in FROST [32]. However, all three works allow aborts, and therefore the distribution of the secret randomness cannot be considered uniform [19].

3 Preliminaries

3.1 General Notation

Let $\lambda \in \mathbb{N}$ denote the security parameter and 1^λ its unary representation. A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is called *negligible* if for all $c \in \mathbb{R}, c > 0$, there exists $k_0 \in \mathbb{N}$ such that $|f(k)| < \frac{1}{k^c}$ for all $k \in \mathbb{N}, k \geq k_0$. For a non-empty set S , let $x \leftarrow_s S$ denote sampling an element of S uniformly at random and assigning it to x . We use $[n]$ to represent the set $\{1, \dots, n\}$ and $[0..n]$ to represent the set $\{0, \dots, n\}$. We represent vectors as $\mathbf{a} = (a_1, \dots, a_n)$.

Let PPT denote probabilistic polynomial time. Algorithms are randomized unless explicitly noted otherwise. Let $y \leftarrow A(x; \omega)$ denote running algorithm A on input x and randomness ω and assigning its output to y . Let $y \leftarrow_s A(x)$ denote $y \leftarrow A(x; \omega)$ for a uniformly random ω . The set of values that have non-zero probability of being output by A on input x is denoted by $[A(x)]$.

Group Generation. Let GrGen be a deterministic polynomial-time algorithm that takes as input a security parameter 1^λ and outputs a group description (\mathbb{G}, p, g) consisting of a group \mathbb{G} of order p , where p is a λ -bit prime, and a generator g of \mathbb{G} .

Polynomial Interpolation. A polynomial $f(x) = a_0 + a_1x + a_2x^2 + \dots + a_t x^t$ of degree t over a field \mathbb{F} can be interpolated by $t + 1$ points. Let $\eta \subseteq [n]$ be the list of $t + 1$ distinct indices corresponding to the x -coordinates $x_i \in \mathbb{F}, i \in \eta$ of these points. Then the Lagrange polynomial $L_i(x)$ has the form:

$$L_i(x) = \prod_{j \in \eta; j \neq i} \frac{x - x_j}{x_i - x_j} \quad (1)$$

Given a set of $t + 1$ points $(x_i, f(x_i))_{i \in [t+1]}$, any point $f(x_\ell)$ on the polynomial f can be determined by Lagrange interpolation as follows:

$$f(x_\ell) = \sum_{k \in \eta} f(x_k) \cdot L_k(x_\ell)$$

3.2 Definitions and Assumptions

Assumption 1 (Discrete Logarithm Assumption (DL)) *Let the advantage of an adversary \mathcal{A} playing the discrete logarithm game Game^{dl} , as defined in Figure 3, be as follows:*

$$\text{Adv}_{\mathcal{A}}^{\text{dl}}(\lambda) = |\Pr[\text{Game}_{\mathcal{A}}^{\text{dl}}(\lambda) = 1]|$$

The discrete logarithm assumption holds if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{\text{dl}}(\lambda)$ is negligible.

Assumption 2 (One-More Discrete Logarithm Assumption (OMDL)) [7] *Let the advantage of an adversary \mathcal{A} playing the $t + 1$ -one-more discrete logarithm game $\text{Game}^{t+1\text{-omdl}}$, as defined in Figure 3, be as follows:*

$$\text{Adv}_{\mathcal{A}}^{t+1\text{-omdl}}(\lambda) = |\Pr[\text{Game}_{\mathcal{A}}^{t+1\text{-omdl}}(\lambda) = 1]|$$

The one-more discrete logarithm assumption holds if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{t+1\text{-omdl}}(\lambda)$ is negligible.

Algebraic One-More Discrete Logarithm Assumption (AOMDL). The algebraic one-more discrete logarithm assumption (AOMDL) was formalized by Jonas, Ruffing, and Seurin [38]. We highlight differences with the OMDL assumption in Figure 3. Note that while the OMDL assumption is not falsifiable (because the adversary is allowed to query for the discrete logarithm of any arbitrary group element), the AOMDL assumption is falsifiable. This is because the input to the AOMDL game is a linear combination of all challenges issued to the adversary, so it is guaranteed that the environment runs in PPT when its respective adversary runs in PPT. Hence, the AOMDL assumption is a strictly weaker assumption than standard OMDL.

An adversary in the AOMDL game is initialized in exactly the same manner as an OMDL adversary, with the same winning condition. The only distinction between the games is the inputs and outputs from the discrete logarithm oracle. In the AOMDL setting, the adversary provides a vector $\mathbf{a} = (a_0, \dots, a_t)$ of coefficients to the discrete logarithm oracle. The oracle then responds with the integer linear combination of these coefficients with respect to the discrete logarithm of the set of challenges (X_0, \dots, X_t) . It is easy to see that the AOMDL

MAIN Game $_{\mathcal{A}}^{\text{dl}}(\lambda)$	MAIN Game $_{\mathcal{A}}^{t+1-\mathbf{a} \text{ omdl}}(\lambda)$	$\mathcal{O}^{\text{dl}}(X, \alpha, \{\beta_i\}_{i=1}^t)$
$(\mathbb{G}, p, g) \leftarrow \text{GrGen}(1^\lambda)$ $x \leftarrow_{\mathcal{S}} \mathbb{Z}_p; X \leftarrow g^x$ $x' \leftarrow_{\mathcal{S}} \mathcal{A}((\mathbb{G}, p, g), X)$ if $x' = x$ return 1 return 0	$(\mathbb{G}, p, g) \leftarrow \text{GrGen}(1^\lambda)$ $Q \leftarrow \emptyset$ $q \leftarrow 0$ for $i \in [0..t]$ do $x_i \leftarrow_{\mathcal{S}} \mathbb{Z}_p; X_i \leftarrow g^{x_i}$ $Q[X_i] = x_i$ $\mathbf{x} \leftarrow (x_0, \dots, x_t)$ $\mathbf{X} \leftarrow (X_0, X_1, \dots, X_t)$ $\mathbf{x}' \leftarrow \mathcal{A}^{\text{O}^{\text{dl}}}((\mathbb{G}, p, g), \mathbf{X})$ if $\mathbf{x}' = \mathbf{x} \wedge q < t + 1$ return 1 return 0	$// X = g^\alpha \prod_{i=0}^t X_i^{\beta_i}$ $q \leftarrow q + 1$ $x \leftarrow \alpha + \sum_{i=0}^t x_i \beta_i$ return x $x \leftarrow \text{dlog}(X)$ return x

Fig. 3. The Discrete Logarithm (DL), One-More Discrete Logarithm (OMDL), and Algebraic One-More Discrete Logarithm (AOMDL) games. \mathbb{G} is a cyclic group with prime order p and generator g . \mathbf{a} is a vector whose elements are in \mathbb{Z}_p . dlog is an algorithm that finds the discrete logarithm relation of the input X and g . The differences between the OMDL and the AOMDL games are highlighted in gray; the key distinction is that in the AOMDL game, the adversary can query for linear combinations of elements which the environment has generated directly; in the OMDL game, the environment must return the discrete logarithm of an arbitrary group element.

game can be used to win the OMDL game, by the adversary simply querying \mathcal{O}^{dl} with a bit vector, to return the discrete logarithm of the term whose coefficient is set to one.

Assumption 3 (Algebraic One-More Discrete Logarithm Assumption)

[38] *Let the advantage of an adversary \mathcal{A} playing the $t + 1$ -algebraic one-more discrete logarithm game $\text{Game}_{\mathcal{A}}^{t+1-\mathbf{a} \text{ omdl}}$, as defined in Figure 3, be as follows:*

$$\text{Adv}_{\mathcal{A}}^{t+1-\mathbf{a} \text{ omdl}}(\lambda) = |\Pr[\text{Game}_{\mathcal{A}}^{t+1-\mathbf{a} \text{ omdl}}(\lambda) = 1]|$$

The algebraic one-more discrete logarithm assumption holds if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{t+1-\mathbf{a} \text{ omdl}}(\lambda)$ is negligible.

Assumption 4 (Algebraic Group Model (AGM) [22])

An adversary is algebraic if for every group element $Z \in \mathbb{G} = \langle g \rangle$ that it outputs, it is required to output a representation $\mathbf{a} = (a_0, a_1, a_2, \dots)$ such that $Z = g^{a_0} \prod Y_i^{a_i}$, where $Y_1, Y_2, \dots \in \mathbb{G}$ are group elements that the adversary has seen thus far.

Definition 1 (Schnorr Signatures [41]). *The Schnorr signature scheme consists of polynomial-time algorithms (Setup, KeyGen, Sign, Verify), defined as follows:*

- **Setup**(1^λ) \rightarrow **par**: *On input the security parameter, run $(\mathbb{G}, p, g) \leftarrow \text{GrGen}(1^\lambda)$ and select a hash function $\text{H} : \{0, 1\}^* \rightarrow \mathbb{Z}_p$. Output public parameters **par** = $((\mathbb{G}, p, g), \text{H})$ (which are given implicitly as input to all other algorithms).*
- **KeyGen**(1^λ) \rightarrow (X, x) : *On input the security parameter, sample a secret key $x \leftarrow_{\$} \mathbb{Z}_p$ and compute a public key $X \leftarrow g^x$. Output (X, x) .*
- **Sign**(x, m) \rightarrow σ : *On input message m , the signer samples a nonce $r \leftarrow_{\$} \mathbb{Z}_p$ and computes a nonce commitment $R \leftarrow g^r$. The signer then computes the challenge $c \leftarrow \text{H}(X, m, R)$ and the response $z \leftarrow r + cx$. Output the signature $\sigma = (R, z)$.*
- **Verify**(X, m, σ) \rightarrow $0/1$: *On input the public key X , a message m , and a signature $\sigma = (R, z)$, the verifier computes $c \leftarrow \text{H}(X, m, R)$ and accepts if $RX^c = g^z$.*

A Schnorr signature is a Sigma protocol zero-knowledge proof of knowledge of the discrete logarithm of the public key, made non-interactive and bound to the message m by the Fiat-Shamir transform [20]. Schnorr signatures are secure under the discrete logarithm assumption in the random oracle model [40].

Definition 2 (Shamir secret sharing [42]). *Shamir secret sharing is an $(n, t+1)$ -threshold secret sharing scheme $\mathcal{SS} = (\text{IssueShares}, \text{Recover})$ that consists of the following algorithms:*

- **IssueShares**($x, n, t + 1$) \rightarrow $\{(1, x_1), \dots, (n, x_n)\}$: *On input a secret x , the number of participants n , and a threshold $t + 1$, perform the following. First, define a polynomial $f(Z) = x + a_1Z + a_2Z^2 + \dots + a_tZ^t$ by sampling t coefficients at random $(a_1, \dots, a_t) \leftarrow_{\$} \mathbb{Z}_p$. Then, set each participant's share $x_i, i \in [n]$, to be the evaluation of $f(i)$:*

$$x_i \leftarrow x + \sum_{j \in [t]} a_j i^j$$

Output $\{(i, x_i)\}_{i \in [n]}$.

- **Recover**($t + 1, \{(i, x_i)\}_{i \in \mathcal{S}}$) \rightarrow \perp/x : *On input a threshold $t + 1$ and a set of shares $\{(i, x_i)\}_{i \in \mathcal{S}}$, output \perp if $\mathcal{S} \not\subseteq [n]$ or if $|\mathcal{S}| < t + 1$. Otherwise, recover x as follows:*

$$x \leftarrow \sum_{i \in \mathcal{S}} \lambda_i x_i$$

where the Lagrange coefficient for the set \mathcal{S} is defined by

$$\lambda_i = \prod_{j \in \mathcal{S}, j \neq i} \frac{j}{i - j}$$

4 Threshold Signature Schemes

We begin with the definition of a threshold signature scheme. We build upon prior game-based definitions in the literature [24, 27, 23], but define an additional algorithm for combining signatures that is separate from the signing rounds. It may be performed by one of the signers or some external party. Note that our definition assumes a three-round signing protocol, but can easily be extended to an arbitrary number of rounds.

We then provide formal definitions of static and adaptive security for threshold signatures. Our security definitions are game based, and can be seen as the threshold analogue to the notion of existential unforgeability under chosen message attack (EUF-CMA) for single-party signatures.

A threshold signature scheme TS whose signing protocol consists of three rounds is a tuple of the following algorithms $\text{TS} = (\text{Setup}, \text{KeyGen}, (\text{Sign}, \text{Sign}', \text{Sign}''), \text{Combine}, \text{Verify})$, defined as follows:

- $\text{Setup}(1^\lambda) \rightarrow \text{par}$: This algorithm generates the public parameters par that are implicitly given as input to all other algorithms.
- $\text{KeyGen}(n, t + 1) \rightarrow (\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_i\}_{i \in [n]})$: A randomized protocol that takes as input the total number of signing parties n and the threshold $t + 1$ representing the minimum number of signing parties. The output is the joint public key \tilde{X} representing the set of n parties, n public keys $\{\tilde{X}_i\}_{i \in [n]}$ representing each party, and the set of secret shares $\{x_i\}_{i \in [n]}$ with respect to a randomly generated secret. It is assumed that participant i is sent its respective share x_i privately.
- $(\text{Sign}, \text{Sign}', \text{Sign}'') \rightarrow \rho_i$: A set of randomized algorithms where each subsequent Sign algorithm represents a single stage in an interactive signing protocol, performed by each signing party in a signing set $\mathcal{S} \subseteq [n]$, $|\mathcal{S}| \geq t + 1$ with respect to a message m . The output from each signing algorithm is a protocol message ρ_i .
- $\text{Combine}(\{(\rho'_i, \rho''_i)\}_{i \in \mathcal{S}}) \rightarrow (m, \sigma)$: A deterministic algorithm that takes as input the set of protocol messages (ρ'_i, ρ''_i) representing the messages sent by each party in Sign and Sign'' for each party in the signing set \mathcal{S} . It outputs σ as the joint signature representing the signing set.
- $\text{Verify}(\tilde{X}, m, \sigma) \rightarrow 0/1$: A deterministic algorithm that takes as input the public key \tilde{X} , a message m , and signature σ . It outputs 0 or 1 indicating whether σ is valid or not.

Remark 1 (Distributed key generation). Our definition assumes a centralized key generation algorithm KeyGen to generate the public key \tilde{X} and set of shares $\{\tilde{X}_i, x_i\}_{i \in [n]}$. However, our scheme and proofs can be adapted to use a fully decentralized distributed key generation protocol (DKG).

Remark 2 (Number of signing rounds). Our definition models a three-round signing scheme by the algorithms $(\text{Sign}, \text{Sign}', \text{Sign}'')$. However, this definition can be adapted to schemes with fewer or more rounds.

Correctness. Correctness requires that for all security parameters λ , for all $t + 1 \leq \ell \leq n$, and for all messages m , if $(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_i\}_{i \in [n]}) \leftarrow^* \text{KeyGen}(n, t + 1)$ and ℓ signers input (x_i, m) to the signing protocol $(\text{Sign}, \text{Sign}', \text{Sign}'')$, then every signer will output a signature share that, when combined with all other shares, results in a signature σ satisfying $\text{Verify}(\tilde{X}, m, \sigma) = 1$.

4.1 Static Security

We present a game-based definition of static security analogous to EUF-CMA for standard signature schemes [29]. The static game is defined formally in Figure 4.1, where it contains all but the dashed boxes. We assume that the adversary can corrupt up to t signers, which we denote by the set cor . We assume the set of honest signers hon to be of size at least one. In addition to the security parameter λ , the game additionally accepts the parameter τ , which determines the ratio of the reconstruction threshold compared to corrupted players. In the static unforgeability game, the adversary can corrupt the maximum possible of t parties, corresponding to $\tau = 1$.

In the static unforgeability game, the challenger generates public parameters par and returns them to the adversary. The adversary must now choose the set of corrupted participants cor , which are fixed for the duration of the protocol. The challenger then runs KeyGen to derive the joint public key \tilde{X} , the individual public key shares $\{\tilde{X}_i\}_{i \in [n]}$, and the signing shares $\{x_i\}_{i \in [n]}$. It returns \tilde{X} , $\{\tilde{X}_i\}_{i \in [n]}$, and the set of corrupt signing shares $\{x_j\}_{j \in \text{cor}}$ to the adversary.

After key generation has concluded, the adversary is able to query honest signers at each step in the signing protocol by querying oracles $\mathcal{O}^{\text{Sign}}, \mathcal{O}^{\text{Sign}'}, \mathcal{O}^{\text{Sign}''}$, and has full power over choosing the set of signers and the message to be signed.

The adversary wins if it can produce a valid forgery $\sigma^* = (R^*, z^*)$ with respect to the joint public key \tilde{X} representing the set of n signers, on a message m^* that has not been previously queried.

The adversary can be *rushing*, meaning that it may wait to produce its outputs after having seen the honest outputs first. The adversary is allowed to be *concurrent*, meaning that it can open simultaneous signing sessions at once or choose not to complete a signing session (e.g., the adversary might query $\mathcal{O}^{\text{Sign}}, \mathcal{O}^{\text{Sign}'}$ but not $\mathcal{O}^{\text{Sign}''}$ before it starts a new session). By modeling a concurrent adversary, we ensure that our notion of security protects against practical attacks which occur in such a setting [8]. Session identifiers are represented as ssid .

Definition 3 (Static Security). Let the advantage of an adversary \mathcal{A} playing the static security game $\text{Game}_{\mathcal{A}}^{\text{UF}}(\lambda, \tau)$, as defined in Figure 4.1, be as follows:

$$\text{Adv}_{\mathcal{A}, \text{TS}}^{\text{st-sec}}(\lambda, \tau) = |\Pr[\text{Game}_{\mathcal{A}, \text{TS}}^{\text{UF}}(\lambda, \tau) = 1]|$$

A threshold signature scheme TS is statically secure if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{\text{st-sec}}(\lambda, \tau)$ is negligible.

<div style="border: 1px solid black; padding: 5px;"> <p>MAIN $\text{Game}_{\mathcal{A}, \text{TS}}^{\text{UF}}(\lambda, \tau)$ $\text{Game}_{\mathcal{A}, \text{TS}}^{\text{adp-UF}}(\lambda, \tau)$</p> <p>$\text{par} \leftarrow \text{Setup}(1^\lambda)$</p> <p>$W \leftarrow \emptyset$ // open signing sessions for $\mathcal{O}^{\text{Sign}}$</p> <p>$W' \leftarrow \emptyset$ // open signing sessions for $\mathcal{O}^{\text{Sign}'}$</p> <p>$W'' \leftarrow \emptyset$ // signing sessions for $\mathcal{O}^{\text{Sign}''}$</p> <p>$Q \leftarrow \emptyset$ // set of queried messages</p> <p>$(n, t + 1, \text{cor}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}(\text{par})$</p> <p>$\text{hon} \leftarrow [n] \setminus \text{cor}$</p> <p>$(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_i\}_{i \in [n]}) \leftarrow \mathcal{K}(\text{KeyGen}(n, t + 1))$</p> <p>$\text{input} \leftarrow (\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_j\}_{j \in \text{cor}}, \text{st}_{\mathcal{A}})$</p> <p>$(m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Sign}}, \mathcal{O}^{\text{Sign}'}, \mathcal{O}^{\text{Sign}''}, \text{Corrupt}}(\text{input})$</p> <p>if $\text{cor} \leq t/\tau$</p> <p style="padding-left: 20px;">// τ determines the ratio</p> <p style="padding-left: 20px;">// of honest and corrupt players</p> <p style="padding-left: 20px;">if $m^* \notin Q \wedge \text{Verify}(\tilde{X}, m^*, \sigma^*) = 1$</p> <p style="padding-left: 40px;">return 1</p> <p>return 0</p> </div> <div style="border: 1px dashed black; padding: 5px; margin-top: 5px;"> <p>$\mathcal{O}^{\text{Corrupt}}(k)$</p> <hr style="border: 0.5px solid black;"/> <p>if $k \notin \text{hon}$ return \perp</p> <p>$\text{cor} \leftarrow \text{cor} \cup \{k\}$</p> <p>$\text{hon} \leftarrow \text{cor} \setminus \{k\}$</p> <p>return $(x_k, \{\text{st}_{k, \ell}\}_{\ell \in [\text{ssid}]})$</p> </div>	<div style="border: 1px solid black; padding: 5px;"> <p>$\mathcal{O}^{\text{Sign}}(k, \text{ssid}, m, \mathcal{S})$</p> <hr style="border: 0.5px solid black;"/> <p>// k denotes the participant identifier</p> <p>$Q \leftarrow Q \cup \{m\}$</p> <p>$W \leftarrow W \cup \{(k, \text{ssid})\}$</p> <p>$(\rho_k, \text{st}_{k, \text{ssid}}) \leftarrow \text{Sign}(m, \mathcal{S})$</p> <p>return ρ_k</p> </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> <p>$\mathcal{O}^{\text{Sign}'}(k, \text{ssid}, m, \mathcal{S}, \{\rho_i\}_{i \in \mathcal{S}})$</p> <hr style="border: 0.5px solid black;"/> <p>if $(k, \text{ssid}) \notin W$ return \perp</p> <p>if $(k, \text{ssid}) \in W'$ return \perp</p> <p>$(\rho'_k, \text{st}_{k, \text{ssid}}) \leftarrow \text{Sign}'(\text{st}_{k, \text{ssid}}, x_k, \{\rho_i\}_{i \in \mathcal{S}})$</p> <p>$W' \leftarrow W' \cup \{(k, \text{ssid})\}$</p> <p>return ρ'_k</p> </div> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> <p>$\mathcal{O}^{\text{Sign}''}(k, \text{ssid}, \{\rho'_i\}_{i \in \mathcal{S}})$</p> <hr style="border: 0.5px solid black;"/> <p>if $(k, \text{ssid}) \notin W'$ return \perp</p> <p>if $(k, \text{ssid}) \in W''$ return \perp</p> <p>$W'' \leftarrow W'' \cup \{(k, \text{ssid})\}$</p> <p>$\rho''_k \leftarrow \text{Sign}''(\text{st}_{k, \text{ssid}}, x_k, \{\rho'_i\}_{i \in \mathcal{S}})$</p> <p>return ρ''_k</p> </div>
---	--

Fig. 4. Static and adaptive unforgeability games for a threshold signature scheme with three signing rounds. The public parameters par are implicitly given as input to all algorithms, and ρ represents protocol messages defined within the construction. The static game contains all but the dashed boxes, and the adaptive game adds the dashed boxes. τ determines the ratio of corrupted players to the reconstruction threshold; for example, $\tau = 1$ when assuming only one honest player at a minimum (honest minority), and $\tau = 2$ when signing requires at least $t/2 + 1$ honest players (honest majority).

4.2 Adaptive Security

We build upon the notion of adaptive security for threshold signature schemes by Libert, Joye, and Yung [33]. We provide a formal game-based definition, which is specified in Figure 4.1. The adaptive security game contains the same inputs and algorithms as the static game, but additionally includes a corruption oracle $\mathcal{O}^{\text{Corrupt}}$.

In the adaptive setting, the adversary is not restricted to choosing its set of corrupt participants cor only at the beginning of the game. Instead, the adversary can at any time choose to corrupt an honest party by querying $\mathcal{O}^{\text{Corrupt}}$, receiving in return the honest party’s secret key and state across all signing sessions.

In addition to producing a valid forgery, the adversary must meet the winning condition that the set of corrupted participants at the end of the experiment is within the expected bound, i.e., less than t/τ , with respect to the corruption ratio τ . If the adversary can corrupt up to $t/2$ parties, then $\tau = 2$; $\tau = 1$ allows corruption of up to t parties.

Definition 4 (Adaptive Security). *Let the advantage of an adversary \mathcal{A} playing the static security game $\text{Game}_{\mathcal{A}, \text{TS}}^{\text{adp-UF}}(\lambda, \tau)$, as defined in Figure 4.1, be as follows:*

$$\text{Adv}_{\mathcal{A}, \text{TS}}^{\text{adp-sec}}(\lambda, \tau) = |\Pr[\text{Game}_{\mathcal{A}, \text{TS}}^{\text{adp-UF}}(\lambda, \tau) = 1]|$$

A threshold signature scheme TS is adaptively secure if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{\text{adp-sec}}(\lambda, \tau)$ is negligible.

5 Schnorr Threshold Signature Scheme Sparkle

Sparkle is a simple three-round Schnorr threshold signature scheme (Fig. 5). We employ a centralized key generation mechanism, but alternatively a distributed key generation protocol (DKG) could be used. The public parameters par generated during setup are provided as input to all other algorithms and protocols. We assume some external mechanism to choose the set of signers $\mathcal{S} \subseteq \{1, \dots, n\}$, where $t + 1 \leq |\mathcal{S}| \leq n$ and \mathcal{S} is ordered to ensure consistency.

Intuitively, Sparkle follows a commit-reveal paradigm for its three-round signing protocol. In the first round, participants commit to their nonces $R_k = g^{r_k}$ by publishing $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$, where m is the message and \mathcal{S} is the signing set. In the second round, participants reveal R_k in the clear. In the third round, participants derive the aggregate nonce \tilde{R} , derive the Schnorr challenge using \tilde{R} , and then produce their signature share z_i . In **Combine**, \tilde{R} and the signature shares are additively combined to produce the Schnorr signature $\sigma = (\tilde{R}, z = \sum_{i \in \mathcal{S}} z_i)$.

This commit-reveal strategy ensures two security properties. First, requiring each participant to publish a commitment in the first round before revealing their R_k prevents a rushing adversary from adaptively choosing its R_j with respect to other participants’. Second, as later seen in our proofs of adaptive security, it allows the environment to effectively handle corruptions at any point in the signing process, without requiring the assumption of secure erasures.

<p>Setup(1^λ)</p> <hr/> $(\mathbb{G}, p, g) \leftarrow \text{GrGen}(1^\lambda)$ <i>//</i> Select two hash functions $\text{H}_{\text{cm}}, \text{H}_{\text{sig}} : \{0, 1\}^* \rightarrow \mathbb{Z}_p$ $\text{par} \leftarrow ((\mathbb{G}, p, g), \text{H}_{\text{cm}}, \text{H}_{\text{sig}})$ return par <hr/> <p>KeyGen($n, t + 1$)</p> <hr/> $x \leftarrow_{\$} \mathbb{Z}_p; \tilde{X} \leftarrow g^x$ $\{(i, x_i)\}_{i \in [n]} \leftarrow \text{SS.IssueShares}(x, n, t + 1)$ <i>//</i> Shamir secret sharing of x for $i \in [n]$ $\tilde{X}_i \leftarrow g^{x_i}$ return $(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_i\}_{i \in [n]})$ <hr/> <p>Sign(m, \mathcal{S})</p> <hr/> <i>//</i> Local signer has index k $r_k \leftarrow_{\$} \mathbb{Z}_p; R_k \leftarrow g^{r_k}$ $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$ $\rho_k \leftarrow \text{cm}_k$ $\text{st}_k \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S})$ return (ρ_k, st_k) <hr/> <p>Sign'($\text{st}_k, x_k, \{\rho_i\}_{i \in \mathcal{S}}$)</p> <hr/> $\text{parse } \{\text{cm}_i\}_{i \in \mathcal{S}} \leftarrow \{\rho_i\}_{i \in \mathcal{S}}$ $\text{parse } (\text{cm}_k, R_k, r_k, m, \mathcal{S}) \leftarrow \text{st}_k$ return \perp if $\text{cm}_k \notin \{\text{cm}_i\}_{i \in \mathcal{S}}$ $\rho'_k \leftarrow R_k$ $\text{st}'_k \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S}, \{\rho_i\}_{i \in \mathcal{S}})$ return (ρ'_k, st'_k)	<p>Sign''($\text{st}'_k, x_k, \{\rho'_i\}_{i \in \mathcal{S}}$)</p> <hr/> <i>//</i> Sign'' must be called once per st'_k $\text{parse } (\text{cm}_k, R_k, r_k, m, \mathcal{S}, \{\rho_i\}_{i \in \mathcal{S}}) \leftarrow \text{st}'_k$ $\text{parse } \{R_i\}_{i \in \mathcal{S}} \leftarrow \{\rho'_i\}_{i \in \mathcal{S}}$ return \perp if $R_k \notin \{R_i\}_{i \in \mathcal{S}}$ for $i \in \mathcal{S}$ do return \perp if $\text{cm}_i \neq \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$ $\tilde{R} \leftarrow \prod_{i \in \mathcal{S}} R_i$ $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$ $z_k \leftarrow r_k + c \lambda_k x_k$ <i>//</i> λ_k is the k^{th} Lagrange coefficient for \mathcal{S} $\rho''_k \leftarrow z_k$ return ρ''_k <hr/> <p>Combine($\{(\rho'_i, \rho''_i)\}_{i \in \mathcal{S}}$)</p> <hr/> $\text{parse } R_i \leftarrow \rho'_i, z_i \leftarrow \rho''_i, i \in \mathcal{S}$ $\tilde{R} \leftarrow \prod_{i \in \mathcal{S}} R_i; z \leftarrow \sum_{i \in \mathcal{S}} z_i$ $\sigma \leftarrow (\tilde{R}, z)$ return σ <hr/> <p>Verify(\tilde{X}, m, σ)</p> <hr/> $\text{parse } (\tilde{R}, z) \leftarrow \sigma$ $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$ if $\tilde{R} \tilde{X}^c = g^z$ return 1 return 0
---	--

Fig. 5. The Sparkle threshold signature scheme. The public parameters par are implicitly given as input to all algorithms and protocols. Sparkle assumes an external mechanism to choose the set $\mathcal{S} \subseteq \{1, \dots, n\}$ of signers, where $t + 1 \leq |\mathcal{S}| \leq n$. \mathcal{S} is required to be ordered to ensure consistency. Note that verification is identical to standard (single-party) Schnorr as in Definition 1.

Parameter Generation. On input the security parameter 1^λ , the setup algorithm runs $(\mathbb{G}, p, g) \leftarrow \text{GrGen}(1^\lambda)$, and selects hash functions $\text{H}_{\text{cm}}, \text{H}_{\text{sig}} : \{0, 1\}^* \rightarrow \mathbb{Z}_p$. Outputs public parameters $\text{par} \leftarrow ((\mathbb{G}, p, g), \text{H}_{\text{cm}}, \text{H}_{\text{sig}})$.

Key Generation. On input the number of signers n and the threshold $t + 1$, it first generates the secret key $x \leftarrow_s \mathbb{Z}_p$. It then performs Shamir secret sharing for x , outputting $\{(i, x_i)\}_{i \in [n]} \leftarrow_s \mathcal{SS}.\text{IssueShares}(x, n, t + 1)$. It then generates the joint public key as $\tilde{X} \leftarrow g^x$, and public keys for each participant as $\tilde{X}_i \leftarrow g^{x_i}$. It then returns $(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_i\}_{i \in [n]})$.

Signing Round 1 (Sign). On input the message m and signing set \mathcal{S} , each participant $P_i, i \in \mathcal{S}$, chooses $r_i \leftarrow_s \mathbb{Z}_p$, computes $R_i \leftarrow g^{r_i}$ and $\text{cm}_i \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$, and outputs their commitment cm_i .

Signing Round 2 (Sign'). On input commitments $\{\text{cm}_j\}_{j \in \mathcal{S}}$, each participant P_i outputs their nonce R_i .

Signing Round 3 (Sign''). On input nonces $\{R_j\}_{j \in \mathcal{S}}$, each participant P_i first checks that the commitments received in the first round are valid, i.e., $\text{cm}_j = \text{H}_{\text{cm}}(m, \mathcal{S}, R_j)$ for all $j \in \mathcal{S}$. If for some j' , $\text{cm}_{j'} \neq \text{H}_{\text{cm}}(m, \mathcal{S}, R_{j'})$, abort. It also aborts if it receives a set of $R_j, j \in \mathcal{S}$ which does not contain a commitment that it issued previously. Otherwise, P_i computes the aggregate nonce $\tilde{R} = \prod_{j \in \mathcal{S}} R_j$. It then computes the hash $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$ and $z_i \leftarrow r_i + c\lambda_i x_i$, where λ_i is the i^{th} Lagrange coefficient corresponding to \mathcal{S} . It outputs z_i .

Combining Signatures. On input nonces $\{R_j\}_{j \in \mathcal{S}}$ and partial signatures $\{z_j\}_{j \in \mathcal{S}}$, the combiner computes $\tilde{R} \leftarrow \prod_{j \in \mathcal{S}} R_j$ and $z \leftarrow \sum_{j \in \mathcal{S}} z_j$, and outputs $\sigma \leftarrow (\tilde{R}, z)$.

Verification. On input the joint public key \tilde{X} , a message m , and a signature $\sigma = (\tilde{R}, z)$, the verifier computes $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$ and accepts if $\tilde{R}\tilde{X}^c = g^z$.

Correctness of Sparkle is straightforward to verify. Note that verification of the threshold signature σ is identical to verification of a standard, key-prefixed Schnorr signature with respect to the aggregate nonce \tilde{R} and joint public key \tilde{X} , as in Definition 1.

6 Static Security Under Standard Assumptions

In this section, we show that Sparkle is statically secure under the discrete logarithm assumption (DL) in the random oracle model (ROM). Static security allows an adversary to control up to t parties, but they must be declared at the beginning of the protocol (Fig. 4.1).

Theorem 1. *Sparkle is statically secure under DL in the ROM.*

We formally prove Theorem 1 in Section 9.

Proof Outline. Let \mathcal{A} be a PPT adversary against the static unforgeability of Sparkle (Fig. 4.1). We construct a PPT reduction \mathcal{B} against the DL assumption (Fig. 3) that uses \mathcal{A} as a subroutine as follows. \mathcal{B} takes as input a DL challenge \hat{X} and aims to output \hat{x} such that $\hat{X} = g^{\hat{x}}$. \mathcal{B} sets the joint public key $\tilde{X} \leftarrow \hat{X}$ and performs a standard simulation of Shamir secret sharing. \mathcal{B} then returns \tilde{X} , public key shares $\{\tilde{X}_i\}_{i \in [n]}$, and secret key shares $\{x_j\}_{j \in \text{cor}}$ to \mathcal{A} .

\mathcal{B} simulates signing without knowing the secret keys $\{x_k\}_{k \in \text{hon}}$ of the honest parties as follows. \mathcal{B} can simulate R_k for all $k \in \text{hon}$ as $R_k \leftarrow g^{z_k} \tilde{X}_k^{-c\lambda_k}$ for random $z_k \leftarrow^s \mathbb{Z}_p$ so that the partial signature z_k output in Round 3 verifies. However, c must equal $\text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$ for aggregate \tilde{R} . Luckily, \mathcal{B} can compute $\tilde{R} = \prod_{i \in \mathcal{S}} R_i$, where \mathcal{S} is the signing set, by extracting R_j for all $j \in \text{cor}$ from \mathcal{A} 's $\text{H}_{\text{cm}}(m, \mathcal{S}, R_j)$ queries in Round 1. So, \mathcal{B} samples a random value for $c \leftarrow^s \mathbb{Z}_p$, computes R_k for all $k \in \text{hon}$ as above, and programs $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$ – all before \mathcal{A} sees the honest R_k 's output at the end of Round 2. This leaves just one problem: \mathcal{B} needs to output $\text{cm}_k = \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$ in Round 1 *before* being able to carry out the above steps. But \mathcal{B} can simply output a random value $\text{cm}_k \leftarrow^s \mathbb{Z}_p$ in Round 1 and program $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$ in Round 2.

\mathcal{B} then *rewinds* \mathcal{A} , programming the single point $c' \leftarrow^s \mathbb{Z}_p$; $c' \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$ on the second iteration of \mathcal{A} . By the local forking lemma [6], \mathcal{A} 's two forgeries $(R^*, m^*, z^*), (R', m', z')$ satisfy $(R^*, m^*) = (R', m')$ with non-negligible probability. Thus, \mathcal{B} can compute $\hat{x} = \frac{z^* - z'}{c^* - c'}$ and win the DL game.

7 Adaptive Security Against up to $t/2$ Corruptions

In this section, we prove the adaptive security of Sparkle against up to $t/2$ corruptions under the algebraic one-more discrete logarithm assumption (AOMDL) in the random oracle model (ROM). The reason the allowed corruption is $t/2$ is that, in order to extract an AOMDL solution, the reduction needs to rewind the adversary once, and there is no guarantee that the adversary will corrupt the same set of parties after the fork as it did during the first iteration of the adversary. When the adversary can corrupt only $t/2$ parties, this causes no issues, as the total number of corruptions over both iterations does not exceed t . If the adversary could corrupt more parties, then the reduction would query its discrete logarithm oracle more than t times and would lose the $t + 1$ -aomdl game.

Theorem 2. *Sparkle is adaptively secure against $t/2$ corruptions under AOMDL in the ROM.*

We formally prove Theorem 2 in Section 9.

Proof Outline. Let \mathcal{A} be a PPT adversary against the adaptive unforgeability of Sparkle (Fig. 4.1). We construct a PPT reduction \mathcal{B} against the $t + 1$ -aomdl assumption (Fig. 3) that uses \mathcal{A} as a subroutine as follows. \mathcal{B} takes as input a $t + 1$ -aomdl challenge (Y_0, Y_1, \dots, Y_t) and aims to output y_i such that $Y_i = g^{y_i}$ for all $i \in [0..t]$ without querying its DL oracle more than t times. For all $i \in [n]$, \mathcal{B} sets the public key share as $\tilde{X}_i = Y_0 Y_1^i \dots Y_t^{i^t}$. The joint public key is $\tilde{X} = Y_0$. \mathcal{B} queries its DL oracle on \tilde{X}_j (with representation $(1, j, \dots, j^t)$) to get x_j for the initial corrupted set cor . \mathcal{B} then returns $(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_j\}_{j \in \text{cor}})$ to \mathcal{A} .

\mathcal{B} 's simulation of signing is similar to the proof of static security. In particular, \mathcal{B} again simulates R_k for honest $k \in \text{hon}$ as $R_k \leftarrow g^{z_k} \tilde{X}_k^{-c\lambda_k}$ in Round 2. If \mathcal{A} corrupts an honest party k after Round 2 has begun, then \mathcal{B} queries its DL

oracle on \tilde{X}_k (with rep. $(1, k, \dots, k^t)$) to get x_k , computes $r_k \leftarrow z_k - c\lambda_k x_k$, and returns x_k along with the honest party's state $\text{st}_{k,\text{ssid}} = \{\text{cm}_k, R_k, r_k, \dots\}$ to \mathcal{A} . However, \mathcal{A} may choose to corrupt some honest k *before* Round 2, or even before it outputs its own cm_j 's in Round 1. In this case, \mathcal{B} samples a random $r_k \leftarrow^* \mathbb{Z}_p$, sets $R_k \leftarrow g^{r_k}$ and programs $\text{cm}_k \leftarrow H_{\text{cm}}(m, \mathcal{S}, R_k)$ for the random cm_k it output in Round 1. It then queries its DL oracle on \tilde{X}_k (with rep. $(1, k, \dots, k^t)$) to get x_k and returns it and $\text{st}_{k,\text{ssid}} = \{\text{cm}_k, R_k, r_k, \dots\}$ to \mathcal{A} . As in the proof of static security, \mathcal{B} rewinds \mathcal{A} in order to extract $x = y_0$ from \mathcal{A} 's two forgeries. Assume w.l.o.g. that \mathcal{A} corrupts t parties over the two iterations. (\mathcal{A} can corrupt up to $t/2$ parties in each iteration, and \mathcal{B} can corrupt the remaining itself.) For simplicity, say the corrupt indices are $\text{cor} = \{1, \dots, t\}$. Then \mathcal{B} has made t DL queries on $g^{x_k} = Y_0 Y_1^k \dots Y_t^{k^t}$. This forms a system of linear equations:

$$\begin{pmatrix} x \\ x_1 \\ \vdots \\ x_t \end{pmatrix} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 1 & 1 & \dots & 1^t \\ \vdots & \vdots & \ddots & \vdots \\ 1 & t & \dots & t^t \end{pmatrix} \begin{pmatrix} y_0 \\ y_1 \\ \vdots \\ y_t \end{pmatrix}$$

This is a Vandermonde matrix and is therefore invertible. \mathcal{B} knows (x, x_1, \dots, x_t) , and so can solve for (y_0, y_1, \dots, y_t) to win the $t + 1$ -aomdl game.

8 Adaptive Security Against t Corruptions

We now prove our strongest adaptivity result: that **Sparkle** is secure against t corruptions. In particular, if exactly $t + 1$ parties engage in signing, all but one of them could be malicious and the unforgeability of **Sparkle** would still hold. We prove this result under the AOMDL assumption in the AGM with random oracles. As a warm-up, we provide a proof of static security under the DL assumption in the AGM and ROM in Appendix A.

Theorem 3. *Sparkle is adaptively secure against t corruptions under the AOMDL assumption in the AGM and ROM.*

We formally prove Theorem 3 in Section 9.

Proof Outline. Let \mathcal{A} be an algebraic adversary against the adaptive unforgeability of **Sparkle** (Fig. 4.1). We construct a PPT reduction \mathcal{B} against the $t + 1$ -aomdl assumption (Fig. 3) that uses \mathcal{A} as a subroutine as follows. \mathcal{B} takes as input a $t + 1$ -aomdl challenge (Y_0, Y_1, \dots, Y_t) and aims to output y_i such that $Y_i = g^{y_i}$ for all $i \in [0..t]$ without querying its DL oracle more than t times. \mathcal{B} simulates key generation, signing, and corruption as in the $t/2$ -adaptive proof, but does not rewind \mathcal{A} , so \mathcal{A} may corrupt a full t parties.

\mathcal{A} 's forgery (m^*, \tilde{R}^*, z^*) satisfies $\tilde{R}^* = g^{z^*} \tilde{X}^{-c^*}$, where \mathcal{A} provided a representation of \tilde{R}^* when querying $c^* = \text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$:

$$\tilde{R}^* = g^{\gamma^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{q_S} R_{i,1}^{\rho_{i,1}^*} \dots R_{i,n}^{\rho_{i,n}^*}$$

for q_S signing queries. Each $R_{i,k}$ verifies as $R_{i,k} = g^{z_{i,k}} \tilde{X}_k^{-c_i \lambda_{i,k}}$, where $c_i = \text{H}_{\text{sig}}(\tilde{X}, m_i, \tilde{R}_i)$. Equating the two expressions for \tilde{R}^* and rearranging, we have:

$$g^{z^*} g^{-\gamma^*} \prod_{i=1}^{q_S} g^{-z_{i,1} \rho_{i,1}^*} \dots g^{-z_{i,n} \rho_{i,n}^*} = \tilde{X}^{c^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{q_S} (\tilde{X}_1^{-c_i \lambda_{i,1}})^{\rho_{i,1}^*} \dots (\tilde{X}_n^{-c_i \lambda_{i,n}})^{\rho_{i,n}^*}$$

\mathcal{B} queries its DL oracle on \tilde{X}_j (with representation $(1, j, \dots, j^t)$) to obtain $\{x_j\}_{j \in \text{cor}}$ for t corrupt parties. For all $k \in [n]$, $\tilde{X}_k = \tilde{X} Y_1^k Y_2^{k^2} \dots Y_t^{k^t}$, and for all $i \in [t]$, $Y_i = \tilde{X}^{L'_{0,i}} \prod_{j \in \text{cor}} g^{x_j L'_{j,i}}$, where $L'_{j,i}$ is the i^{th} coefficient of the Lagrange polynomial $L'_j(Z)$ for the set $0 \cup \text{cor}$. Plugging these in and rearranging, we have:

$$g^{\eta^*} = \tilde{X}^{c^* + \xi^*} \prod_{k=1}^n \tilde{X}^{\mu_k^*} g^{\nu_k^*}$$

where $\eta^* = z^* - \gamma^* - \sum_{i=1}^{q_S} (z_{i,1} \rho_{i,1}^* + \dots + z_{i,n} \rho_{i,n}^*)$, $\mu_k^* = 1 + \sum_{i=1}^t L'_{0,i} k^i$, and $\nu_k^* = \sum_{i=1}^t \left(\sum_{j \in \text{cor}} x_j L'_{j,i} \right) k^i$. Then:

$$x = \frac{\eta^* - \sum_{k=1}^n \nu_k^*}{c^* + \xi^* + \sum_{k=1}^n \mu_k^*}$$

\mathcal{A} fixed \tilde{R}^* and thus η^* , $\{\nu_i^*\}_{i \in [n]}$, ξ^* , $\{\mu_i^*\}_{i \in [n]}$ as it queried $\text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$ to receive random c^* . Thus, the denominator is nonzero with overwhelming probability and \mathcal{B} can solve for x . \mathcal{B} can then compute (y_0, y_1, \dots, y_t) as in the $t/2$ -adaptive proof to win the $t+1$ -aomdl game.

9 Static and Adaptive Security Proofs

In this section, we provide formal proofs of static and adaptive security for Sparkle. In particular, we prove the following: (1) static security (against up to t corruptions) (Theorem 1) under DL+ROM, (2) adaptive security against up to $t/2$ corruptions (Theorem 2) under AOMDL+ROM, and (3) adaptive security against up to t corruptions (Theorem 3) under AOMDL+AGM+ROM.

9.1 Proof of Static Security

Proof. (of Theorem 1.) Let \mathcal{A} be a PPT adversary attempting to break the static unforgeability of Sparkle (Fig. 4.1) that makes up to q_H queries to H_{cm} and H_{sig} ,

and q_S queries to its signing oracles. Without loss of generality, we assume \mathcal{A} queries H_{sig} on its forgery (\tilde{X}, m^*, R^*) . Then, let $q = q_H + q_S + 1$. We construct a PPT reduction \mathcal{B} against the DL assumption (Fig. 3) that uses \mathcal{A} as a subroutine such that

$$\text{Adv}_{\mathcal{A}}^{st\text{-}sec}(\lambda, 1) \leq \sqrt{q \text{Adv}_{\mathcal{B}}^{\text{dl}}(\lambda) + \text{negl}(\lambda)}$$

Here, $\tau = 1$ to allow \mathcal{A} to corrupt $t/\tau = t$ parties. The reduction \mathcal{B} runs \mathcal{A} two times. On the second iteration, \mathcal{B} programs H_{sig} to output a different random value on a single point so that it can extract a discrete logarithm solution from \mathcal{A} 's two forgeries. \mathcal{B} perfectly simulates $\text{Game}_{\mathcal{A}}^{\text{UF}}(\lambda, 1)$. However, \mathcal{B} can only extract a discrete logarithm if \mathcal{A} 's forgery (m^*, R^*, z^*) at the end of each iteration verifies and includes the same nonce R^* . By the local forking lemma [6], this occurs with probability less than $\frac{1}{q}(\text{Adv}_{\mathcal{A}}^{st\text{-}sec}(\lambda, 1))^2$.

The Reduction \mathcal{B} : We define the reduction \mathcal{B} playing game $\text{Game}_{\mathcal{B}}^{\text{dl}}(\lambda)$ as follows. \mathcal{B} is responsible for simulating key generation and oracle responses for queries to $\mathcal{O}^{\text{Sign}}, \mathcal{O}^{\text{Sign}'}, \mathcal{O}^{\text{Sign}''}, H_{\text{cm}}$, and H_{sig} . Let $Q_{\text{cm}}, Q_{\text{sig}}$ be the set of $H_{\text{cm}}, H_{\text{sig}}$ queries and their responses, respectively. \mathcal{B} initializes them to the empty set and maintains them across both iterations of the adversary. \mathcal{B} may program the random oracles $H_{\text{cm}}, H_{\text{sig}}$. Let Q be the set of messages that have been queried in $\mathcal{O}^{\text{Sign}}$ as in $\text{Game}_{\mathcal{A}}^{\text{UF}}(\lambda, 1)$. \mathcal{B} initializes Q to the empty set. At the beginning of the second iteration of \mathcal{A} , \mathcal{B} resets Q to the empty set.

DL Input. \mathcal{B} takes as input the group description $\mathcal{G} = (\mathbb{G}, p, g)$ and a DL challenge \tilde{X} . \mathcal{B} aims to output \hat{x} such that $\tilde{X} = g^{\hat{x}}$.

Static Corruption. \mathcal{B} runs $\mathcal{A}()$. \mathcal{A} chooses the total number of potential signers n , the threshold $t + 1, 1 \leq t + 1 \leq n$, and the set of corrupt parties $\text{cor} \leftarrow \{j\}, |\text{cor}| \leq t$, which are fixed for the rest of the protocol. \mathcal{B} sets $\text{hon} \leftarrow [n] \setminus \text{cor}$ and must reveal the secret keys of the corrupt parties to \mathcal{A} , which \mathcal{B} does in the next step.

Simulating KeyGen. \mathcal{B} simulates the key generation algorithm (Fig. 5) using its DL challenge \tilde{X} as follows.

1. \mathcal{B} sets the joint public key $\tilde{X} \leftarrow \tilde{X}$.
2. \mathcal{B} simulates a Shamir secret sharing of the discrete logarithm of \tilde{X} by performing the following steps. (See Section 3 for notation.) Assume without loss of generality that $|\text{cor}| = t$.
 - (a) \mathcal{B} samples t random values $x_j \leftarrow_s \mathbb{Z}_p$ for $j \in \text{cor}$.
 - (b) Let f be the polynomial whose constant term is the challenge $f(0) = \hat{x}$ and for which $f(j) = x_j$ for all $j \in \text{cor}$. \mathcal{B} computes the $t + 1$ Lagrange polynomials $\{L'_0(Z), \{L'_j(Z)\}_{j \in \text{cor}}\}$ relating to the set (of x -coordinates) $0 \cup \text{cor}$.
 - (c) For all $1 \leq i \leq t$, \mathcal{B} computes

$$Y_i = \tilde{X}^{L'_{0,i}} \prod_{j \in \text{cor}} g^{x_j L'_{j,i}}$$

where $L'_{j,i}$ is the i^{th} coefficient of $L'_j(Z) = L'_{j,0} + L'_{j,1}Z + \dots + L'_{j,t}Z^t$.
(d) For $1 \leq i \leq n$, \mathcal{B} computes

$$\tilde{X}_i = \tilde{X} Y_1^i Y_1^{i^2} \dots Y_t^{i^t}$$

which is implicitly equal to $g^{f(i)}$.

The joint public key is $\tilde{X} = g^{f(0)} = \dot{X}$ with corresponding secret key $x = \dot{x}$. \mathcal{B} runs $\mathcal{A}^{\mathcal{O}^{\text{Sign, Sign}', \text{Sign}''}}(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_j\}_{j \in \text{cor}})$.

Simulating Random Oracle Queries. \mathcal{B} handles \mathcal{A} 's random oracle queries throughout the protocol by lazy sampling, as follows.

H_{cm} : When \mathcal{A} queries H_{cm} on (m, \mathcal{S}, R) , \mathcal{B} checks whether $(m, \mathcal{S}, R, \text{cm}) \in Q_{\text{cm}}$ and, if so, returns cm . Else, \mathcal{B} samples $\text{cm} \leftarrow_{\$} \mathbb{Z}_p$, appends $(m, \mathcal{S}, R, \text{cm})$ to Q_{cm} , and returns cm .

H_{sig} : When \mathcal{A} queries H_{sig} on (X, m, R) , \mathcal{B} checks whether $(X, m, R, c) \in Q_{\text{sig}}$ and, if so, returns c . Else, \mathcal{B} samples $c \leftarrow_{\$} \mathbb{Z}_p$, appends (X, m, R, c) to Q_{sig} , and returns c .

Simulating Sparkle Signing. \mathcal{B} handles \mathcal{A} 's signing queries as follows.

Round 1 ($\mathcal{O}^{\text{Sign}}$): In the first round of signing for session ssid , each party P_i in the signing set \mathcal{S} sends a commitment cm_i . When \mathcal{A} queries $\mathcal{O}^{\text{Sign}}$ on (k, m, \mathcal{S}) for honest $k \in \text{hon}$, \mathcal{B} samples $\text{cm}_k \leftarrow_{\$} \mathbb{Z}_p$, appends $(m, \mathcal{S}, \cdot, \text{cm}_k)$ to Q_{cm} , sets $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, \cdot, \cdot, m, \mathcal{S})$, and returns cm_k to \mathcal{A} .

Round 2 ($\mathcal{O}^{\text{Sign}'}$): In the second round of signing for session ssid , each party P_i in the signing set \mathcal{S} takes as input the set of commitments $\{\text{cm}_i\}_{i \in \mathcal{S}}$ and reveals its nonce commitment R_i such that $\text{cm}_i = H_{\text{cm}}(m, \mathcal{S}, R_i)$. When \mathcal{A} queries $\mathcal{O}^{\text{Sign}'}$ on $(k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$ for $k \in \text{hon}$, then for all $j \in \mathcal{S} \cap \text{cor}$, \mathcal{B} looks up cm_j for a record $(m, \mathcal{S}, R_j, \text{cm}_j) \in Q_{\text{cm}}$.

If there exists $j' \in \mathcal{S}$ such that $(m', \mathcal{S}', \cdot, \text{cm}_{j'}) \in Q_{\text{cm}}$ and $(m', \mathcal{S}') \neq (m, \mathcal{S})$, or if there exists some $j' \in \mathcal{S}$ for which no record $(m, \mathcal{S}, \cdot, \text{cm}_{j'}) \in Q_{\text{cm}}$ exists, then \mathcal{B} chooses $R_k \leftarrow_{\$} g^{r_k}$ randomly, updates $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S})$, programs $\text{cm}_k \leftarrow H_{\text{cm}}(m, \mathcal{S}, R_k)$ and returns R_k .

Else if this is the first query in the signing session, \mathcal{B} samples $c \leftarrow_{\$} \mathbb{Z}_p$. For all $k \in \text{hon}$, \mathcal{B} samples $z_k \leftarrow_{\$} \mathbb{Z}_p$, computes $R_k \leftarrow g^{z_k} \tilde{X}_k^{-c \lambda_k}$ (λ_k is the Lagrange coefficient for the set \mathcal{S}), and programs $\text{cm}_k \leftarrow H_{\text{cm}}(m, \mathcal{S}, R_k)$ (updating $(m, \mathcal{S}, R_k, \text{cm}_k) \in Q_{\text{cm}}$). Then \mathcal{B} computes $\tilde{R} = \prod_{i \in \mathcal{S}} R_i$ and programs $c \leftarrow H_{\text{sig}}(\tilde{X}, m, \tilde{R})$. (However, if \mathcal{A} has already queried H_{sig} on $(\tilde{X}, m, \tilde{R})$, then \mathcal{B} aborts.)

Finally, \mathcal{B} sets $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c \lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$, and returns R_k to \mathcal{A} . If this is not the first query, \mathcal{B} looks up $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c \lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$ and returns R_k .

Round 3 ($\mathcal{O}^{\text{Sign}''}$): In the third round of signing for session ssid , each party P_i in the signing set \mathcal{S} produces a partial signature on the message m . When \mathcal{A} queries

$\mathcal{O}^{\text{Sign}''}$ on $(k, \{R_i\}_{i \in \mathcal{S}})$ for $k \in \text{hon}$, \mathcal{B} looks up $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$, checks whether $\text{cm}_i = \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$ for all $i \in \mathcal{S}$ and returns \perp if not. If $\text{cm}_i = \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$ but \mathcal{A} never queried H_{cm} on input (m, \mathcal{S}, R_i) , \mathcal{B} aborts. Else, \mathcal{B} sets $\text{st}_{k, \text{ssid}} \leftarrow ()$ and returns z_k .

Output. At the end of the signing rounds, \mathcal{A} produces a forgery $(m^*, \sigma^*) = (m^*, (R^*, z^*))$. \mathcal{A} wins if $\text{Verify}(\tilde{X}, m^*, \sigma^*) = 1$ and $m^* \notin Q$.

Extracting the Discrete Logarithm of \tilde{X} . \mathcal{B} 's simulation of key generation and signing is perfect, and \mathcal{B} aborts with negligible probability. Indeed, \mathcal{B} aborts in Round 3 if \mathcal{A} reveals R_j such that $\text{cm}_j = \text{H}_{\text{cm}}(m, \mathcal{S}, R_j)$ but \mathcal{A} never queried H_{cm} on (m, \mathcal{S}, R_j) . This requires \mathcal{A} to have guessed cm_j ahead of time, which occurs with negligible probability $1/p$.

\mathcal{B}_2 also aborts in Round 2 if \mathcal{A} had previously queried H_{sig} on $(\tilde{X}, m, \tilde{R})$. In that case, \mathcal{B} had returned a random $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$, so the reduction fails. However, this implies that \mathcal{A} guessed R_k before \mathcal{B} revealed it, which occurs with negligible probability $1/p$.

It remains to show that \mathcal{B} can extract the discrete logarithm of \tilde{X} from \mathcal{A} 's two valid forgeries. \mathcal{A} 's first forgery (m^*, R^*, z^*) satisfies $R^* \tilde{X}^{c^*} = g^{z^*}$, where $c^* = \text{H}_{\text{sig}}(\tilde{X}, m^*, R^*)$. Here, z^* does not suffice for \mathcal{B} to extract the discrete logarithm of \tilde{X} because it does not necessarily know the discrete logarithm of R^* . Thus, \mathcal{B} chooses $c' \leftarrow \mathbb{Z}_p$ and programs H_{sig} to output c' on input (\tilde{X}, m^*, R^*) . \mathcal{B} resets Q to the empty set, but the sets $Q_{\text{cm}}, Q_{\text{sig}}$ are kept for the second iteration of the adversary. \mathcal{B} then runs \mathcal{A} again on the same random coins.

After the second iteration, suppose \mathcal{A} terminates with (m', R', z') . If $(m', R') = (m^*, R^*)$ and \mathcal{A} 's forgeries both verify, then \mathcal{B} returns $\hat{x} = \frac{z^* - z'}{c^* - c'}$ such that $\tilde{X} = g^{\hat{x}}$ and wins $\text{Game}_{\mathcal{B}}^{\text{dl}}(\lambda)$. If $(m', R') \neq (m^*, R^*)$ or \mathcal{A} 's forgery does not verify, then \mathcal{B} must abort. By the local forking lemma [6], this happens with probability less than $\frac{1}{q}(\text{Adv}_{\mathcal{A}}^{\text{st-sec}}(\lambda, 1))^2$. If \mathcal{A} succeeds having not queried H_{cm} on (m, \mathcal{S}, R_j) , then \mathcal{B} aborts. This occurs with probability less than $\frac{q_H}{p}$.

Thus,

$$\frac{1}{q}(\text{Adv}_{\mathcal{A}}^{\text{st-sec}}(\lambda, 1))^2 \leq \text{Adv}_{\mathcal{B}}^{\text{dl}}(\lambda) + \text{negl}(\lambda)$$

9.2 Proof of Adaptive Security for up to $t/2$ Corruptions

Proof. (of Theorem 2.) Let \mathcal{A} be a PPT adversary attempting to break the adaptive unforgeability of Sparkle (Fig. 4.1) that makes up to q_H queries to H_{cm} and H_{sig} , and q_S queries to its signing oracles. Without loss of generality, we assume \mathcal{A} queries H_{sig} on its forgery (\tilde{X}, m^*, R^*) . Then, let $q = q_H + q_S + 1$. We construct a PPT reduction \mathcal{B} against the $t + 1$ -aomdl assumption (Fig. 3) that uses \mathcal{A} as a subroutine such that

$$\text{Adv}_{\mathcal{A}}^{\text{adp-sec}}(\lambda, 2) \leq \sqrt{q \text{Adv}_{\mathcal{B}}^{t+1\text{-aomdl}}(\lambda) + \text{negl}(\lambda)}$$

Here, $\tau = 2$ to restrict \mathcal{A} to corrupt $t/2$ parties. The reduction \mathcal{B} runs \mathcal{A} two times. Over the two iterations, \mathcal{B} makes no more than t queries to its discrete

logarithm oracle \mathcal{O}^{dl} and aims to output $t + 1$ discrete logarithms that constitute a valid solution to the AOMDL challenge. If \mathcal{B} makes fewer than t queries while responding to \mathcal{A} 's oracle queries, then it makes the additional queries necessary to extract an AOMDL solution. On the second iteration of \mathcal{A} , \mathcal{B} programs H_{sig} to output a different random value on a single point so that it can extract one of the $t + 1$ discrete logarithm solutions from \mathcal{A} 's two forgeries. \mathcal{B} perfectly simulates $\text{Game}_{\mathcal{A}}^{\text{adp-UF}}(\lambda, 2)$. However, \mathcal{B} can only extract a solution if \mathcal{A} 's forgery (m^*, R^*, z^*) at the end of each iteration verifies and includes the same nonce R^* . By the local forking lemma [6], this occurs with probability less than $\frac{1}{q}(\text{Adv}_{\mathcal{A}}^{\text{adp-sec}}(\lambda, 2))^2$.

The Reduction \mathcal{B} : We define the reduction \mathcal{B} playing game $\text{Game}_{\mathcal{B}}^{t+1\text{-aomdl}}(\lambda)$ as follows. \mathcal{B} is responsible for simulating key generation and oracle responses for queries to $\mathcal{O}^{\text{Sign}}, \mathcal{O}^{\text{Sign}'}, \mathcal{O}^{\text{Sign}''}, \mathcal{O}^{\text{Corrupt}}, \text{H}_{\text{cm}}$, and H_{sig} . Let $\text{Q}_{\text{cm}}, \text{Q}_{\text{sig}}$ be the set of $\text{H}_{\text{cm}}, \text{H}_{\text{sig}}$ queries and their responses, respectively. \mathcal{B} initializes them to the empty set and maintains them across both iterations of the adversary. \mathcal{B} may program the random oracles $\text{H}_{\text{cm}}, \text{H}_{\text{sig}}$. Let Q be the set of messages that have been queried in $\mathcal{O}^{\text{Sign}}$ as in $\text{Game}_{\mathcal{A}}^{\text{adp-UF}}(\lambda, 2)$. \mathcal{B} initializes Q to the empty set. At the beginning of the second iteration of \mathcal{A} , \mathcal{B} resets Q to the empty set.

AOMDL Input. \mathcal{B} takes as input the group description $\mathcal{G} = (\mathbb{G}, p, g)$ and an AOMDL challenge of $t + 1$ values (Y_0, \dots, Y_t) . As in $\text{Game}_{\mathcal{B}}^{t+1\text{-aomdl}}(\lambda)$, \mathcal{B} has access to a discrete logarithm oracle \mathcal{O}^{dl} , which it may query up to t times. \mathcal{B} aims to output (y_0, \dots, y_t) such that $Y_i = g^{y_i}$ for all $0 \leq i \leq t$.

Initial Corruption. \mathcal{B} runs $\mathcal{A}()$. \mathcal{A} chooses the total number of potential signers n , the threshold $t + 1$, $1 \leq t + 1 \leq n$, and the initial set of corrupt parties $\text{cor} \leftarrow \{j\}, |\text{cor}| \leq t/2$. \mathcal{B} sets $\text{hon} \leftarrow [n] \setminus \text{cor}$ and must reveal the secret keys of the corrupt parties to \mathcal{A} , which \mathcal{B} does in the next step.

Simulating KeyGen. \mathcal{B} simulates the key generation algorithm (Fig. 5) using its AOMDL challenge (Y_0, \dots, Y_t) as follows. For all $1 \leq i \leq n$, \mathcal{B} sets the public key share as

$$\tilde{X}_i = Y_0 Y_1^i \cdots Y_t^{i^t}$$

which is implicitly equal to $g^{f(i)}$. The joint public key is $\tilde{X} = g^{f(0)} = Y_0$ with corresponding secret key $x = y_0$. \mathcal{B} obtains the initial corrupt secret key shares by querying $x_j = f(j) \leftarrow \mathcal{O}^{\text{dl}}(\tilde{X}_j)$ (with representation $(1, j, \dots, j^t)$) for all $j \in \text{cor}$. \mathcal{B} runs $\mathcal{A}^{\mathcal{O}^{\text{Sign}}, \text{Sign}', \text{Sign}'', \text{Corrupt}}(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_j\}_{j \in \text{cor}})$.

Simulating Random Oracle Queries. \mathcal{B} handles \mathcal{A} 's random oracle queries throughout the protocol by lazy sampling, as follows.

H_{cm} : When \mathcal{A} queries H_{cm} on (m, \mathcal{S}, R) , \mathcal{B} checks whether $(m, \mathcal{S}, R, \text{cm}) \in \text{Q}_{\text{cm}}$ and, if so, returns cm . Else, \mathcal{B} samples $\text{cm} \leftarrow_{\mathcal{S}} \mathbb{Z}_p$, appends $(m, \mathcal{S}, R, \text{cm})$ to Q_{cm} , and returns cm .

$\underline{H}_{\text{sig}}$: When \mathcal{A} queries H_{sig} on (X, m, R) , \mathcal{B} checks whether $(X, m, R, c) \in \mathcal{Q}_{\text{sig}}$ and, if so, returns c . Else, \mathcal{B} samples $c \leftarrow_{\$} \mathbb{Z}_p$, appends (X, m, R, c) to \mathcal{Q}_{sig} , and returns c .

Simulating Sparkle Signing. \mathcal{B} handles \mathcal{A} 's signing queries as follows.

Round 1 ($\mathcal{O}^{\text{Sign}}$): In the first round of signing for session ssid , each party P_i in the signing set \mathcal{S} sends a commitment cm_i . When \mathcal{A} queries $\mathcal{O}^{\text{Sign}}$ on (k, m, \mathcal{S}) for honest $k \in \text{hon}$, \mathcal{B} samples $\text{cm}_k \leftarrow_{\$} \mathbb{Z}_p$, appends $(m, \mathcal{S}, \cdot, \text{cm}_k)$ to \mathcal{Q}_{cm} , sets $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, \cdot, \cdot, m, \mathcal{S})$, and returns cm_k to \mathcal{A} .

Round 2 ($\mathcal{O}^{\text{Sign}'}$): In the second round of signing for session ssid , each party P_i in the signing set \mathcal{S} takes as input the set of commitments $\{\text{cm}_i\}_{i \in \mathcal{S}}$ and reveals its nonce commitment R_i such that $\text{cm}_i = \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$. When \mathcal{A} queries $\mathcal{O}^{\text{Sign}'}$ on $(k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$ for $k \in \text{hon}$, then for all $j \in \mathcal{S} \cap \text{cor}$, \mathcal{B} looks up cm_j for a record $(m, \mathcal{S}, R_j, \text{cm}_j) \in \mathcal{Q}_{\text{cm}}$.

If there exists $j' \in \mathcal{S}$ such that $(m', \mathcal{S}', \cdot, \text{cm}_{j'}) \in \mathcal{Q}_{\text{cm}}$ and $(m', \mathcal{S}') \neq (m, \mathcal{S})$, or if there exists some $j' \in \mathcal{S}$ for which no record $(m, \mathcal{S}, \cdot, \text{cm}_{j'}) \in \mathcal{Q}_{\text{cm}}$ exists, then \mathcal{B} chooses $R_k \leftarrow g^{r_k}$ randomly, updates $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S})$, programs $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$ and returns R_k .

Else if this is the first query in the signing session, \mathcal{B} samples $c \leftarrow_{\$} \mathbb{Z}_p$. For all $k \in \text{hon}$, \mathcal{B} samples $z_k \leftarrow_{\$} \mathbb{Z}_p$, computes $R_k \leftarrow g^{z_k} \tilde{X}_k^{-c\lambda_k}$ (λ_k is the Lagrange coefficient for the set \mathcal{S}), and programs $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$ (updating $(m, \mathcal{S}, R_k, \text{cm}_k) \in \mathcal{Q}_{\text{cm}}$). Then \mathcal{B} computes $\tilde{R} = \prod_{i \in \mathcal{S}} R_i$ and programs $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$. (However, if \mathcal{A} has already queried H_{sig} on $(\tilde{X}, m, \tilde{R})$, then \mathcal{B} aborts.)

Finally, \mathcal{B} sets $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_k\}_{k \in \mathcal{S}})$, and returns R_k to \mathcal{A} . If this is not the first query, \mathcal{B} looks up $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_k\}_{k \in \mathcal{S}})$ and returns R_k .

Round 3 ($\mathcal{O}^{\text{Sign}''}$): In the third round of signing for session ssid , each party P_i in the signing set \mathcal{S} produces a partial signature on the message m . When \mathcal{A} queries $\mathcal{O}^{\text{Sign}''}$ on $(k, \{R_i\}_{i \in \mathcal{S}})$ for $k \in \text{hon}$, \mathcal{B} looks up $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$, checks whether $\text{cm}_i = \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$ for all $i \in \mathcal{S}$ and returns \perp if not. If $\text{cm}_i = \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$ but \mathcal{A} never queried H_{cm} on input (m, \mathcal{S}, R_i) , \mathcal{B} aborts. Else, \mathcal{B} sets $\text{st}_{k, \text{ssid}} \leftarrow ()$ and returns z_k .

Simulating Corruption Queries ($\mathcal{O}^{\text{Corrupt}}$): \mathcal{A} may at any time corrupt an honest party k by querying $\mathcal{O}^{\text{Corrupt}}(k)$. Upon receiving a corruption query, \mathcal{B} first checks that $k \in \text{hon}$, returning \perp if not. Otherwise, \mathcal{B} queries its DL oracle \mathcal{O}^{dl} on $\tilde{X}_k = g^{f(k)}$ (with representation $(1, k, \dots, k^t)$) to obtain the secret key $x_k = f(k)$. Then, for each $\text{st}_{k, \text{ssid}}$, \mathcal{B} does the following:

- If $\text{st}_{k, \text{ssid}} = (\text{cm}_k, \cdot, \cdot, m, \mathcal{S})$, then \mathcal{B} chooses $r_k \leftarrow_{\$} \mathbb{Z}_p$, sets $R_k \leftarrow g^{r_k}$, and programs $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$. It then updates $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S})$.
- If $\text{st}_{k, \text{ssid}} = (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$, then \mathcal{B} computes $r_k = z_k - c\lambda_k x_k$ (now that it knows x_k) and updates $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$.

- If $\text{st}_{k, \text{ssid}} = ()$, then return $()$. This case occurs if signing session ssid has already been completed (i.e., a valid signature was issued).

Finally, \mathcal{B} sets $\text{hon} \leftarrow \text{hon} \setminus \{k\}$ and $\text{cor} \leftarrow \text{cor} \cup \{k\}$ and returns $x_k, \{\text{st}_{k, \ell}\}_{\ell \in [\text{ssid}]}$ to \mathcal{A} .

Output. At the end of the signing rounds, \mathcal{A} produces a forgery $(m^*, \sigma^*) = (m^*, (R^*, z^*))$. \mathcal{A} wins if $\text{Verify}(\tilde{X}, m^*, \sigma^*) = 1, m^* \notin Q$ and $|\text{cor}| \leq t/2$.

Extracting the Discrete Logarithm of Y_0 . \mathcal{B} 's simulation of key generation and signing is perfect, and \mathcal{B} aborts with negligible probability. Indeed, \mathcal{B} aborts in Round 3 if \mathcal{A} reveals R_j such that $\text{cm}_j = \text{H}_{\text{cm}}(m, \mathcal{S}, R_j)$ but \mathcal{A} never queried H_{cm} on (m, \mathcal{S}, R_j) . This requires \mathcal{A} to have guessed cm_j ahead of time, which occurs with negligible probability $1/p$.

\mathcal{B}_2 also aborts in Round 2 if \mathcal{A} had previously queried H_{sig} on $(\tilde{X}, m, \tilde{R})$. In that case, \mathcal{B} had returned a random $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$, so the reduction fails. However, this implies that \mathcal{A} guessed R_k before \mathcal{B} revealed it, which occurs with negligible probability $1/p$.

Next, we show that \mathcal{B} can extract the discrete logarithm of Y_0 from \mathcal{A} 's two valid forgeries. \mathcal{A} 's first forgery (m^*, R^*, z^*) satisfies $R^* \tilde{X}^{c^*} = g^{z^*}$, where $c^* = \text{H}_{\text{sig}}(\tilde{X}, m^*, R^*)$. Here, z^* does not suffice for \mathcal{B} to extract the discrete logarithm of $\tilde{X} = Y_0$ because it does not necessarily know the discrete logarithm of R^* . Thus, \mathcal{B} chooses $c' \leftarrow \mathbb{Z}_p$ and programs H_{sig} to output c' on input (\tilde{X}, m^*, R^*) . \mathcal{B} resets Q to the empty set, but the sets $\mathbf{Q}_{\text{cm}}, \mathbf{Q}_{\text{sig}}$ are kept for the second iteration of the adversary. \mathcal{B} then runs \mathcal{A} again on the same random coins.

After the second iteration, suppose \mathcal{A} terminates with (m', R', z') . If $(m', R') = (m^*, R^*)$ and \mathcal{A} 's forgeries both verify, then \mathcal{B} returns $y_0 = x = \frac{z^* - z'}{c^* - c'}$ such that $Y_0 = \tilde{X} = g^x$. If $(m', R') \neq (m^*, R^*)$ or \mathcal{A} 's forgery does not verify, then \mathcal{B} must abort. By the local forking lemma [6], this happens with probability less than $\frac{1}{q} (\text{Adv}_{\mathcal{A}}^{\text{adp-sec}}(\lambda, 2))^2$. If \mathcal{A} succeeds having not queried H_{cm} on (m, \mathcal{S}, R_j) , then \mathcal{B} aborts. This occurs with probability less than $\frac{qH}{p}$.

Thus,

$$\frac{1}{q} (\text{Adv}_{\mathcal{A}}^{\text{adp-sec}}(\lambda, 2))^2 \leq \Pr[\mathcal{B} \text{ extracts } y_0] + \text{negl}(\lambda)$$

If \mathcal{B} extracts y_0 , then we use this to extract a full AOMDL solution as follows.

Extracting an AOMDL Solution. The reduction \mathcal{B} must now extract the remaining y_1, \dots, y_t such that $Y_i = g^{y_i}$.

Assume without loss of generality that \mathcal{A} makes t corruptions over the two iterations. (If not, \mathcal{B} can corrupt the remaining number at the end, by querying its DL oracle until it reaches t secret keys.) Recall that \mathcal{B} set $\tilde{X}_i = Y_0 Y_1^i \dots Y_t^{i^t}$

and made t DL queries $g^{x_{i_1}}, \dots, g^{x_{i_t}}$:

$$\begin{aligned} g^{x_{i_1}} &= Y_0 Y_1^{i_1} \dots Y_t^{i_1} \\ &\vdots \\ g^{x_{i_t}} &= Y_0 Y_1^{i_t} \dots Y_t^{i_t} \end{aligned}$$

Recall also that $\tilde{X} = Y_0$. This forms the following system of linear equations:

$$\begin{aligned} x &= y_0 \\ x_{i_1} &= y_0 + i_1 y_1 \dots + i_1^t y_t \\ &\vdots \\ x_{i_t} &= y_0 + i_t y_1 \dots + i_t^t y_t \end{aligned}$$

Equivalently,

$$\begin{pmatrix} x \\ x_{i_1} \\ \vdots \\ x_{i_t} \end{pmatrix} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 1 & i_1 & \dots & i_1^t \\ \vdots & \vdots & \ddots & \vdots \\ 1 & i_t & \dots & i_t^t \end{pmatrix} \begin{pmatrix} y_0 \\ y_1 \\ \vdots \\ y_t \end{pmatrix}$$

\mathcal{B} knows all of the values on the left-hand side. The matrix

$$V = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 1 & i_1 & \dots & i_1^t \\ \vdots & \vdots & \ddots & \vdots \\ 1 & i_t & \dots & i_t^t \end{pmatrix}$$

is a Vandermonde matrix and is therefore invertible. Thus, \mathcal{B} can solve for (y_0, y_1, \dots, y_t) and win the $t+1$ -aomdl game.

9.3 Proof of Adaptive Security for up to t Corruptions

Proof. (of Theorem 3.) Let \mathcal{A} be an algebraic adversary attempting to break the adaptive unforgeability of Sparkle (Fig. 4.1). We construct a PPT reduction \mathcal{B} such that whenever \mathcal{A} outputs a valid forgery, \mathcal{B} breaks the $t+1$ -aomdl assumption (Fig. 3). More formally, we have

$$\text{Adv}_{\mathcal{A}}^{\text{adp-sec}}(\lambda, 1) \leq \text{Adv}_{\mathcal{B}}^{t+1\text{-aomdl}}(\lambda) + \text{negl}(\lambda)$$

Here, $\tau = 1$ to allow \mathcal{A} to corrupt $t/\tau = t$ parties.

The Reduction \mathcal{B} : We define the reduction \mathcal{B} playing game $\text{Game}_{\mathcal{B}}^{t+1\text{-aomdl}}(\lambda)$ as follows. \mathcal{B} is responsible for simulating key generation and oracle responses for queries to $\mathcal{O}^{\text{Sign}}$, $\mathcal{O}^{\text{Sign}'}$, $\mathcal{O}^{\text{Sign}''}$, $\mathcal{O}^{\text{Corrupt}}$, H_{cm} , and H_{sig} . Let Q_{cm} , Q_{sig} be the set of H_{cm} , H_{sig} queries and their responses, respectively. \mathcal{B} may program the random

oracles $H_{\text{cm}}, H_{\text{sig}}$. Let Q be the set of messages that have been queried to $\mathcal{O}^{\text{Sign}}$ as in game $\text{Game}_{\mathcal{A}}^{\text{adp-UF}}(\lambda, 1)$. \mathcal{B} initializes $Q_{\text{cm}}, Q_{\text{sig}}, Q$ to the empty set.

AOMDL Input. \mathcal{B} takes as input the group description $\mathcal{G} = (\mathbb{G}, p, g)$ and an AOMDL challenge of $t + 1$ values (Y_0, \dots, Y_t) . As in $\text{Game}_{\mathcal{B}}^{t+1\text{-aomdl}}(\lambda)$, \mathcal{B} has access to a discrete logarithm oracle \mathcal{O}^{dl} , which it may query up to t times. \mathcal{B} aims to output (y_0, \dots, y_t) such that $Y_i = g^{y_i}$ for all $0 \leq i \leq t$.

Initial Corruption. \mathcal{B} runs $\mathcal{A}()$. \mathcal{A} chooses the total number of potential signers n , the threshold $t + 1$, $1 \leq t + 1 \leq n$, and the initial set of corrupt parties $\text{cor} \leftarrow \{j\}, |\text{cor}| \leq t$. \mathcal{B} sets $\text{hon} \leftarrow [n] \setminus \text{cor}$ and must reveal the secret keys of the corrupt parties to \mathcal{A} , which \mathcal{B} does in the next step.

Simulating KeyGen. \mathcal{B} simulates the key generation algorithm (Fig. 5) using its AOMDL challenge (Y_0, \dots, Y_t) as follows. For all $1 \leq i \leq n$, \mathcal{B} sets the public key share as

$$\tilde{X}_i = Y_0 Y_1^i \dots Y_t^i$$

which is implicitly equal to $g^{f(i)}$. The joint public key is $\tilde{X} = g^{f(0)} = Y_0$ with corresponding secret key $x = y_0$. \mathcal{B} obtains the initial corrupt secret key shares by querying $x_j = f(j) \leftarrow \mathcal{O}^{\text{dl}}(\tilde{X}_j)$ (with representation $(1, j, \dots, j^t)$) for all $j \in \text{cor}$. \mathcal{B} runs $\mathcal{A}^{\mathcal{O}^{\text{Sign}}, \text{Sign}', \text{Sign}'', \text{Corrupt}}(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_j\}_{j \in \text{cor}})$.

Simulating Random Oracle Queries. \mathcal{B} handles \mathcal{A} 's random oracle queries throughout the protocol by lazy sampling, as follows.

H_{cm} : When \mathcal{A} queries H_{cm} on (m, \mathcal{S}, R) , \mathcal{B} checks whether $(m, \mathcal{S}, R, \text{cm}) \in Q_{\text{cm}}$ and, if so, returns cm . Else, \mathcal{B} samples $\text{cm} \leftarrow_{\mathcal{S}} \mathbb{Z}_p$, appends $(m, \mathcal{S}, R, \text{cm})$ to Q_{cm} , and returns cm .

H_{sig} : When \mathcal{A} queries H_{sig} on (X, m, R) , \mathcal{B} checks whether $(X, m, R, c) \in Q_{\text{sig}}$ and, if so, returns c . Else, \mathcal{B} samples $c \leftarrow_{\mathcal{S}} \mathbb{Z}_p$, appends (X, m, R, c) to Q_{sig} , and returns c .

Simulating Sparkle Signing. \mathcal{B} handles \mathcal{A} 's signing queries as follows.

Round 1 ($\mathcal{O}^{\text{Sign}}$): In the first round of signing for session ssid , each party P_i in the signing set \mathcal{S} sends a commitment cm_i . When \mathcal{A} queries $\mathcal{O}^{\text{Sign}}$ on (k, m, \mathcal{S}) for honest $k \in \text{hon}$, \mathcal{B} samples $\text{cm}_k \leftarrow_{\mathcal{S}} \mathbb{Z}_p$, appends $(m, \mathcal{S}, \cdot, \text{cm}_k)$ to Q_{cm} , sets $\text{st}_{k, \text{ssid}} \leftarrow (\text{cm}_k, \cdot, \cdot, m, \mathcal{S})$, and returns cm_k .

Round 2 ($\mathcal{O}^{\text{Sign}'}$): In the second round of signing for session ssid , each party P_i in the signing set \mathcal{S} takes as input the set of commitments $\{\text{cm}_i\}_{i \in \mathcal{S}}$ and reveals its nonce commitment R_i such that $\text{cm}_i = H_{\text{cm}}(m, \mathcal{S}, R_i)$. When \mathcal{A} queries $\mathcal{O}^{\text{Sign}'}$ on $(k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$ for $k \in \text{hon}$, then for all $j \in \mathcal{S} \cap \text{cor}$, \mathcal{B} looks up cm_j for a record $(m, \mathcal{S}, R_j, \text{cm}_j) \in Q_{\text{cm}}$.

If there exists $j' \in \mathcal{S}$ such that $(m', \mathcal{S}', \cdot, \text{cm}_{j'}) \in Q_{\text{cm}}$ and $(m', \mathcal{S}') \neq (m, \mathcal{S})$, or if there exists some $j' \in \mathcal{S}$ for which no record $(m, \mathcal{S}, \cdot, \text{cm}_{j'}) \in Q_{\text{cm}}$ exists, then

\mathcal{B} chooses $R_k \leftarrow_{\mathcal{S}} g^{r_k}$ randomly, updates $\text{st}_{k,\text{ssid}} \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S})$, programs $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$ and returns R_k .

Else if this is the first query in the signing session, \mathcal{B} samples $c \leftarrow_{\mathcal{S}} \mathbb{Z}_p$. For all $k \in \text{hon}$, \mathcal{B} samples $z_k \leftarrow_{\mathcal{S}} \mathbb{Z}_p$, computes $R_k \leftarrow g^{z_k} \tilde{X}_k^{-c\lambda_k}$ (λ_k is the Lagrange coefficient for the set \mathcal{S}), and programs $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$ (updating $(m, \mathcal{S}, R_k, \text{cm}_k) \in \text{Q}_{\text{cm}}$). Then \mathcal{B} computes $\tilde{R} = \prod_{i \in \mathcal{S}} R_i$ and programs $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$. (However, if \mathcal{A} has already queried H_{sig} on $(\tilde{X}, m, \tilde{R})$, then \mathcal{B} aborts.)

Finally, \mathcal{B} sets $\text{st}_{k,\text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$, and returns R_k to \mathcal{A} . If this is not the first query, \mathcal{B} looks up $\text{st}_{k,\text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$ and returns R_k .

Round 3 ($\mathcal{O}^{\text{Sign}''}$): In the third round of signing for session ssid , each party P_i in the signing set \mathcal{S} produces a partial signature on the message m . When \mathcal{A} queries $\mathcal{O}^{\text{Sign}''}$ on $(k, \{R_i\}_{i \in \mathcal{S}})$ for $k \in \text{hon}$, \mathcal{B} looks up $\text{st}_{k,\text{ssid}} \leftarrow (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$, checks whether $\text{cm}_i = \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$ for all $i \in \mathcal{S}$ and returns \perp if not. If $\text{cm}_i = \text{H}_{\text{cm}}(m, \mathcal{S}, R_i)$ but \mathcal{A} never queried H_{cm} on input (m, \mathcal{S}, R_i) , \mathcal{B} aborts. Else, \mathcal{B} sets $\text{st}_{k,\text{ssid}} \leftarrow ()$, and returns z_k .

Simulating Corruption Queries ($\mathcal{O}^{\text{Corrupt}}$): \mathcal{A} may at any time corrupt an honest party k by querying $\mathcal{O}^{\text{Corrupt}}(k)$. Upon receiving a corruption query, \mathcal{B} first checks that $k \in \text{hon}$, returning \perp if not. Otherwise, \mathcal{B} queries its DL oracle \mathcal{O}^{dl} on $\tilde{X}_k = g^{f(k)}$ (with representation $(1, k, \dots, k^t)$) to obtain the secret key $x_k = f(k)$. Then, for each $\text{st}_{k,\text{ssid}}$, \mathcal{B} does the following:

- If $\text{st}_{k,\text{ssid}} = (\text{cm}_k, \cdot, \cdot, m, \mathcal{S})$, then \mathcal{B} chooses $r_k \leftarrow_{\mathcal{S}} \mathbb{Z}_p$, sets $R_k \leftarrow g^{r_k}$, and programs $\text{cm}_k \leftarrow \text{H}_{\text{cm}}(m, \mathcal{S}, R_k)$. It then updates $\text{st}_{k,\text{ssid}} \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S})$.
- If $\text{st}_{k,\text{ssid}} = (\text{cm}_k, R_k, z_k, c\lambda_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$, then \mathcal{B} computes $r_k = z_k - c\lambda_k x_k$ (now that it knows x_k) and updates $\text{st}_{k,\text{ssid}} \leftarrow (\text{cm}_k, R_k, r_k, m, \mathcal{S}, \{\text{cm}_i\}_{i \in \mathcal{S}})$.
- If $\text{st}_{k,\text{ssid}} = ()$, then return $()$. This case occurs if signing session ssid has already been completed (i.e., a valid signature was issued).

Finally, \mathcal{B} sets $\text{hon} \leftarrow \text{hon} \setminus \{k\}$ and $\text{cor} \leftarrow \text{cor} \cup \{k\}$ and returns $x_k, \{\text{st}_{k,\ell}\}_{\ell \in [\text{ssid}]}$ to \mathcal{A} .

Output. At the end of the signing rounds, \mathcal{A} produces a forgery $(m^*, \sigma^*) = (m^*, (R^*, z^*))$. \mathcal{A} wins if $\text{Verify}(\tilde{X}, m^*, \sigma^*) = 1$, $m^* \notin Q$, and $|\text{cor}| \leq t$.

Extracting an AOMDL Solution. \mathcal{B} 's simulation of key generation and signing is perfect, and \mathcal{B} aborts with negligible probability. Indeed, \mathcal{B} aborts in Round 3 if \mathcal{A} reveals R_j such that $\text{cm}_j = \text{H}_{\text{cm}}(m, \mathcal{S}, R_j)$ but \mathcal{A} never queried H_{cm} on (m, \mathcal{S}, R_j) . This requires \mathcal{A} to have guessed cm_j ahead of time, which occurs with negligible probability $1/p$.

\mathcal{B}_2 also aborts in Round 2 if \mathcal{A} had previously queried H_{sig} on $(\tilde{X}, m, \tilde{R})$. In that case, \mathcal{B} had returned a random $c \leftarrow \text{H}_{\text{sig}}(\tilde{X}, m, \tilde{R})$, so the reduction fails. However, this implies that \mathcal{A} guessed R_k before \mathcal{B} revealed it, which occurs with negligible probability $1/p$.

It remains to show that \mathcal{B} can extract an AOMDL solution from \mathcal{A} 's output. Assume without loss of generality that \mathcal{A} makes t corruptions over the course of the protocol. (If not, \mathcal{B} can corrupt the remaining number at the end, by querying its DL oracle until it reaches t secret keys.) Recall that \mathcal{B} set $\tilde{X}_i = Y_0 Y_1^i \cdots Y_t^{i^t}$ and made t DL queries $g^{x_{i_1}}, \dots, g^{x_{i_t}}$:

$$\begin{aligned} g^{x_{i_1}} &= Y_0 Y_1^{i_1} \cdots Y_t^{i_1^t} \\ &\vdots \\ g^{x_{i_t}} &= Y_0 Y_1^{i_t} \cdots Y_t^{i_t^t} \end{aligned}$$

Recall also that $\tilde{X} = Y_0$. This forms the following system of linear equations:

$$\begin{aligned} x &= y_0 \\ x_{i_1} &= y_0 + i_1 y_1 \cdots + i_1^t y_t \\ &\vdots \\ x_{i_t} &= y_0 + i_t y_1 \cdots + i_t^t y_t \end{aligned}$$

Equivalently,

$$\begin{pmatrix} x \\ x_{i_1} \\ \vdots \\ x_{i_t} \end{pmatrix} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 1 & i_1 & \cdots & i_1^t \\ \vdots & \vdots & \ddots & \vdots \\ 1 & i_t & \cdots & i_t^t \end{pmatrix} \begin{pmatrix} y_0 \\ y_1 \\ \vdots \\ y_t \end{pmatrix} \quad (2)$$

\mathcal{B} knows all of the values $\{x_j\}_{j \in \text{cor}} = \{x_{i_1}, \dots, x_{i_t}\}$ on the left-hand side, but not x . However, \mathcal{B} can compute x as follows.

Extracting the Discrete Logarithm of $\tilde{X} = Y_0$. \mathcal{A} 's forgery satisfies:

$$\tilde{R}^* = g^{z^*} \tilde{X}^{-c^*} \quad (3)$$

where $c^* = \text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$. On the other hand, when \mathcal{A} made its query $\text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$, it provided a representation of \tilde{R}^* in basis $(g, \tilde{X}, \tilde{X}_1, \dots, \tilde{X}_n, \{R_{i,1}, \dots, R_{i,n}\}_{i \in [q_S]})$:

$$\tilde{R}^* = g^{\gamma^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \cdots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{q_S} R_{i,1}^{\rho_{i,1}^*} \cdots R_{i,n}^{\rho_{i,n}^*}$$

where q_S is the number of signing queries that \mathcal{A} makes. Each $R_{i,k}$ satisfies $R_{i,k} = g^{z_{i,k}} \tilde{X}_k^{-c_i \lambda_{i,k}}$, where $c_i = \text{H}_{\text{sig}}(\tilde{X}, m_i, \tilde{R}_i)$. Thus,

$$\tilde{R}^* = g^{\gamma^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \cdots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{q_S} (g^{z_{i,1}} \tilde{X}_1^{-c_i \lambda_{i,1}})^{\rho_{i,1}^*} \cdots (g^{z_{i,n}} \tilde{X}_n^{-c_i \lambda_{i,n}})^{\rho_{i,n}^*}$$

Equating this with Eq. (3), we have:

$$g^{z^*} \tilde{X}^{-c^*} = g^{\gamma^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{qs} (g^{z_{i,1}} \tilde{X}_1^{-c_i \lambda_{i,1}})^{\rho_{i,1}^*} \dots (g^{z_{i,n}} \tilde{X}_n^{-c_i \lambda_{i,n}})^{\rho_{i,n}^*}$$

Rearranging, we have:

$$g^{z^*} g^{-\gamma^*} \prod_{i=1}^{qs} g^{-z_{i,1} \rho_{i,1}^*} \dots g^{-z_{i,n} \rho_{i,n}^*} = \tilde{X}^{c^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{qs} (\tilde{X}_1^{-c_i \lambda_{i,1}})^{\rho_{i,1}^*} \dots (\tilde{X}_n^{-c_i \lambda_{i,n}})^{\rho_{i,n}^*} \quad (4)$$

Let $\eta^* = z^* - \gamma^* - \sum_{i=1}^{qs} (z_{i,1} \rho_{i,1}^* + \dots + z_{i,n} \rho_{i,n}^*)$ and $\zeta_k^* = \xi_k^* - \sum_{i=1}^{qs} c_i \lambda_{i,k} \rho_{i,k}^*$ for all $1 \leq k \leq n$. Then Eq. (4) can be rewritten as:

$$g^{\eta^*} = \tilde{X}^{c^* + \xi^*} \tilde{X}_1^{\zeta_1^*} \dots \tilde{X}_n^{\zeta_n^*} \quad (5)$$

Recall that $\tilde{X}_i = \tilde{X} Y_1^i Y_2^{i^2} \dots Y_t^{i^t}$ for all $i \in [n]$ and that $Y_i = \tilde{X}^{L'_{0,i}} \prod_{j \in \text{cor}} g^{x_j L'_{j,i}}$ for all $i \in [t]$. Thus,

$$\tilde{X}_k = \tilde{X} \prod_{i=1}^t \left(\tilde{X}^{L'_{0,i}} \prod_{j \in \text{cor}} g^{x_j L'_{j,i}} \right)^{k^i}$$

Let $\mu_k^* = 1 + \sum_{i=1}^t L'_{0,i} k^i$ and $\nu_k^* = \sum_{i=1}^t \left(\sum_{j \in \text{cor}} x_j L'_{j,i} \right) k^i$. Then \tilde{X}_k can be rewritten as:

$$\tilde{X}_k = \tilde{X}^{\mu_k^*} g^{\nu_k^*}$$

and Eq. (5) can be rewritten as:

$$g^{\eta^*} = \tilde{X}^{c^* + \xi^*} \prod_{k=1}^n \tilde{X}^{\mu_k^*} g^{\nu_k^*}$$

Rearranging, we have:

$$g^{\eta^* - \sum_{k=1}^n \nu_k^*} = \tilde{X}^{c^* + \xi^* + \sum_{k=1}^n \mu_k^*}$$

and

$$\dot{x} = \frac{\eta^* - \sum_{k=1}^n \nu_k^*}{c^* + \xi^* + \sum_{k=1}^n \mu_k^*}$$

\mathcal{A} fixed \tilde{R}^* and thus $\eta^*, \{\nu_i^*\}_{i \in [n]}, \xi^*, \{\mu_i^*\}_{i \in [n]}$ as it queried $\mathbf{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$ to receive random c^* . Thus, the denominator is nonzero with overwhelming probability and \mathcal{B} can solve for x .

The matrix

$$V = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 1 & i_1 & \cdots & i_1^t \\ \vdots & \vdots & \ddots & \vdots \\ 1 & i_t & \cdots & i_t^t \end{pmatrix}$$

in Equation 2 is a Vandermonde matrix and is therefore invertible. Thus, \mathcal{B} can solve for (y_0, y_1, \dots, y_t) and win the $t + 1$ -aomdl game.

Acknowledgements. Elizabeth Crites was supported by Input Output through their funding of the Blockchain Technology Lab at the University of Edinburgh.

References

- [1] M. Abdalla, M. Barbosa, J. Katz, J. Loss, and J. Xu. “Algebraic Adversaries in the Universal Composability Framework”. In: *Advances in Cryptology - ASIACRYPT 2021 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part III*. Ed. by M. Tibouchi and H. Wang. Vol. 13092. Lecture Notes in Computer Science. Springer, 2021, pp. 311–341.
- [2] J. F. Almansa, I. Damgård, and J. B. Nielsen. “Simplified Threshold RSA with Adaptive and Proactive Security”. In: *EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006*. Ed. by S. Vaudenay. Vol. 4004. LNCS. Springer, 2006, pp. 593–611.
- [3] R. Bacho and J. Loss. “On the Adaptive Security of the Threshold BLS Signature Scheme”. In: *IACR Cryptol. ePrint Arch.* (2022), p. 534. URL: <https://eprint.iacr.org/2022/534>.
- [4] B. Bauer, G. Fuchsbauer, and A. Plouviez. “The One-More Discrete Logarithm Assumption in the Generic Group Model”. In: *ASIACRYPT 2021, Singapore, December 6-10, 2021*. Ed. by M. Tibouchi and H. Wang. Vol. 13093. LNCS. Springer, 2021, pp. 587–617.
- [5] M. Bellare, E. Crites, C. Komlo, M. Maller, S. Tessaro, and C. Zhu. *Better than advertised security for non-interactive threshold signatures*. CRYPTO 2022. To appear. 2022.
- [6] M. Bellare, W. Dai, and L. Li. “The Local Forking Lemma and Its Application to Deterministic Encryption”. In: *ASIACRYPT 2019, Kobe, Japan, December 8-12, 2019*. Ed. by S. D. Galbraith and S. Moriai. Vol. 11923. LNCS. Springer, 2019, pp. 607–636.
- [7] M. Bellare, C. Namprempre, D. Pointcheval, and M. Semanko. “The One-More-RSA-Inversion Problems and the Security of Chaum’s Blind Signature Scheme”. In: *J. Cryptol.* 16.3 (2003), pp. 185–215.
- [8] F. Benhamouda, T. Lepoint, J. Loss, M. Orrù, and M. Raykova. “On the (in)security of ROS”. In: *EUROCRYPT 2021, Zagreb, Croatia, October 17-21, 2021*. Ed. by A. Canteaut and F. Standaert. Vol. 12696. LNCS. Springer, 2021, pp. 33–53.

- [9] A. Boldyreva. “Threshold Signatures, Multisignatures and Blind Signatures Based on the Gap-Diffie-Hellman-Group Signature Scheme”. In: *PKC 2003, Miami, FL, USA, January 6-8, 2003*. Ed. by Y. Desmedt. Vol. 2567. LNCS. Springer, 2003, pp. 31–46.
- [10] D. Boneh, M. Drijvers, and G. Neven. “Compact Multi-signatures for Smaller Blockchains”. In: *ASIACRYPT 2018, Brisbane, QLD, Australia, December 2-6, 2018*. Ed. by T. Peyrin and S. D. Galbraith. Vol. 11273. LNCS. Springer, 2018, pp. 435–464.
- [11] L. Brandão and M. Davidson. *Notes on Threshold EdDSA/Schnorr Signatures*. <https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8214B.ipd.pdf>. 2022.
- [12] L. Brandão and R. Peralta. *NIST First Call for Multi-Party Threshold Schemes*. <https://nvlpubs.nist.gov/nistpubs/ir/2023/NIST.IR.8214C.ipd.pdf>. 2023.
- [13] R. Canetti. “Universally Composable Security: A New Paradigm for Cryptographic Protocols”. In: *FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*. IEEE Computer Society, 2001, pp. 136–145.
- [14] R. Canetti, U. Feige, O. Goldreich, and M. Naor. “Adaptively Secure Multi-Party Computation”. In: *STOC '96, Philadelphia, Pennsylvania, USA, May 22-24, 1996*. Ed. by G. L. Miller. ACM, 1996, pp. 639–648.
- [15] R. Canetti, R. Gennaro, S. Goldfeder, N. Makriyannis, and U. Peled. “UC Non-Interactive, Proactive, Threshold ECDSA with Identifiable Aborts”. In: *IACR Cryptol. ePrint Arch.* (2021), p. 60. URL: <https://eprint.iacr.org/2021/060>.
- [16] R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. “Adaptive Security for Threshold Cryptosystems”. In: *CRYPTO '99, Santa Barbara, California, USA, August 15-19, 1999*. Ed. by M. J. Wiener. Vol. 1666. LNCS. Springer, 1999, pp. 98–115.
- [17] D. Connolly, C. Komlo, I. Goldberg, and C. Wood. *Two-Round Threshold Schnorr Signatures with FROST*. 2022. URL: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-frost/>.
- [18] M. Drijvers et al. “On the Security of Two-Round Multi-Signatures”. In: *SP 2019, San Francisco, CA, USA, May 19-23, 2019*. IEEE, 2019, pp. 1084–1101.
- [19] B. Edgington. *Upgrading Ethereum*. 2023. URL: https://eth2book.info/bellatrix/part2/building_blocks/randomness/.
- [20] A. Fiat and A. Shamir. “How to Prove Yourself: Practical Solutions to Identification and Signature Problems”. In: *CRYPTO 1986, Santa Barbara, California, USA, 1986*. Ed. by A. M. Odlyzko. Vol. 263. LNCS. Springer, 1986, pp. 186–194.
- [21] M. Fischlin. “Communication-Efficient Non-interactive Proofs of Knowledge with Online Extractors”. In: *CRYPTO 2005, Santa Barbara, California, USA, August 14-18, 2005*. Ed. by V. Shoup. Vol. 3621. LNCS. Springer, 2005, pp. 152–168.
- [22] G. Fuchsbauer, E. Kiltz, and J. Loss. “The Algebraic Group Model and its Applications”. In: *CRYPTO 2018, Santa Barbara, CA, USA, August 19-23, 2018*. Ed. by H. Shacham and A. Boldyreva. Vol. 10992. LNCS. Springer, 2018, pp. 33–62.
- [23] R. Gennaro and S. Goldfeder. “Fast Multiparty Threshold ECDSA with Fast Trustless Setup”. In: *CCS 2018, Toronto, ON, Canada, October 15-19, 2018*. Ed. by D. Lie, M. Mannan, M. Backes, and X. Wang. ACM, 2018, pp. 1179–1194.
- [24] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. “Robust Threshold DSS Signatures”. In: *Inf. Comput.* 164.1 (2001), pp. 54–84.

- [25] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. “Secure Applications of Pedersen’s Distributed Key Generation Protocol”. In: *CT-RSA 2003, San Francisco, CA, USA, April 13-17, 2003*. Ed. by M. Joye. Vol. 2612. LNCS. Springer, 2003, pp. 373–390.
- [26] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. “Secure Distributed Key Generation for Discrete-Log Based Cryptosystems”. In: *J. Cryptol.* 20.1 (2007), pp. 51–83.
- [27] R. Gennaro, T. Rabin, S. Jarecki, and H. Krawczyk. “Robust and Efficient Sharing of RSA Functions”. In: *J. Cryptol.* 20.3 (2007), p. 393.
- [28] C. Gentry and D. Wichs. “Separating succinct non-interactive arguments from all falsifiable assumptions”. In: *STOC 2011, San Jose, CA, USA, 6-8 June 2011*. Ed. by L. Fortnow and S. P. Vadhan. ACM, 2011, pp. 99–108.
- [29] S. Goldwasser, S. Micali, and R. L. Rivest. “A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks”. In: *SIAM J. Comput.* 17.2 (1988), pp. 281–308.
- [30] S. Jarecki and A. Lysyanskaya. “Adaptively Secure Threshold Cryptography: Introducing Concurrency, Removing Erasures”. In: *EUROCRYPT 2000, Bruges, Belgium, May 14-18, 2000*. Ed. by B. Preneel. Vol. 1807. LNCS. Springer, 2000, pp. 221–242.
- [31] N. Koblitz and A. Menezes. “Another look at non-standard discrete log and Diffie-Hellman problems”. In: *J. Math. Cryptol.* 2.4 (2008), pp. 311–326.
- [32] C. Komlo and I. Goldberg. “FROST: Flexible Round-Optimized Schnorr Threshold Signatures”. In: *SAC 2020, Halifax, NS, Canada (Virtual Event), October 21-23, 2020*. Ed. by O. Dunkelman, M. J. J. Jr., and C. O’Flynn. Vol. 12804. LNCS. Springer, 2020, pp. 34–65.
- [33] B. Libert, M. Joye, and M. Yung. “Born and raised distributively: Fully distributed non-interactive adaptively-secure threshold signatures with short shares”. In: *Theoretical Computer Science* 645 (2016), pp. 1–24.
- [34] Y. Lindell. “Simple Three-Round Multiparty Schnorr Signing with Full Simulatability”. In: *IACR Cryptol. ePrint Arch.* (2022), p. 374. URL: <https://eprint.iacr.org/2022/374>.
- [35] A. Lysyanskaya and C. Peikert. “Adaptive Security in the Threshold Setting: From Cryptosystems to Signature Schemes”. In: *ASIACRYPT 2001, Gold Coast, Australia, December 9-13, 2001*. Ed. by C. Boyd. Vol. 2248. LNCS. Springer, 2001, pp. 331–350.
- [36] N. Makriyannis. *On the Classic Protocol for MPC Schnorr Signatures*. Cryptology ePrint Archive, Paper 2022/1332. <https://eprint.iacr.org/2022/1332>. 2022. URL: <https://eprint.iacr.org/2022/1332>.
- [37] G. Maxwell, A. Poelstra, Y. Seurin, and P. Wuille. “Simple Schnorr multi-signatures with applications to Bitcoin”. In: *Des. Codes Cryptogr.* 87.9 (2019), pp. 2139–2164.
- [38] J. Nick, T. Ruffing, and Y. Seurin. “MuSig2: Simple Two-Round Schnorr Multi-signatures”. In: *CRYPTO 2021, Virtual Event, August 16-20, 2021*. Ed. by T. Malkin and C. Peikert. Vol. 12825. LNCS. Springer, 2021, pp. 189–221.
- [39] A. Nicolosi, M. N. Krohn, Y. Dodis, and D. Mazières. “Proactive Two-Party Signatures for User Authentication”. In: *Proceedings of the Network and Distributed System Security Symposium, NDSS 2003, San Diego, California, USA*. The Internet Society, 2003. URL: <https://www.ndss-symposium.org/ndss2003/proactive-two-party-signatures-user-authentication/>.

- [40] D. Pointcheval and J. Stern. “Security Arguments for Digital Signatures and Blind Signatures”. In: *J. Cryptol.* 13.3 (2000), pp. 361–396.
- [41] C. Schnorr. “Efficient Signature Generation by Smart Cards”. In: *J. Cryptol.* 4.3 (1991), pp. 161–174.
- [42] A. Shamir. “How to Share a Secret”. In: *Commun. ACM* 22.11 (1979), pp. 612–613.
- [43] D. R. Stinson and R. Strobl. “Provably Secure Distributed Schnorr Signatures and a (t, n) Threshold Scheme for Implicit Certificates”. In: *ACISP 2001, Sydney, Australia, July 11-13, 2001*. Ed. by V. Varadharajan and Y. Mu. Vol. 2119. LNCS. Springer, 2001, pp. 417–434.

A Proof of Static Security in the AGM

We additionally provide a proof of static security for **Sparkle** under the discrete logarithm (DL) assumption in the algebraic group model (AGM) and random oracle model (ROM) without any tightness loss. This security reduction serves as a warm-up to the proof of full adaptive security in Section 8.

Theorem 4. *Sparkle is statically secure under DL in the AGM+ROM.*

Proof. Let \mathcal{A} be an algebraic adversary attempting to break the static unforgeability of **Sparkle** (Fig. 4.1) We construct a PPT reduction \mathcal{B} against the DL assumption (Fig. 3) that uses \mathcal{A} as a subroutine such that

$$\text{Adv}_{\mathcal{A}}^{\text{st-sec}}(\lambda, 1) \leq \text{Adv}_{\mathcal{B}}^{\text{dl}}(\lambda) + \text{negl}(\lambda)$$

Here, $\tau = 1$ to allow \mathcal{A} to corrupt $t/\tau = t$ parties.

The Reduction \mathcal{B} : We define the reduction \mathcal{B} playing game $\text{Game}_{\mathcal{B}}^{\text{dl}}(\lambda)$ as follows. \mathcal{B} is responsible for simulating key generation and oracle responses for queries to $\mathcal{O}^{\text{Sign}}, \mathcal{O}^{\text{Sign}'}, \mathcal{O}^{\text{Sign}''}, \text{H}_{\text{cm}}$, and H_{sig} . Let $\text{Q}_{\text{cm}}, \text{Q}_{\text{sig}}$ be the set of $\text{H}_{\text{cm}}, \text{H}_{\text{sig}}$ queries and their responses, respectively. \mathcal{B} may program the random oracles $\text{H}_{\text{cm}}, \text{H}_{\text{sig}}$. Let Q be the set of messages that have been queried in $\mathcal{O}^{\text{Sign}}$ as in $\text{Game}_{\mathcal{A}}^{\text{UF}}(\lambda, 1)$. \mathcal{B} initializes $\text{Q}_{\text{cm}}, \text{Q}_{\text{sig}}, Q$ to the empty set.

DL Input. \mathcal{B} takes as input the group description $\mathcal{G} = (\mathbb{G}, p, g)$ and a DL challenge \dot{X} . \mathcal{B} aims to output \dot{x} such that $\dot{X} = g^{\dot{x}}$.

Static Corruption. \mathcal{B} runs $\mathcal{A}()$. \mathcal{A} chooses the total number of potential signers n , the threshold $t + 1, 1 \leq t + 1 \leq n$, and the set of corrupt parties $\text{cor} \leftarrow \{j\}, |\text{cor}| \leq t$, which are fixed for the rest of the protocol. \mathcal{B} sets $\text{hon} \leftarrow [n] \setminus \text{cor}$ and must reveal the secret keys of the corrupt parties to \mathcal{A} , which \mathcal{B} does in the next step.

Simulating KeyGen. \mathcal{B} simulates the key generation algorithm (Fig. 5) using its DL challenge \dot{X} as follows.

1. \mathcal{B} sets the joint public key $\tilde{X} \leftarrow \dot{X}$.

2. \mathcal{B} simulates a Shamir secret sharing of the discrete logarithm of \tilde{X} by performing the following steps. (See Section 3 for notation.) Assume without loss of generality that $|\text{cor}| = t$.

- (a) \mathcal{B} samples t random values $x_j \leftarrow^* \mathbb{Z}_p$ for $j \in \text{cor}$.
- (b) Let f be the polynomial whose constant term is the challenge $f(0) = \dot{x}$ and for which $f(j) = x_j$ for all $j \in \text{cor}$. \mathcal{B} computes the $t + 1$ Lagrange polynomials $\{L'_0(Z), \{L'_j(Z)\}_{j \in \text{cor}}\}$ relating to the set (of x -coordinates) $0 \cup \text{cor}$.
- (c) For all $1 \leq i \leq t$, \mathcal{B} computes

$$Y_i = \tilde{X}^{L'_{0,i}} \prod_{j \in \text{cor}} g^{x_j L'_{j,i}} \quad (6)$$

where $L'_{j,i}$ is the i^{th} coefficient of $L'_j(Z) = L'_{j,0} + L'_{j,1}Z + \dots + L'_{j,t}Z^t$.

- (d) For $1 \leq i \leq n$, \mathcal{B} computes

$$\tilde{X}_i = \tilde{X} Y_1^i Y_2^{i^2} \dots Y_t^{i^t} \quad (7)$$

which is implicitly equal to $g^{f(i)}$.

The joint public key is $\tilde{X} = g^{f(0)} = \dot{X}$ with corresponding secret key \dot{x} . \mathcal{B} runs $\mathcal{A}^{\mathcal{O}^{\text{Sign}, \text{Sign}', \text{Sign}''}}(\tilde{X}, \{\tilde{X}_i\}_{i \in [n]}, \{x_j\}_{j \in \text{cor}})$.

Simulating Random Oracle and Signing Queries. \mathcal{B} simulates random oracle and signing queries as in the static proof under DL+ROM (Section 9.1).

Output. At the end of the signing rounds, \mathcal{A} produces a forgery $(m^*, \sigma^*) = (m^*, (R^*, z^*))$. \mathcal{A} wins if $\text{Verify}(\tilde{X}, m^*, \sigma^*) = 1$ and $m^* \notin Q$.

Extracting the Discrete Logarithm of \dot{X} . \mathcal{A} 's forgery satisfies:

$$\tilde{R}^* = g^{z^*} \tilde{X}^{-c^*} \quad (8)$$

where $c^* = \text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$. On the other hand, when \mathcal{A} made its query $\text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$, it provided a representation of \tilde{R}^* in basis $(g, \tilde{X}, \tilde{X}_1, \dots, \tilde{X}_n, \{R_{i,1}, \dots, R_{i,n}\}_{i \in [q_S]})$:

$$\tilde{R}^* = g^{\gamma^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{q_S} R_{i,1}^{\rho_{i,1}^*} \dots R_{i,n}^{\rho_{i,n}^*}$$

where q_S is the number of signing queries that \mathcal{A} makes. Each $R_{i,k}$ satisfies $R_{i,k} = g^{z_{i,k}} \tilde{X}_k^{-c_i \lambda_{i,k}}$, where $c_i = \text{H}_{\text{sig}}(\tilde{X}, m_i, \tilde{R}_i)$. Thus,

$$\tilde{R}^* = g^{\gamma^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{q_S} (g^{z_{i,1}} \tilde{X}_1^{-c_i \lambda_{i,1}})^{\rho_{i,1}^*} \dots (g^{z_{i,n}} \tilde{X}_n^{-c_i \lambda_{i,n}})^{\rho_{i,n}^*}$$

Equating this with Eq. (8), we have:

$$g^{z^*} \tilde{X}^{-c^*} = g^{\gamma^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{qs} (g^{z_{i,1}} \tilde{X}_1^{-c_i \lambda_{i,1}})^{\rho_{i,1}^*} \dots (g^{z_{i,n}} \tilde{X}_n^{-c_i \lambda_{i,n}})^{\rho_{i,n}^*}$$

Rearranging, we have:

$$g^{z^*} g^{-\gamma^*} \prod_{i=1}^{qs} g^{-z_{i,1} \rho_{i,1}^*} \dots g^{-z_{i,n} \rho_{i,n}^*} = \tilde{X}^{c^*} \tilde{X}^{\xi^*} \tilde{X}_1^{\xi_1^*} \dots \tilde{X}_n^{\xi_n^*} \prod_{i=1}^{qs} (\tilde{X}_1^{-c_i \lambda_{i,1}})^{\rho_{i,1}^*} \dots (\tilde{X}_n^{-c_i \lambda_{i,n}})^{\rho_{i,n}^*}$$

Let $\eta^* = z^* - \gamma^* - \sum_{i=1}^{qs} (z_{i,1} \rho_{i,1}^* + \dots + z_{i,n} \rho_{i,n}^*)$ and $\zeta_k^* = \xi_k^* - \sum_{i=1}^{qs} c_i \lambda_{i,k} \rho_{i,k}^*$ for all $k \in [n]$. Then this can be rewritten as:

$$g^{\eta^*} = \tilde{X}^{c^* + \xi^*} \tilde{X}_1^{\zeta_1^*} \dots \tilde{X}_n^{\zeta_n^*} \quad (9)$$

Recall that $\tilde{X}_i = \tilde{X} Y_1^i Y_2^{i^2} \dots Y_t^i$ for all $i \in [n]$ (Eq. (7)) and that $Y_i = \tilde{X}^{L'_{0,i}} \prod_{j \in \text{cor}} g^{x_j L'_{j,i}}$ for all $i \in [t]$ (Eq. (6)). Thus,

$$\tilde{X}_k = \tilde{X} \prod_{i=1}^t \left(\tilde{X}^{L'_{0,i}} \prod_{j \in \text{cor}} g^{x_j L'_{j,i}} \right)^{k^i}$$

Let $\mu_k^* = 1 + \sum_{i=1}^t L'_{0,i} k^i$ and $\nu_k^* = \sum_{i=1}^t \left(\sum_{j \in \text{cor}} x_j L'_{j,i} \right) k^i$. Then \tilde{X}_k can be rewritten as:

$$\tilde{X}_k = \tilde{X}^{\mu_k^*} g^{\nu_k^*}$$

and Eq. (9) can be rewritten as:

$$g^{\eta^*} = \tilde{X}^{c^* + \xi^*} \prod_{k=1}^n \tilde{X}_k^{\mu_k^*} g^{\nu_k^*}$$

Rearranging, we have:

$$g^{\eta^* - \sum_{k=1}^n \nu_k^*} = \tilde{X}^{c^* + \xi^* + \sum_{k=1}^n \mu_k^*}$$

and

$$\dot{x} = \frac{\eta^* - \sum_{k=1}^n \nu_k^*}{c^* + \xi^* + \sum_{k=1}^n \mu_k^*}$$

A fixed \tilde{R}^* and thus $\eta^*, \{\nu_k^*\}_{k \in [n]}, \xi^*, \{\mu_k^*\}_{k \in [n]}$ as it queried $\text{H}_{\text{sig}}(\tilde{X}, m^*, \tilde{R}^*)$ to receive random c^* . Thus, the denominator is nonzero with overwhelming probability and \mathcal{B} can solve for \dot{x} to win the DL game.