

From One-Time to Two-Round Reusable Multi-Signatures without Nested Forking

Lior Rotem*

Gil Segev[†]

Eylon Yogev[‡]

Abstract

Multi-signature schemes are gaining significant interest due to their blockchain applications. Of particular interest are two-round schemes in the plain public-key model that offer key aggregation, and whose security is based on the hardness of the DLOG problem. Unfortunately, despite substantial recent progress, the security proofs of the proposed schemes provide rather insufficient concrete guarantees (especially for 256-bit groups). This frustrating situation has so far been approached either by relying on the security of seemingly-stronger assumptions or by considering restricted classes of attackers (e.g., algebraic attackers, which are assumed to provide an algebraic justification of each group element that they produce).

We present a complementing approach by constructing multi-signature schemes that satisfy two relaxed notions of security, whose applicability nevertheless ranges from serving as drop-in replacements to enabling expressive smart contract validation procedures. Our first notion, *one-time unforgeability*, extends the analogous single-signer notion by considering attackers that obtain a single signature for some message and set of signers of their choice. We construct a *non-interactive* one-time scheme based on any ring-homomorphic one-way function, admitting efficient instantiations based on the DLOG and RSA assumptions. Aggregated verification keys and signatures consist of two group elements and a single group element, respectively, and our security proof consists of a *single* application of the forking lemma (thus avoiding the substantial security loss exhibited by the proposed two-round schemes). Additionally, we demonstrate that our scheme naturally extends to a t -time scheme, where aggregated verification keys consist of $t + 1$ group elements, while aggregated signatures still consist of a single group element.

Our second notion, *single-set unforgeability*, considers attackers that obtain *any polynomial number* of signatures but are restricted to a single set of signers of their choice. We transform any non-interactive one-time scheme into a two-round single-set scheme via a novel forking-free construction that extends the seminal Naor-Yung tree-based approach to the multi-signer setting. Aggregated verification keys are essentially identical to those of the underlying one-time scheme, and the length of aggregated signatures is determined by that of the underlying scheme while scaling linearly with the length of messages (noting that long messages can always be hashed using a collision-resistant function). Instantiated with our one-time scheme, we obtain

*Computer Science Department, Stanford University, 353 Jane Stanford Way, Stanford, CA 94305, USA. Email: lrotem@cs.stanford.edu. Supported by a research grant from Protocol Labs.

[†]School of Computer Science and Engineering, Hebrew University of Jerusalem, Jerusalem 91904, Israel. Email: segev@cs.huji.ac.il. Supported by the Israel Science Foundation (Grant No. 1336/22) and by the European Union (ERC, FTRC, 101043243). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

[‡]Department of Computer Science, Bar-Ilan University, Israel. Email: eylon.yogev@biu.ac.il. Supported by the Israel Science Foundation (Grant No. 2302/22), European Research Union (ERC, CRYPTOPROOF, 101164375), and by an Alon Young Faculty Fellowship. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

aggregated verification keys and signatures whose lengths are completely independent of the number of signers.

Contents

1	Introduction	1
1.1	Our Contributions	2
1.2	Overview of Our Approach	4
1.3	Related Work	7
2	Preliminaries	8
2.1	Ring-Homomorphic One-Way Functions	9
2.2	The Forking Lemma	9
2.3	Additional Cryptographic Primitives	10
3	One-Time and Single-Set Security for Multi-Signature Schemes	13
3.1	One-Time Unforgeability	14
3.2	Single-Set Unforgeability	15
4	One-Time Multi-Signatures via Ring-Homomorphic One-Way Functions	16
5	From One-Time to Single-Set Multi-Signatures	22
	References	34
A	Extension: A t-Time Multi-Signature Scheme	40
B	From One-Time Multi-Signatures to Collision-Resistant Hashing	42

1 Introduction

A multi-signature scheme [IN83, BN06] enables any set of signers, within a large and potentially permissionless system, to jointly produce a compact signature on a given message. Research on the design and analysis of multi-signature schemes has recently gained significant renewed interest, as such schemes were found particularly suitable for blockchain applications. These range from drop-in replacements for standard signatures (e.g., [BDN18, MPS⁺19]), to smart contracts with expressive multi-owner validation procedures (e.g., [BWG⁺21, Arg22, Sta23]).

Two breakthroughs: Plain PK model and key aggregation. The high-level of suitability exhibited by multi-signatures to blockchain applications follows mostly due to two major breakthroughs. First, Bellare and Neven [BN06] showed that multi-signature schemes can provide security in the plain public-key model, capturing a realistic and permissionless environment, while not compromising on practicality. Specifically, in this model, each signer locally produces their signing and verification keys, without engaging in an interactive key-generation process with other signers or with a registration authority, and without augmenting verification keys with proofs of knowledge that need to be individually verified by all other signers. Second, Boneh, Drijvers and Neven [BDN18] and Maxwell, Poelstra, Seurin and Wuille [MPS⁺19] showed that multi-signature schemes can support non-interactive aggregation of verification keys (in the plain public-key model). Therefore, once an aggregated verification key has been verified to correspond to a particular set of signers, any dependence on the number of signers during all future signature verifications may be completely eliminated. This essentially turns the verification of multi-signatures as practical as (and even fully compatible with) that of standard signatures.

Following up on earlier constructions (e.g., [OO91, LHL94, MOR01, Bol03, LOS⁺06, BGO⁺07, RY07] and the references therein), this significant progress has led to a host of recent multi-signature schemes [DEF⁺19, NRS⁺20, AB21, BD21, NRS21, BTT22, DOT⁺22, FSZ22, LK23, PW23, TZ23]. Of particular interest in the blockchain setting are two-round schemes in the plain public-key model whose security is based on the hardness of the discrete logarithm (DLOG) problem in prime-order groups, as such schemes may admit practical implementations over standard elliptic curves, such as Secp256k1 or Curve25519.

Concrete security guarantees in 256-bit groups? As observed by Bellare and Dai [BD21] (and most recently also by Pan and Wagner [PW23]), despite the substantial recent efforts, the existing proofs that base the security of the proposed two-round schemes in the plain public-key model on the hardness of the DLOG problem provide rather insufficient concrete guarantees in 256-bit groups. As common for DLOG-based signatures, these proofs of security are based on the classic forking lemma [PS00, BN06], but instead of relying on a single application of the lemma, they rely on two nested applications – leading to a substantial loss in the provable concrete security.

At a high level, under the widely-accepted assumption that the success probability of any t -time algorithm in solving the DLOG problem in a group of order p is at most t^2/p [Sho97], proofs of security that rely on two nested applications of the forking lemma seem limited to bounding the success probability of t -time attackers with roughly $(t^2/p)^{1/4}$. For a 256-bit prime p , such a bound falls short of providing sufficient concrete security guarantees, especially when compared to the $(t^2/p)^{1/2}$ bound resulting from a single application of the forking lemma (as is the case, for example, with Schnorr signatures, and more generally with signature schemes obtained from identification protocols via the Fiat-Shamir transform [FS86, Sch91, AAB⁺02, KMP16]).¹

¹For simplicity, in the above discussion we did not include additional factors that depend on the number of random-

This frustrating situation has so far been approached for DLOG-based schemes in the plain public-key model via two relaxations: Either relying on the security of a recently-introduced stronger variant of the DLOG assumption (the interactive XIDL assumption introduced by Bellare and Dai [BD21]), or by proving security with respect to restricted classes of attackers [AB21, BD21, NRS21, LK23] (most notably, algebraic attackers, which are assumed to provide an algebraic justification of each group element that they produce [FKL18, AHK20, BFL20, FPS20, MTT19, RS20]). On the one hand, these relaxations have indeed led to tighter concrete security bounds. On the other hand, however, the extent to which the concrete security bounds resulting from these relaxations capture the security of the relevant schemes relative to hardness of the DLOG problem, is naturally rather limited.

1.1 Our Contributions

We present a complementing approach for designing multi-signature schemes in the plain public-key model and obtaining a better understanding of their security: Instead of relying on recently-introduced assumptions or considering restricted classes of attackers via the algebraic group model, we construct multi-signature schemes that satisfy relaxed notions of security. Our approach relaxes the full-fledged notion of security for multi-signature schemes in the plain public-key model by restricting attackers’ abilities to obtain signatures of their choice. This leads us to formalizing two notions of security, *one-time unforgeability* and *single-set unforgeability*, and to construct schemes that satisfy them. Although our notions are not as strong as the full-fledged one, their applicability nevertheless ranges from serving as drop-in replacements to enabling expressive smart contract validation procedures, as we discuss in Section 1.2.

One-time multi-signatures. Our notion of one-time unforgeability naturally extends the analogous single-signer notion to the multi-signer setting by considering attackers that obtain a single signature for some message and set of signers of their choice. We construct a multi-signature scheme satisfying this notion of security in the random-oracle model based on any ring-homomorphic one-way function, admitting instantiations based on the DLOG and RSA assumptions [CD98, CFG⁺15].²

Our scheme’s aggregated verification keys consist of two group elements, and when compared to the known two-round multi-signature schemes that satisfy the full-fledged notion of security for such schemes based on the hardness of the DLOG problem without relying on the algebraic group model [BD21, NRS21, TZ23], our scheme offers:³ (1) non-interactive signing, (2) aggregate signatures that consist of a single group element, and (3) security proof that consists of a single application of the forking lemma and thus avoids the substantial security loss resulting from two nested applications (without restricting adversaries to algebraic ones). In particular, when relying on the hardness of the DLOG problem, we recover the above-discussed $(t^2/p)^{1/2}$ bound, similarly to single-signer Schnorr signatures. Although here we focus mainly on the dependence on the order p of the group, we note that our concrete bound includes additional terms that depend on the number of random-oracle queries issued by attackers. In addition, we demonstrate that our scheme naturally extends to a t -time scheme, where aggregated verification keys consist of $t + 1$ group elements, while aggregated signatures still consist of a single group element.

oracle queries and signing queries issued by attackers, but rather focused mainly on the dependence on the order p of the group.

²We note that a notion of one-time unforgeability in the context of *aggregate* signatures was introduced by Boneh and Kim [BK20], as we discuss in Section 1.2.

³It is not clear how to compare the efficiency and concrete security guarantees of our one-time scheme to those of schemes that satisfy the full-fledged notion of security for multi-signature schemes. See Section 1.3 for more details.

Single-set multi-signatures. Our notion of single-set unforgeability considers attackers that obtain *any polynomial number* of signatures for messages of their choice, but are restricted to requesting all of these signatures with respect to a *single set* of signers. In this context, a single set of signers corresponds to a single vector of verification keys, which may be adversarially chosen based on the public parameters of the scheme and on the honestly-generated verification key that is attacked. As we discuss in Section 1.2, this notion already suffices for various applications of multi-signatures, such as validating blockchain transactions in a wide range of settings.

We show that any one-time multi-signature scheme with non-interactive signing can be transformed into a scheme satisfying our notion of single-set unforgeability, where the signing process of the resulting scheme consists of two rounds. Our transformation is obtained via a tree-based construction that relies on standard cryptographic tools, most notably on a non-interactive zero-knowledge proof system (we rely on these standard tools for overcoming the challenges that arise when extending the seminal Naor-Yung tree-based signature scheme [NY89] to the multi-signer setting).⁴

The length of the resulting scheme’s verification keys and aggregated verification keys is independent of the number of signers, and the keys themselves are essentially identical to those of the underlying one-time scheme. The length of the resulting scheme’s signatures and aggregated signatures is also independent of the number of signers, and determined by that of the underlying scheme while scaling linearly with the length of messages ($\ell\lambda$ -bit signatures for ℓ -bit messages, where λ is the security parameter, as in the Naor-Yung transformation⁵). Instantiated with our one-time scheme, we obtain verification keys and signatures whose lengths are completely independent of the number of signers.

Our transformation demonstrates that at least for short messages there is no inherent and significant security loss when transforming one-time multi-signatures into single-set multi-signatures, and that our security loss essentially matches that of transforming single-user one-time signatures into re-usable ones.⁶ Finally, We note that in the single-signer setting, tree-based signatures that utilize one-time signatures were initially mostly of foundational interest, whereas additional substantial efforts have demonstrated their practical applicability (see, for example, [BHH⁺15, AE18, BHK⁺19, HK22, KHR⁺22] and the references therein).

One-time multi-signatures: Structure vs. hardness. Revisiting our one-time multi-signature scheme, it is quite noticeable that whereas one-time single-signer signatures can be constructed based on any one-way function [Lam79], our multi-signer scheme is based on the more structured notion of a ring-homomorphic one-way function. Although such functions admit realizations based on standard number-theoretic assumptions [CD98, CFG⁺15], this raises the question of whether one-time multi-signatures require more structured forms of cryptographic hardness when compared to one-time single-signer signatures.

Addressing this fundamental question, we prove that one-time multi-signature schemes satisfying a natural property (which is satisfied by our one-time scheme) cannot be constructed in a fully black-box manner based on one-way functions. An interesting question for further research is whether this

⁴We emphasize that we rely on standard non-interactive zero-knowledge proofs, which can be realized based on well-studied falsifiable assumptions, and that we do not rely on succinct non-interactive arguments of knowledge (SNARKs) [Mic00, Gro10, GW11, BCI⁺13, BCS16]. In particular, we do not have any requirements regarding the length of the resulting proofs (they are not included in our signatures, and only play an intermediate role) and do not assume any form of proofs of knowledge.

⁵Without loss of generality, $\ell \leq \lambda$ as otherwise longer messages can first be hashed using a collision-resistant function.

⁶The work of Blazy, Kakvi, Kiltz, and Pan [BKK⁺15] presented a construction of re-usable signatures with a tight security reduction. However, their starting point was not a one-time signature scheme, but rather a Chameleon hash function, which is a significant more structured object.

can be circumvented via a seemingly less-natural construction.

1.2 Overview of Our Approach

In this section we first briefly discuss the applicability of schemes that satisfy our notions of security. Then, we present a high-level overview of our constructions.

The applicability of one-time and single-set multi-signatures. Our relaxed notions naturally serve as intermediate notions for obtaining a better understanding of the concrete security of full-fledged multi-signatures. At the same time, a direct application of one-time multi-signatures is for validating one-time transactions, such Bitcoin UTXOs [Nak09] (as described by Boneh and Kim [BK20] in the somewhat incomparable context of one-time aggregated signatures, which we discuss below). Specifically, a UTXO is spent after validating one or more signatures with respect to the verification keys committed in the UTXO. Once spent, it cannot be spent again, and the funds are transferred to a different UTXO. Thus, using a one-time multi-signature scheme for UTXOs can eliminate any dependence on the number of signers during verification, and to reduce both communication and storage cost. This comes at the cost of not using the same signing key for more than one UTXO, and this can be managed, for example, by deriving any number of one-time signing keys via a single master key for a pseudorandom function (thus, the one-time signing keys need not be stored, but can instead be reproduced upon demand).

A somewhat less direct application of one-time multi-signatures is for validating standard (i.e., reusable) transactions via account abstraction (e.g., [BWG⁺21, Arg22, Sta23]). At a high level, account abstraction (among its various features) enables smart contracts to offer arbitrary validation logic. Already in the single-signer setting, consider a smart contract whose storage includes a one-time verification key, and whenever the user provides a transaction they provide a signature with respect to the currently-stored verification key both on the provided input to the smart contract and on a newly-generated one-time verification key that the contract will store instead of its current one. This is motivated by the textbook path-based construction of signatures from one-time signatures [KL21], which is typically presented as a warm-up for the classic Naor-Yung tree-based construction [NY89]. Unlike the textbook construction, here the verification time and signature length do not scale with the number of previously-generated signatures, since each newly-generated verification key replaces its predecessor in the contract’s storage.⁷ Such a mechanism extends to the multi-signer setting in various ways using a one-time multi-signature scheme, where each transaction additionally updates the stored aggregated verification key. Here, the advantages of using a one-time multi-signature scheme with non-interactive signing and concrete security guarantees based on the hardness of the DLOG problem may be significant.

Finally, for our notion of single-set multi-signatures, where an adversary may observe any polynomial number of signatures for a single set of signers, account abstraction is again useful. Specifically, for any smart contract whose transactions require validating multiple signers, as long as each signer allocates a contract-specific signing key that is not used for any other purpose, then single-set security suffices (and there is no need for any key updates as with one-time multi-signatures). As discussed above, such contract-specific keys can be derived via a single master key, and thus do not have to be explicitly stored.

⁷It should be noted that, over time, such a mechanism may run into various synchronization issues since each signing key can be used only once. Such issues can be dealt with either by using a t -time multi-signature scheme, or by using a recovery mode (again, enabled by account abstraction) that allows key updates via a standard multi-signature scheme. Assuming that such issues are hopefully not-too-frequent, the overall efficiency of the validation procedure would be determined by that of the one-time scheme.

Our one-time multi-signature scheme. The starting point of our one-time scheme is the Schnorr-based one-time signature scheme designed by Bellare and Shoup [BS08], which was extended by Boneh and Kim [BK20] to DLOG-based and lattice-based one-time *aggregate* signature schemes. Recall that aggregate signature schemes enable to aggregate signatures on *any* set of messages, whereas multi-signature schemes enable to aggregate signatures on the same message. As a result, the schemes of Boneh and Kim do not support aggregation of verification keys, and their verification time scales linearly with the number of signers. We use a different aggregation method that is tailored to aggregating signatures on the same message. Once such an aggregated key has been verified to correspond to a particular set of signers (e.g., in a preliminary phase as discussed above for blockchain transactions), this enables us to guarantee that the signing and verification operations are independent of the number of signers.

Our scheme is based on the abstract notion of a ring-homomorphic one-way function, introduced by Catalano, Fiore, Gennaro and Vamvourellis [CFG⁺15]. At a high level, we consider an efficiently-computable homomorphism $f : \mathcal{X} \rightarrow \mathcal{Y}$ for cyclic groups \mathcal{X} and \mathcal{Y} that allows computing linear operations “in the exponent” over a ring $\mathbb{K} = \mathbb{Z}_q$ for some prime q (in the DLOG-based instantiation, the prime q corresponds to the order of the cyclic groups⁸, but in the RSA-based instantiation this is not the case – see Section 2.1 for a formal definition).

Each signer in our scheme samples $x, r \leftarrow \mathcal{X}$, and sets $\text{sk} = (x, r)$ and $\text{vk} = (X, R) = (f(x), f(r))$ as their signing key and verification key, respectively. Then, a vector $\vec{\text{vk}} = ((X_1, R_1), \dots, (X_n, R_n))$ of verification keys, corresponding to n signers, is aggregated by computing $\text{aggvk} = (\text{agg}X, \text{agg}R) = (\prod_{i=1}^n X_i^{a_i}, \prod_{i=1}^n R_i^{a_i})$, where $(a_1, \dots, a_n) = \text{H}_1(\vec{\text{vk}})$ for a hash function H_1 . For any message m and vector $\vec{\text{vk}}$ of verification keys, each signer computes a signature $\sigma_i = r_i + \text{H}_0(m, \text{aggvk}) \cdot x_i$ for a hash function H_0 , and signatures are similarly aggregated by computing $\text{agg}\sigma = \sum_{i=1}^n a_i \cdot \sigma_i$. In turn, this enables to verify an aggregate signature $\text{agg}\sigma$ with respect to an aggregated verification key $\text{aggvk} = (\text{agg}X, \text{agg}R)$ by checking whether $f(\text{agg}\sigma) = (\text{agg}X)^{\text{H}_0(m, \text{aggvk})} \cdot \text{agg}R$.

Our security proof models the hash function H_0 and H_1 as random oracles, and reduces the task of breaking the one-time unforgeability of the scheme to that of breaking the one-wayness property of the ring-homomorphic function (see Definition 2.1). Specifically, given any attacker for our scheme, we construct an inverter that is given $X = f(x)$ for a randomly chosen x , and outputs a pair (x', d) such that $f(x') = X^d$ and $d \neq 0$ (note that, in the DLOG-setting d can always be efficiently inverted, and thus without loss of generality $d = 1$, but in the RSA-setting this is not the case). At a very high level, our proof programs the random oracle H_0 for embedding the value X in the honestly-generated verification key given as input to the attacker, while generating R in a way that would be consistent with the response to the attacker’s single signing query. Then, the proof relies on a single application of the forking lemma using the random oracle H_1 for resampling the value of a_i that corresponds to position of the honestly-generated verification key in the aggregated verification key with respect to which the attacker produces a forgery. Given two such forgeries, we are then able to efficiently produce x' and d as required. The proof naturally contains a variety of challenges for implementing this high-level idea, and we refer the reader to Section 4 for a complete and formal description.

From one-time to single-set multi-signatures. As mentioned above, our approach is inspired by the Naor-Yung transformation of a one-time signature scheme into a reusable one [NY89]. Recall that, in the Naor-Yung transformation, the signer implicitly holds a binary tree of exponential size, where each node of the tree is associated with a pair of one-time signing and verification keys. Specifically, for signing ℓ -bit messages, the tree has 2^ℓ leaves, where each leaf and each internal node

⁸In the DLOG-based instantiation, the homomorphism is simply the group exponentiation operation relative to a given generator.

$\alpha \in \{0, 1\}^{\leq \ell} \cup \{\varepsilon\}$ is associated with a pair $(\text{sk}_\alpha, \text{vk}_\alpha)$ of one-time keys.⁹ The keys sk_ε and vk_ε corresponding to the root ε serve as the signing key and verification key, respectively, and all other keys (and randomness that may be needed for using them) do not have to be explicitly generated or stored, but rather can be produced whenever needed using a pseudorandom function whose key is additionally included in the signing key.

For signing a message $m \in \{0, 1\}^\ell$, the signer first uses the signing key sk_m associated with the leaf correspond to the binary string m for signing the message m itself. Then, it uses the signing keys positioned on the path connecting the root ε to the leaf m for certifying the path: For every $j \in \{0, \dots, \ell - 1\}$, the signer uses the key $\text{sk}_{m|_j}$ associated with the internal node corresponding to the binary string $m|_j$ to sign the concatenation of the two verification keys $\text{vk}_{m|_j0}$ and $\text{vk}_{m|_j1}$ corresponding to its two children in the tree.¹⁰ This makes sure that each signing key is used at most once, thus enabling to rely on the one-time security of the underlying scheme.

Equipped with the Naor-Yung transformation, let us attempt extending it to the multi-signer setting. As in the single-signer setting, suppose that each signer implicitly holds a tree, where each node $\alpha \in \{0, 1\}^{\leq \ell} \cup \{\varepsilon\}$ is associated with a pair $(\text{sk}_\alpha, \text{vk}_\alpha)$ of keys for a one-time *multi-signature* scheme. The keys sk_ε and vk_ε corresponding to the root ε again serve as the signing key and verification key, respectively, and all other keys are similarly generated using a pseudorandom function upon demand. Note that this structure enables to aggregate the root verification keys $\text{vk}_\varepsilon^{(1)}, \dots, \text{vk}_\varepsilon^{(n)}$ of any n signers by using the key-aggregation algorithm of the underlying one-time scheme. More generally, it enables to aggregate not only the root keys, but to implicitly define an *aggregated tree*. In the aggregated tree, each node $\alpha \in \{0, 1\}^{\leq \ell} \cup \{\varepsilon\}$ is associated with an aggregated verification key aggvk_α that is obtained by aggregating the one-time verification keys $\text{vk}_\alpha^{(1)}, \dots, \text{vk}_\alpha^{(n)}$ associated with the node α in the n individual trees.

This observation leads to the following elegant, yet insecure, two-round signing protocol. For signing a message m with respect to signers with root verification keys $\text{vk}_\varepsilon^{(1)}, \dots, \text{vk}_\varepsilon^{(n)}$, each signer first sends all other signers the 2ℓ verification keys on the path leading from the signer's root to the leaf m . At this point, all signers know the verification keys $\text{vk}_{m|_j b}^{(i)}$ for all $i \in [n]$, $j \in \{0, \dots, \ell - 1\}$ and $b \in \{0, 1\}$. This enables each signer to compute the aggregated verification keys $\text{aggvk}_{m|_j b}$ along the path from the root to the leaf m in the aggregated tree. Now, for each level $j \in \{0, \dots, \ell - 1\}$, each signer uses their one-time signing key $\text{sk}_{m|_j}^{(i)}$ to non-interactively compute a signature $\sigma_{m|_j}^{(i)}$ on the concatenation of the two aggregated verification keys $\text{aggvk}_{m|_j 0}$ and $\text{aggvk}_{m|_j 1}$ with respect to the signer set $\text{vk}_{m|_j}^{(1)}, \dots, \text{vk}_{m|_j}^{(n)}$. Finally, each signer uses their one-time signing key $\text{sk}_m^{(i)}$ to compute a signature $\sigma_m^{(i)}$ on the message m with respect to the signer set $\text{vk}_m^{(1)}, \dots, \text{vk}_m^{(n)}$. These signatures are then aggregated for each level $j \in \{1, \dots, \ell\}$ to produce a signature that consists of ℓ aggregated one-time signatures. Note that both the length of the resulting signature and the time required to verify it scale linearly with the length of the message (as in the Naor-Yung scheme), but are completely independent of the number n of signers.

At this point, we would like to argue that since each one-time signing key is used at most once, then we can rely on the security of the underlying one-time multi-signature scheme to claim that our tree-based scheme is secure against attackers issuing any polynomial number of signing queries. This argument fails, however, if the attacker can request signatures with respect to more than one set of signers. The reason is that different sets of signers induce different aggregated trees. Consider an adversary that requests a signature from some signer i with respect to a set \mathcal{S} , and then requests

⁹We denote by $\{0, 1\}^{\leq \ell}$ the set of all binary strings of length at most ℓ , and by ε the empty string.

¹⁰We denote by $m|_j$ the leftmost j bits of a binary string m (where $m|_0 = \varepsilon$), and we denote by $m|_j b$ the binary string obtained by concatenating the strings $m|_j$ and b .

a signature from the same signer i with respect to a different set \mathcal{S}' . The first signature includes a signature with respect to the one-time signing key $\text{sk}_\varepsilon^{(i)}$ on information derived from the aggregated tree corresponding to \mathcal{S} , and the second signature includes a signature with respect to the same one-time signing key $\text{sk}_\varepsilon^{(i)}$ on information derived from the aggregated tree corresponding to \mathcal{S}' . Since the underlying scheme is only guaranteed to be one-time secure, the security of the tree-based construction breaks down.

Still, one might hope that the construction is secure as long as the attacker issues signature queries with respect to a single set of signers, eliminating the issue we just described. This coincides with our single-set notion of security. However, this is not the case. This added restriction is insufficient because a root verification key may be used by malicious signers for different trees. That is, a malicious signer can still send different verification keys in the first round of two invocations of the signing protocol. This again results in two distinct aggregated trees (even though the two invocations of the signing protocol share the same set of signers).

We resolve this additional challenge by identifying signers not only with their root verification key, but also with a commitment containing a key for a pseudorandom function from which their entire tree is derived. Now, in the first round of the signing protocol, each signer sends all one-time verification keys on the path from their root to the respective leaf, while augmenting each such key with a non-interactive zero-knowledge proof asserting that it has been generated correctly. From a foundational standpoint, such proofs can be based on the existence of trapdoor functions¹¹ [FLS90]. For practical instantiations, these NIZK proofs can be based on one of the many recent efficient protocols¹². Crucially, these proofs are not included as part of the resulting signature, but rather only serve as “proofs of semi-honest behavior” that enable each signer to continue to the second round of the signing protocol. This describes the high-level intuitive structure of our scheme, and we refer the reader to Section 5 for a complete and formal description.

1.3 Related Work

DDH-based multi-signatures. Whereas most of the work on multi-signatures in prime-order groups focused on DLOG-based schemes, several schemes were suggested also based on the DDH assumption (e.g. [LYG19, FH21, TSS⁺23, PW23]). As in the single-signer setting, DDH-based signatures may lead to tighter reductions. Although, when implemented in concrete groups, such schemes typically do not offer the same efficiency guarantees as DLOG-based ones.

DLOG-based two-round multi-signatures. Existing DLOG-based two-round multi-signature schemes can be roughly divided into three categories: (1) schemes whose security is established in the algebraic group model [AB21, BD21, NRS21, LK23], (2) schemes whose security is established based on interactive variants of the DLOG problem (without relying on the algebraic group model) [BD21, NRS21], and (3) schemes whose security is established based on the standard DLOG problem (again, without relying on the algebraic group model) [BD21, NRS21, TZ23]. As discussed in Section 1.1, it is challenging to compare the efficiency and concrete security guarantees of our one-time scheme to those of the existing DLOG-based schemes, as these schemes satisfy the full-fledged notion of security for multi-signature schemes. If any comparison can be made, it would seem essential for focus on category 3, since the schemes in categories 1 and 2 seem to inherently avoid nested applications of the forking lemma.

¹¹Technically speaking, these have to be certifiably injective and doubly-enhanced [BY96, Gol11, GR13, CL18].

¹²See [GS08, Gro16, BBB⁺18, BSBH⁺18, GWC19, CHM⁺20, XZS22, GLS⁺23] and the many references therein for a highly non-exhaustive list of examples.

Focusing on category 3, the schemes of Bellare and Dai [BD21], Nick, Ruffing and Seurin [NRS21], and Tessaro and Zhu [TZ23] rely on two nested applications of the forking lemma (whereas we rely on a single application), their verification keys consist of a single group element (whereas our verification keys consist of two group elements), and their signatures consist of two, two and three group elements, respectively (whereas our signatures consist of a single group element).

Synchronized multi-signatures. A substantially different tree-based approach, both in terms of its goals and in terms of its structure, was recently presented by Fleischhacker, Simkin and Zhang [FSZ22]. They constructed a lattice-based multi-signature scheme in the *synchronized* model, where it is assumed that signers share a global notion of time, and the signing algorithm takes the current time step as an additional input. Most notably, it is additionally assumed that no signer signs more than one message per time step, and the goal is to aggregate signatures for the same message and *same time step, without knowing the set of signers in advance*. Thus, a signature may be aggregated together with those of any subset of other signers. As noted by Fleischhacker et al. such flexibility seems particularly useful in the blockchain setting for the task of confirming newly-generated blocks: Block validators may be synchronized by the number of the block that they sign, each validator does not sign more than one block in each time period, and validators do not know which of the other potential validators will actually participate.

In contrast to the synchronized model, we design our scheme in the plain public key model [BN06], where there is no global notion of time (or any other form of synchronization), and we aim at aggregating signatures not for a given time step but rather for a given set of signers that is specified during the signing process. In particular, our scheme guarantees that a signature may be aggregated only with a specific set of signers that is provided to the signing algorithm as input, while keeping the verification time independent of the number of signers.

From the technical perspective, as discussed above, our approach relies on trees of exponential size, which are never explicitly constructed in their entirety. Each of the exponential number of leaves corresponds to a message-dedicated verification key for a one-time multi-signature scheme, and the path leading to each leaf is constructed upon demand using a pseudorandom function. In contrast, for supporting T time periods, Fleischhacker et al. explicitly sample T key pairs, and then compute a homomorphic Merkle tree with the corresponding T verification keys as its leaves. As a result, their scheme seems limited to supporting only a polynomial number T of time periods, and the efficiency of their key-generation algorithm scales linearly with T .

Finally, we note that since messages in their scheme are not signed with respect to a given set of signers (but rather a signature can be aggregated together with those of any subset of other signers), the notion of security required from their one-time primitive does not explicitly consider multiple signers. Specifically, Fleischhacker et al. introduce a notion of single-signer one-time key-homomorphic signatures, whereas we explicitly introduce a notion of one-time multi-signatures satisfying security guarantees tailored to the multi-signer setting, and our tree-based construction can rely on any multi-signature scheme that satisfies it.

2 Preliminaries

In this section we present the basic notions and the cryptographic primitives and tools that are used in this work. For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$. For a distribution X we denote by $x \leftarrow X$ the process of sampling a value x from the distribution X . Similarly, for a set \mathcal{X} we denote by $x \leftarrow \mathcal{X}$ the process of sampling a value x from the uniform distribution over \mathcal{X} .

2.1 Ring-Homomorphic One-Way Functions

Our construction of a one-time multi-signature scheme relies on the notion of ring-homomorphic one-way functions, introduced by Catalano, Fiore, Gennaro and Vamvourellis [CFG⁺15]. This notion was presented by Catalano et al. as part of their framework for algebraic one-way functions, which is closely-related to the notion of group-homomorphic one-way functions introduced by Cramer and Damgård [CD98]. The notion considers function families $F = (\text{Setup}, \text{Eval})$ where for any $\lambda \in \mathbb{N}$ and for any function index Ind produced by $F.\text{Setup}(1^\lambda)$ it holds that $F.\text{Eval}(\text{Ind}, \cdot) : \mathcal{X}_{\text{Ind}} \rightarrow \mathcal{Y}_{\text{Ind}}$ is an efficiently-computable homomorphism for cyclic groups \mathcal{X}_{Ind} and \mathcal{Y}_{Ind} that allows computing linear operations “in the exponent” over a ring \mathbb{K}_{Ind} . For our purposes, we consider the specific case where $\mathbb{K}_{\text{Ind}} = \mathbb{Z}_q$ for some prime $q = q(\text{Ind})$, and note that already this case captures the known constructions based on the hardness of the discrete-logarithm problem (i.e., the exponentiation function in a cyclic group) and RSA problem (i.e., the RSA function) as we discuss below [CD98, CFG⁺15].

As formalize by Catalano et al. [CFG⁺15], such a function family F is *ring homomorphic* if there exists an efficient algorithm Eval' such that for any $\lambda \in \mathbb{N}$, Ind produced by $\text{Setup}(1^\lambda)$, generators $h_1, \dots, h_m \in \mathcal{X}_{\text{Ind}}$, vector of elements $\mathbf{W} = (W_1, \dots, W_\ell) \in \mathcal{X}_{\text{Ind}}^\ell$ where $W_i = h_1^{\omega_i^{(1)}} \cdots h_m^{\omega_i^{(m)}} \cdot R_i$, for some $R_i \in \mathcal{X}_{\text{Ind}}$ and some integers $\omega_i^{(j)} \in \mathbb{Z}$ (note that this decomposition may not be unique), and vector of integers $\boldsymbol{\alpha} \in \mathbb{Z}^\ell$, it holds that

$$\text{Eval}'(\text{Ind}, \mathbf{A}, \mathbf{W}, \boldsymbol{\Omega}, \boldsymbol{\alpha}) = h_1^{\langle \boldsymbol{\omega}^{(1)}, \boldsymbol{\alpha} \rangle} \cdots h_m^{\langle \boldsymbol{\omega}^{(m)}, \boldsymbol{\alpha} \rangle} \prod_{i=1}^{\ell} R_i^{\alpha_i}$$

where $\mathbf{A} = (A_1, \dots, A_m) \in \mathcal{Y}_{\text{Ind}}^m$ is such that $A_i = F_{\text{Ind}}(h_i)$, $\boldsymbol{\Omega} = \left(\omega_i^{(j)} \right)_{i,j} \in \mathbb{Z}^{\ell \times m}$, and each product $\langle \boldsymbol{\omega}^{(j)}, \boldsymbol{\alpha} \rangle$ in the exponent is computed over the ring \mathbb{K}_{Ind} .

In terms of one-wayness, Catalano et al. formalized the following notion of *flexible one-wayness*, asking that given $X = F_{\text{Ind}}(x)$ for a uniformly distributed $x \in \mathcal{X}_{\text{Ind}}$, it should be infeasible to output $x' \in \mathcal{X}_{\text{Ind}}$ and $d \in \mathbb{K}_{\text{Ind}}$ such that $F_{\text{Ind}}(x') = X^d$ and $d \neq 0_{\mathbb{K}_{\text{Ind}}}$.

Definition 2.1. A ring-homomorphic function family $F = (\text{Setup}, \text{Eval})$ is *flexible one-way* if for every probabilistic polynomial-time algorithm A there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}_{F,A}^{\text{hom-ow}}(\lambda) \stackrel{\text{def}}{=} \Pr \left[A(\text{Ind}, X) = (x', d) \text{ s.t. } \left[F_{\text{Ind}}(x') = X^d \text{ and } d \neq 0_{\mathbb{K}_{\text{Ind}}} \right] \right] \leq \nu(\lambda),$$

where $\text{Ind} \leftarrow \text{Setup}(1^\lambda)$, $x \leftarrow \mathcal{X}_{\text{Ind}}$ and $X = F_{\text{Ind}}(x)$.

Note that if $\mathbb{K}_{\text{Ind}} = \mathbb{Z}_q$, where q is the order of \mathcal{X}_λ , then the above definition is equivalent to the standard notion of one-wayness. In particular, the hardness of computing discrete logarithms in cyclic groups is equivalent to the flexible one-wayness of the group exponentiation function in such groups. In addition, Catalano et al. [CFG⁺15] showed that the RSA function $x \rightarrow x^e \bmod N$ in the subgroup $\mathbb{QR}_N \subset \mathbb{Z}_N^*$ of quadratic residues (where N is the product of two “safe primes” and thus \mathbb{QR}_N is cyclic) is flexible one-way based on the RSA assumption. In this case, $\mathbb{K}_{\text{Ind}} = \mathbb{Z}_e$ for any prime $e \geq 3$.

2.2 The Forking Lemma

The proof of security for our one-time multi-signature scheme relies on the “forking lemma” of Bellare and Neven [BN06] (following Pointcheval and Stern [PS00]). Let $q \geq 1$ be an integer, and let \mathcal{H}, \mathcal{X}

and \mathcal{Y} be a sets such that $|\mathcal{H}| \geq 2$. Let A be a randomized algorithm that on input $(x, \vec{h}) \in \mathcal{X} \times \mathcal{H}^q$ returns either a pair $(i, y) \in [q] \times \mathcal{Y}$ or the dedicated symbol \perp . Let F_A be an algorithm that takes inputs in \mathcal{X} and returns either an output $(y, y') \in \mathcal{Y}^2$ or the dedicated symbol \perp , and is defined as follows:

1. Sample random coins $\rho \leftarrow \{0, 1\}^*$ for A .
2. Sample $h_1, \dots, h_q \leftarrow \mathcal{H}$ and compute $\text{out}_1 = A(x, h_1, \dots, h_q; \rho)$.
3. If $\text{out}_1 = \perp$ then return \perp , and otherwise let $\text{out}_1 = (i, y)$.
4. Sample $h'_1, \dots, h'_q \leftarrow \mathcal{H}$ and compute $\text{out}_2 = A(x, h_1, \dots, h_{i-1}, h'_i, \dots, h_q; \rho)$.
5. If $\text{out}_2 = \perp$ then return \perp , and otherwise let $\text{out}_2 = (i', y')$.
6. If $i' = i$ and $h_i \neq h'_i$ then return (i, y, y') , and otherwise return \perp .

The following lemma, due to Bellare and Neven [BN06], relates the probability that F_A successfully provides an output (other than \perp) to the corresponding probability of A .

Lemma 2.2 ([BN06]). *For any algorithm A and for any distribution D over \mathcal{X} it holds that*

$$\Pr_{x \leftarrow D} [F_A(x) \neq \perp] \geq \epsilon \cdot \left(\frac{\epsilon}{q} - \frac{1}{|\mathcal{H}|} \right),$$

where

$$\epsilon = \Pr_{\substack{x \leftarrow D \\ \vec{h} \leftarrow \mathcal{H}^q}} [A(x, \vec{h}) \neq \perp].$$

2.3 Additional Cryptographic Primitives

In this section we formally define additional basic cryptographic primitives that are used as building blocks in our construction of a single-set multi-signature scheme: Pseudorandom functions, collision-resistant hash functions, non-interactive statistically-binding commitment schemes, and non-interactive zero-knowledge proofs.

Pseudorandom functions. For integers ℓ_{in} and ℓ_{out} we denote by $\text{Func}_{\ell_{\text{in}}, \ell_{\text{out}}}$ the set of all functions mapping ℓ_{in} -bit inputs to ℓ_{out} -bit outputs. We rely on the following standard notion of pseudorandom functions (e.g., [Gol01]):

Definition 2.3. Let $\ell_{\text{in}} = \ell_{\text{in}}(\lambda)$ and $\ell_{\text{out}} = \ell_{\text{out}}(\lambda)$ be functions of the security parameter $\lambda \in \mathbb{N}$. A *pseudorandom function* is a pair of polynomial-time algorithms $\Pi = (\text{KG}, \text{Eval})$ with the following two properties:

1. **Eval** is a deterministic algorithm such that for any $\lambda \in \mathbb{N}$ and for any \mathbf{k} produced by $\text{KG}(1^\lambda)$ it holds that $\text{Eval}(\mathbf{k}, \cdot) : \{0, 1\}^{\ell_{\text{in}}(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{out}}(\lambda)}$.
2. For every probabilistic polynomial-time algorithm A there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}_{\Pi, A}^{\text{prf}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[A^{\text{Eval}(\mathbf{k}, \cdot)}(1^\lambda) = 1 \right] - \Pr \left[A^{f(\cdot)}(1^\lambda) = 1 \right] \right| \leq \nu(\lambda),$$

where the probability is taken over the choices of $\mathbf{k} \leftarrow \text{KG}(1^\lambda)$, and $f \leftarrow \text{Func}_{\ell_{\text{in}}(\lambda), \ell_{\text{out}}(\lambda)}$, and over the internal randomness of A .

Collision-resistant hash functions. We rely on the standard notion of collision-resistant hash functions (e.g., [Gol01]):

Definition 2.4. Let $\ell_{\text{in}} = \ell_{\text{in}}(\lambda)$ and $\ell_{\text{out}} = \ell_{\text{out}}(\lambda)$ be functions of the security parameter $\lambda \in \mathbb{N}$. A *collision-resistant hash family* is a pair of polynomial-time algorithms $\Pi = (\text{KG}, \text{Eval})$ with the following two properties:

1. Eval is a deterministic algorithm such that for any $\lambda \in \mathbb{N}$ and for any \mathbf{k} produced by $\text{KG}(1^\lambda)$ it holds that $\text{Eval}(\mathbf{k}, \cdot) : \{0, 1\}^{\ell_{\text{in}}(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{out}}(\lambda)}$.
2. For every probabilistic polynomial-time algorithm A there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}_{\Pi, A}^{\text{crh}}(\lambda) \stackrel{\text{def}}{=} \Pr \left[A(1^\lambda, \mathbf{k}) = (x, x') \text{ s.t. } \begin{array}{l} x \neq x' \in \{0, 1\}^{\ell_{\text{in}}(\lambda)} \text{ and} \\ \text{Eval}(\mathbf{k}, x) = \text{Eval}(\mathbf{k}, x') \end{array} \right],$$

where the probability is taken over the choice of $\mathbf{k} \leftarrow \text{KG}(1^\lambda)$ and over the internal randomness of A .

Non-interactive statistically-binding commitment schemes. We rely on the following standard notion of a non-interactive statistically-binding commitment scheme [Nao91, Gol01]:

Definition 2.5. Let $\epsilon_{\text{binding}} = \epsilon_{\text{binding}}(\lambda)$ be a function of the security parameter $\lambda \in \mathbb{N}$. A *non-interactive $\epsilon_{\text{binding}}$ -statistically-binding commitment scheme* for a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ is a triplet of probabilistic polynomial-time algorithms $\Pi = (\text{Setup}, \text{Commit}, \text{Verify})$ with the following properties:

1. **Perfect completeness:** For every $\lambda \in \mathbb{N}$ and $m \in \mathcal{M}_\lambda$ it holds that

$$\Pr \left[\text{Verify}(\text{crs}, \text{com}, \text{decom}, m) = 1 \mid \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda) \\ (\text{com}, \text{decom}) \leftarrow \text{Commit}(\text{crs}, m) \end{array} \right] = 1$$

where the probability is taken over the internal randomness of Setup , Commit and Verify .

2. **Statistical binding:** For every $\lambda \in \mathbb{N}$ it holds that

$$\Pr \left[\begin{array}{l} \exists \text{com}, m \neq m', \text{decom}, \text{decom}' \text{ s.t.} \\ \text{Verify}(\text{crs}, \text{com}, \text{decom}, m) = \text{Verify}(\text{crs}, \text{com}, \text{decom}', m') = 1 \end{array} \right] \leq \epsilon_{\text{binding}}(\lambda)$$

where the probability is taken over the choice of $\text{crs} \leftarrow \text{Setup}(1^\lambda)$.

3. **Computational hiding:** For every probabilistic polynomial-time algorithm $A = (A_1, A_2)$ there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}_{\Pi, A}^{\text{hiding}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Exp}_{\Pi, A}^{\text{hiding}}(0, \lambda) = 1 \right] - \Pr \left[\text{Exp}_{\Pi, A}^{\text{hiding}}(1, \lambda) = 1 \right] \right| \leq \nu(\lambda)$$

for all sufficiently large n , where for each $b \in \{0, 1\}$ the experiment $\text{Exp}_{\Pi, A}^{\text{hiding}}(b, \lambda)$ is defined as:

- (a) $\text{crs} \leftarrow \text{Setup}(1^\lambda)$.
- (b) $(m_0, m_1, \text{st}) \leftarrow A_1(1^\lambda, \text{crs})$.
- (c) $(\text{com}, \text{decom}) \leftarrow \text{Commit}(\text{crs}, m_b)$.
- (d) $b' \leftarrow A_2(\text{crs}, \text{com}, \text{st})$.
- (e) Output b' .

Non-interactive zero-knowledge proof systems. We rely on the following standard notion of a non-interactive simulation-sound zero-knowledge proof system [BFM88, BSM⁺91, FLS90, Sah99, SCO⁺01, Gro06, GOS06, Lin06]:

Definition 2.6. A *non-interactive simulation-sound zero-knowledge proof system* for a language $\mathcal{L} = \{\mathcal{L}_\lambda\}_{\lambda \in \mathbb{N}}$ with a witness relation $R(\mathcal{L}) = \{R_\lambda\}_{\lambda \in \mathbb{N}}$ is a 5-tuple of probabilistic polynomial-time algorithms $\Pi = (\text{Setup}, \text{P}, \text{V}, \text{Sim}_1, \text{Sim}_2)$ with the following properties:

1. **Perfect completeness:** For every $\lambda \in \mathbb{N}$ and $(x, w) \in R_\lambda$ it holds that

$$\Pr \left[\text{V}(1^\lambda, \text{crs}, x, \pi) = 1 \mid \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda) \\ \pi \leftarrow \text{P}(1^\lambda, \text{crs}, x, w) \end{array} \right] = 1$$

where the probability is taken over the internal randomness of **Setup**, **P** and **V**.

2. **Adaptive zero knowledge:** For every probabilistic polynomial-time algorithm A there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}_{\Pi, A}^{\text{zk}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr \left[\text{Exp}_{\Pi, A}^{\text{zk}}(\lambda) = 1 \right] - \Pr \left[\text{Exp}_{\Pi, A, \text{Sim}_1, \text{Sim}_2}^{\text{zk}}(\lambda) = 1 \right] \right| \leq \nu(\lambda)$$

for all sufficiently large n , where the experiment $\text{Exp}_{\Pi, A}^{\text{zk}}(\lambda)$ is defined as:

- (a) $\text{crs} \leftarrow \text{Setup}(1^\lambda)$.
- (b) $b \leftarrow A^{\text{P}(1^\lambda, \text{crs}, \cdot)}(1^\lambda, \text{crs})$.
- (c) Output b .

and the experiment $\text{Exp}_{\Pi, A, \text{Sim}_1, \text{Sim}_2}^{\text{zk}}(\lambda)$ is defined as:

- (a) $(\text{crs}, \tau) \leftarrow \text{Sim}_1(1^\lambda)$.
- (b) $b \leftarrow A^{\text{Sim}'_2(1^\lambda, \tau, \cdot)}(1^\lambda, \text{crs})$, where $\text{Sim}'_2(1^\lambda, \tau, x, w) = \text{Sim}_2(1^\lambda, \tau, x, w)$ if $(x, w) \in R_\lambda$ and $\text{Sim}'_2(1^\lambda, \tau, x, w) = \perp$ otherwise.
- (c) output b .

3. **Simulation soundness:** For every probabilistic polynomial-time algorithm A there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}_{\Pi, A}^{\text{ss}}(\lambda) \stackrel{\text{def}}{=} \Pr \left[\text{Exp}_{\Pi, A}^{\text{ss}}(\lambda) = 1 \right] \leq \nu(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$, where the experiment $\text{Exp}_{\Pi, A}^{\text{ss}}(\lambda)$ is defined as:

- (a) $(\text{crs}, \tau) \leftarrow \text{Sim}_1(1^\lambda)$.
- (b) $(x, \pi) \leftarrow A^{\text{Sim}_2(1^\lambda, \tau, \cdot)}(1^\lambda, \text{crs})$.
- (c) Denote by Q the set of Sim_2 's query-response pairs.
- (d) Output 1 if and only if $x \notin \mathcal{L}_\lambda$, $(x, \pi) \notin Q$, and $\text{V}(1^\lambda, \text{crs}, x, \pi) = 1$.

3 One-Time and Single-Set Security for Multi-Signature Schemes

In this section we formalize our two notions of security for multi-signature schemes. These notions are obtained by relaxing the notion of security for multi-signature schemes, formalized by Bellare and Neven [BN06] and recently refined by Bellare and Dai [BD21], by restricting adversaries' abilities to obtain signatures of their choice. In what follows we briefly recall the syntax and correctness requirement of multi-signature schemes, and then formally present our notions of security in Sections 3.1 and 3.2.

A multi-signature scheme is a six-tuple $\Pi = (\text{Setup}, \text{KG}, \text{KAgg}, \text{Sign}, \text{SAgg}, \text{Verify})$ of polynomial-time algorithms. The setup algorithm Setup receives as input the unary representation of the security parameter $\lambda \in \mathbb{N}$ and outputs public parameters pp . The key-generation algorithm KG receives as input the public parameters pp , and outputs a signing key sk and a verification key vk . The key-aggregation algorithm KAgg is a deterministic algorithm that takes as input the public parameters pp and a vector of verification keys $\vec{\text{vk}}$, and outputs an aggregated verification key aggvk .¹³ For schemes with non-interactive signing, the signing algorithm Sign receives as input the public parameters pp , a signing key sk , a vector $\vec{\text{vk}}$ of verification keys, and a message m that is taken from a message space \mathcal{M} , and outputs a signature σ . For schemes with interactive signing, the signing algorithm defines an interactive protocol by additionally receiving as input at each round the relevant party's internal state the communication produced by all other parties. The signature-aggregation algorithm SAgg is a deterministic algorithm that takes as input the public parameters pp , a vector of verification keys $\vec{\text{vk}}$, and a vector of signatures $\vec{\sigma}$, and outputs an aggregated signature σ . Finally, the verification algorithm Verify receives as input the public parameters pp , an aggregated verification key aggvk , a message m and an aggregated signature σ , and outputs either 0 or 1.

In terms of correctness, we consider the following standard requirement, which we formalize for simplicity for schemes with non-interactive signing and then discuss its extension to schemes with interactive signing:

Definition 3.1. A multi-signature scheme $\Pi = (\text{Setup}, \text{KG}, \text{KAgg}, \text{Sign}, \text{SAgg}, \text{Verify})$ with non-interactive signing over a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ is *perfectly correct* if for any polynomial $n = n(\cdot)$, security parameter $\lambda \in \mathbb{N}$, and message $m \in \mathcal{M}_\lambda$ it holds that

$$\Pr \left[\text{Verify} \left(\text{pp}, \text{KAgg} \left(\text{pp}, \vec{\text{vk}} \right), m, \text{SAgg} \left(\text{pp}, \vec{\text{vk}}, \vec{\sigma} \right) \right) = 1 \right] = 1$$

for $\vec{\text{vk}} = (\text{vk}_1, \dots, \text{vk}_n)$ and $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$, where the probability is taken over the choice of $\text{pp} \leftarrow \text{Setup}(1^\lambda)$, and over the choices of $(\text{sk}_i, \text{vk}_i) \leftarrow \text{KG}(\text{pp})$ and $\sigma_i \leftarrow \text{Sign} \left(\text{pp}, \text{sk}_i, \vec{\text{vk}}, m \right)$ for every $i \in [n]$.

The above definition extends to schemes with interactive signing by letting $(\sigma_1, \dots, \sigma_n)$ denote the local output of each party in the interactive signing protocol, where each party $i \in [n]$ is provided with $(\text{pp}, \text{sk}_i, \vec{\text{vk}}, m)$ as input. We refer the reader to the work of Bellare and Dai [BD21] for a formal treatment of the correctness requirement for schemes with interactive signing.

A standard relaxation of the above definition allows for a negligible error probability. Concretely, our single-set multi-signature scheme presented in Section 5 provides perfect correctness whenever the vector $\vec{\text{vk}}$ consists of distinct verification keys. Since essentially any notion of security for signature schemes guarantees that collisions among honestly-generated verification keys may occur only with a negligible probability, this yields at most a negligible error probability.

¹³For our framework we view collections of verification keys as vectors and not sets. Note that any set can be uniquely transformed into a vector by determining an order among its elements (e.g., lexicographic order).

3.1 One-Time Unforgeability

For presenting the notion of one-time unforgeability for multi-signature schemes, we focus on multi-signature schemes with non-interactive signing, and note that the following treatment naturally extends to schemes with interactive signing (our one-time multi-signature scheme in Section 4 has non-interactive signing).

This notion of security captures attacks in which an adversary may obtain a single signature of their choice. Specifically, we consider a security experiment in which the adversary first receives the public parameters \mathbf{pp} of the scheme and an honestly-generated verification key \mathbf{vk} . Then, the adversary may issue a single signing query for some message m with respect to some set of signers indicated via a vector $\vec{\mathbf{vk}} = (\mathbf{vk}_1, \dots, \mathbf{vk}_n)$ of verification keys. Both the message m and the vector $\vec{\mathbf{vk}}$ of verification keys may be arbitrarily chosen (in polynomial time) by the adversary based on the public parameters and the honest verification key. Next, the adversary obtains a signature $\sigma \leftarrow \text{Sign}(\mathbf{pp}, \mathbf{sk}, \vec{\mathbf{vk}}, m)$ produced using the signing key \mathbf{sk} corresponding to the honest verification key \mathbf{vk} (note that if $\mathbf{vk} \notin \vec{\mathbf{vk}}$ then the signing algorithm may be defined to output \perp), and the adversary's goal is to output a non-trivial forgery $(\vec{\mathbf{vk}}^*, m^*, \text{agg}\sigma^*)$. The non-triviality of the forgery is reflected in the fact that $\mathbf{vk} \in \vec{\mathbf{vk}}^*$ and $(\vec{\mathbf{vk}}^*, m^*) \neq (\vec{\mathbf{vk}}, m)$ (i.e., it could not have been directly obtained as the result of the signing query), and that the aggregated signature $\text{agg}\sigma^*$ verifies correctly for the message m^* with respect to the aggregate verification key corresponding to $\vec{\mathbf{vk}}^*$.

Definition 3.2. Let $t = t(\lambda)$ and $\epsilon = \epsilon(\lambda)$ be functions of the security parameter $\lambda \in \mathbb{N}$. A non-interactive multi-signature scheme $\Pi = (\text{Setup}, \text{KG}, \text{KAgg}, \text{Sign}, \text{SAgg}, \text{Verify})$ is *one-time* (t, ϵ) -unforgeable if for any algorithm $A = (A_1, A_2)$ that runs in time at most t it holds that

$$\text{Adv}_{\Pi, A}^{1\text{TimeMS}}(\lambda) \stackrel{\text{def}}{=} \Pr \left[\text{Exp}_{\Pi, A}^{1\text{TimeMS}}(\lambda) = 1 \right] \leq \epsilon(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$, where the experiment $\text{Exp}_{\Pi, A}^{1\text{TimeMS}}(\lambda)$ is defined as follows:

1. $\mathbf{pp} \leftarrow \text{Setup}(1^\lambda)$.
2. $(\mathbf{sk}, \mathbf{vk}) \leftarrow \text{KG}(\mathbf{pp})$.
3. $(\vec{\mathbf{vk}}, m, \text{st}) \leftarrow A_1(1^\lambda, \mathbf{pp}, \mathbf{vk})$.
4. $\sigma \leftarrow \text{Sign}(\mathbf{pp}, \mathbf{sk}, \vec{\mathbf{vk}}, m)$.
5. $(\vec{\mathbf{vk}}^*, m^*, \text{agg}\sigma^*) \leftarrow A_2(\text{st}, \sigma)$.
6. If the following conditions are satisfied then output 1 and otherwise output 0:
 - (a) $\mathbf{vk} \in \vec{\mathbf{vk}}^*$.
 - (b) $(\vec{\mathbf{vk}}^*, m^*) \neq (\vec{\mathbf{vk}}, m)$
 - (c) $\text{Verify}(\mathbf{pp}, \text{KAgg}(\vec{\mathbf{vk}}^*), m^*, \text{agg}\sigma^*) = 1$.

In some cases we omit the parameters t and ϵ corresponding to the running time and success probability of adversaries, respectively, and consider polynomial-time adversaries with negligible success probabilities. In addition, when considering schemes whose security is analyzed in the random-oracle model [BR93], we augment all algorithms (including the adversary) with access to the random oracle, introduce an additional parameter $q_{\mathbb{H}}$ that upper bounds the number of direct random-oracle queries issued by the adversary, and consider all probabilities also over the randomness of the oracle.

3.2 Single-Set Unforgeability

Our notion of single-set unforgeability for multi-signature schemes considers adversaries that obtain any polynomial number of signatures of their choice but are restricted to requesting all of these signatures with respect to a single set of signers. In this context, a single set of signers corresponds to a single vector of verification keys, which may be adversarially chosen based on the public parameters of the scheme and on the honestly-generated verification key that the adversary is attacking.

Looking ahead, our construction of a multi-signature scheme that provides single-set security consists of a two-round signing protocol: Each party sends one message to all other parties, receives one message from all other parties, and then produces their output. Following Bellare and Dai [BD21], for formalizing the security of schemes with an interactive signing protocol, adversaries are provided access to a stateful signing oracle that enables to initiate sessions and to continue previously-initiated ones. Specifically, given a multi-signature scheme Π with a two-round signing protocol $\text{Sign} = (\text{Sign}_1, \text{Sign}_2)$, we denote by $\mathcal{O}_{\text{Sign}}(\text{pp}, \text{sk}, \vec{\text{vk}}, \cdot)$ the corresponding stateful signing oracle that is provided to the adversary, where pp denotes the public parameters, sk denotes the signing key corresponding to the honestly-generated verification key vk given as input to the adversary, and $\vec{\text{vk}}$ denotes the vector of verification keys corresponding to the single set of signers chosen by the adversary. For this oracle, an adversary may issue two types of queries:

- Session-initiation queries: On query a message m , the oracle first assigns a unique session identifier sid (e.g., in an incremental manner) and computes $(\text{msg}_1, \text{st}) \leftarrow \text{Sign}_1(\text{pp}, \text{sk}, \vec{\text{vk}}, m)$. Then, it locally stores the pair (sid, st) and outputs $(\text{sid}, \text{msg}_1)$.
- Communication queries: On query a pair $(\text{sid}, \text{msg}_2)$, the oracle first retrieves the stored pair (sid, st) , and then computes and outputs $\text{out} \leftarrow \text{Sign}_2(\text{st}, \text{msg}_2)$. If no stored pair exists for the session identifier sid then the oracle outputs \perp .

Definition 3.3. Let $t = t(\lambda)$, $q_{\text{sign}} = q_{\text{sign}}(\lambda)$, and $\epsilon = \epsilon(\lambda)$ be functions of the security parameter $\lambda \in \mathbb{N}$. A non-interactive multi-signature scheme $\Pi = (\text{Setup}, \text{KG}, \text{KAgg}, \text{Sign}, \text{SAgg}, \text{Verify})$ is *single-set* $(t, q_{\text{sign}}, \epsilon)$ -unforgeable if for any algorithm $A = (A_1, A_2)$ that runs in time at most t and issues at most q_{sign} signing queries, it holds that

$$\text{Adv}_{\Pi, A}^{\text{SSetMS}}(\lambda) \stackrel{\text{def}}{=} \Pr \left[\text{Exp}_{\Pi, A}^{\text{SSetMS}}(\lambda) = 1 \right] \leq \epsilon(\lambda)$$

for all sufficiently large $\lambda \in \mathbb{N}$, where the experiment $\text{Exp}_{\Pi, A}^{\text{SSetMS}}(\lambda)$ is defined as follows:

1. $\text{pp} \leftarrow \text{Setup}(1^\lambda)$.
2. $(\text{sk}, \text{vk}) \leftarrow \text{KG}(\text{pp})$.
3. $(\vec{\text{vk}}, \text{st}) \leftarrow A_1(1^\lambda, \text{pp}, \text{vk})$.
4. $(\vec{\text{vk}}^*, m^*, \text{agg}\sigma^*) \leftarrow A_2^{\mathcal{O}_{\text{Sign}}(\text{pp}, \text{sk}, \vec{\text{vk}}, \cdot)}(\text{st})$.
5. If the following conditions are satisfied then output 1 and otherwise output 0:
 - (a) $\text{vk} \in \vec{\text{vk}}^*$.
 - (b) $\vec{\text{vk}}^* \neq \vec{\text{vk}}$ or A_2 did not query $\mathcal{O}_{\text{Sign}}(\text{pp}, \text{sk}, \vec{\text{vk}}, \cdot)$ with m^* .
 - (c) $\text{Verify}(\text{pp}, \text{KAgg}(\vec{\text{vk}}^*), m^*, \text{agg}\sigma^*) = 1$.

As noted in Section 3.1, when considering schemes whose security is analyzed in the random-oracle model [BR93], we augment all algorithms (including the adversary) with access to the random oracle, introduce an additional parameter q_{H} that upper bounds the number of direct random-oracle queries issued by the adversary, and consider all probabilities also over the randomness of the oracle.

4 One-Time Multi-Signatures via Ring-Homomorphic One-Way Functions

In this section we describe our construction of a one-time multi-signature scheme and prove its security. Our construction is parameterized by the security parameter $\lambda \in \mathbb{N}$ and by an integer $n = n(\lambda)$ determining an upper bound on the size of supported signer sets. The construction relies on the following building blocks:

- A ring-homomorphic one-way function $F = (F.\text{Setup}, F.\text{Eval})$. As formalized in Section 2.1 following Catalano et al. [CFG⁺15], recall that for any $\lambda \in \mathbb{N}$ and for any function index Ind produced by $F.\text{Setup}(1^\lambda)$ it holds that $F.\text{Eval}(\text{Ind}, \cdot) : \mathcal{X}_{\text{Ind}} \rightarrow \mathcal{Y}_{\text{Ind}}$ is a homomorphism for cyclic groups \mathcal{X}_{Ind} and \mathcal{Y}_{Ind} that allows computing linear operations “in the exponent” over a ring $\mathbb{K}_{\text{Ind}} = \mathbb{Z}_q$ for some prime $q = q(\text{Ind}) \geq K(\lambda)$, where q is at most the order of the cyclic groups. For simplifying our notation, for any function index Ind we let $F_{\text{Ind}}(\cdot) = F.\text{Eval}(\text{Ind}, \cdot)$.
- Hash functions H_0 and H_1 that will be modeled, for the security analysis, as random oracles. For any function index Ind of the ring-homomorphic function F we assume that $H_0 : \mathcal{M}_\lambda \times \mathcal{Y}_{\text{Ind}} \times \mathcal{Y}_{\text{Ind}} \rightarrow \mathbb{K}_{\text{Ind}}$ and $H_1 : (\mathcal{Y}_{\text{Ind}} \times \mathcal{Y}_{\text{Ind}})^n \rightarrow \mathbb{K}_{\text{Ind}}^n$, where $\mathcal{M} = \mathcal{M}_\lambda$ is the supported message space of the constructed multi-signature scheme (we can choose, for example, $\mathcal{M}_\lambda = \{0, 1\}^\lambda$). In terms of input lengths, this means that we assume sufficiently long input lengths for H_0 and H_1 . In terms of output lengths, recall that $\mathbb{K}_{\text{Ind}} = \mathbb{Z}_q$ for some prime q , where q is at most the order of the cyclic groups. In practice, this can be realized in a standard manner by producing sufficiently-long outputs using a cryptographic hash function, and then reducing the results modulo q to obtain a statistically-small error.

In what follows we first describe our one-time scheme, denoted $1T$, and then prove its correctness and security. For simplicity and for avoiding the introduction of additional notation, when describing the scheme and proving its security we assume that all signer sets are of size $n = n(\lambda)$, but we note that an upper bound on the size of supported signer sets would suffice. In addition, we note that, in practice our scheme does not essentially require any setup. Specifically, when realizing the ring-homomorphic one-way function via group exponentiation in a prime-order group [CD98, CFG⁺15], the function index Ind consists of the description of the group, which is typically fixed (e.g., standard 256-bit curves such as Secp256k1 or Curve25519).

The scheme $1T = (\text{Setup}, \text{KG}, \text{KAgg}, \text{Sign}, \text{SAgg}, \text{Verify})$

Setup(1^λ). On input 1^λ the setup algorithm samples $\text{Ind} \leftarrow F.\text{Setup}(1^\lambda)$ and outputs $\text{pp} = \text{Ind}$.

KG(pp). On input pp as above, the key-generation algorithm samples $x, r \leftarrow \mathcal{X}_{\text{Ind}}$, computes $X = F_{\text{Ind}}(x)$ and $R = F_{\text{Ind}}(r)$, and then outputs $(\text{sk}, \text{vk}) = ((x, r), (X, R))$.

KAgg($\text{pp}, \vec{\text{vk}}$). On input pp as above and $\vec{\text{vk}} = (\text{vk}_1, \dots, \text{vk}_n)$, where $\text{vk}_i = (X_i, R_i)$ for every $i \in [n]$, the key-aggregation algorithm computes

$$(a_1, \dots, a_n) = H_1(\vec{\text{vk}}) \in \mathbb{K}_{\text{Ind}}^n$$

$$\text{agg}X = \prod_{i=1}^n X_i^{a_i} \in \mathcal{Y}_{\text{Ind}}$$

$$\text{agg}R = \prod_{i=1}^n R_i^{a_i} \in \mathcal{Y}_{\text{Ind}},$$

and outputs $\text{aggvk} = (\text{agg}X, \text{agg}R)$.

Sign $(\mathbf{pp}, \mathbf{sk}, \vec{\mathbf{vk}}, m)$. On input \mathbf{pp} as above, $\mathbf{sk} = (x, r)$, $\vec{\mathbf{vk}}$ and $m \in \mathcal{M}_\lambda$, the signing algorithm is defined as follows:

1. If $(F_{\text{Ind}}(x), F_{\text{Ind}}(r)) \notin \vec{\mathbf{vk}}$ then output \perp .
2. Otherwise, compute $\text{aggvk} = \text{KAgg}(\mathbf{pp}, \vec{\mathbf{vk}})$, and output

$$\sigma = r + H_0(m, \text{aggvk}) \cdot x \in \mathcal{X}_{\text{Ind}}.$$

SAgg $(\mathbf{pp}, \vec{\mathbf{vk}}, \vec{\sigma})$. On input \mathbf{pp} as above, $\vec{\mathbf{vk}}$ and $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$, the signature-aggregation algorithm computes $(a_1, \dots, a_n) = H_1(\vec{\mathbf{vk}}) \in \mathbb{K}_{\text{Ind}}$ and outputs $\text{agg}\sigma = \sum_{i=1}^n a_i \cdot \sigma_i \in \mathcal{X}_{\text{Ind}}$.

Verify $(\mathbf{pp}, \text{aggvk}, m, \text{agg}\sigma)$. On input \mathbf{pp} as above, $\text{aggvk} = (\text{agg}X, \text{agg}R)$, m and $\text{agg}\sigma$, if

$$F_{\text{Ind}}(\text{agg}\sigma) = (\text{agg}X)^{H_0(m, \text{aggvk})} \cdot \text{agg}R$$

then the verification algorithm outputs 1 and otherwise it outputs 0.

Correctness. Fix any $\lambda \in \mathbb{N}$, any public parameters $\mathbf{pp} = \text{Ind}$, and any vector of verification keys $\vec{\mathbf{vk}} = (\mathbf{vk}_1, \dots, \mathbf{vk}_n)$. For every $i \in [n]$ let $\mathbf{vk}_i = (X_i, R_i)$, where $X_i = F_{\text{Ind}}(x_i)$ and $R_i = F_{\text{Ind}}(r_i)$. Then, for any message m it holds that

$$\begin{aligned} F_{\text{Ind}}(\text{agg}\sigma) &= F_{\text{Ind}}\left(\sum_{i=1}^n a_i \cdot r_i + H_0(m, \text{aggvk}) \cdot \sum_{i=1}^n a_i \cdot x_i\right) \\ &= \left(\prod_{i=1}^n F_{\text{Ind}}(x_i)^{a_i}\right)^{H_0(m, \text{aggvk})} \cdot \left(\prod_{i=1}^n F_{\text{Ind}}(r_i)^{a_i}\right) \\ &= (\text{agg}X)^{H_0(m, \text{aggvk})} \cdot \text{agg}R, \end{aligned}$$

where $(a_1, \dots, a_n) = H_1(\vec{\mathbf{vk}})$, and thus the scheme provides perfect correctness.

Security. The following theorem establishes the security of our scheme based on that of the underlying ring-homomorphic one-way function (recall Definition 2.1):

Theorem 4.1. *Let A be a probabilistic polynomial-time algorithm that issues $q_{H_0} = q_{H_0}(\lambda)$ and $q_{H_1} = q_{H_1}(\lambda)$ queries to the oracles H_0 and H_1 , respectively. Then, there exists a probabilistic polynomial-time algorithm I such that for every $\lambda \in \mathbb{N}$ it holds that*

$$\text{Adv}_{\text{IT}, A}^{\text{1TimeMS}}(\lambda) \leq (q_{H_0})^2 \cdot q_{H_1} \cdot n \cdot \left(\sqrt{\text{Adv}_{F, I}^{\text{hom-ow}}(\lambda) + \frac{(q_{H_0})^2 + (q_{H_1})^2}{K(\lambda)}} + \frac{1}{K(\lambda)} \right)$$

Proof of Theorem 4.1. Let A be a probabilistic polynomial-time algorithm that issues a single signing query, and issues q_{H_0} and q_{H_1} queries to the oracles H_0 and H_1 , respectively. Without loss of generality, we assume that A does no query either H_0 and H_1 more than once with the same input. We denote by $(\vec{\mathbf{vk}}, m)$ the signing query issued by A , and by $(\vec{\mathbf{vk}}^*, m^*, \text{agg}\sigma^*)$ the output provided by A . By increasing the number of A 's queries to H_0 by at most 2, we assume that A always queries H_0 with (m, aggvk) and (m^*, aggvk^*) , where $\text{aggvk} = \text{KAgg}(\mathbf{pp}, \vec{\mathbf{vk}})$ and $\text{aggvk}^* = \text{KAgg}(\mathbf{pp}, \vec{\mathbf{vk}}^*)$. Similarly, by increasing the number of A 's queries to H_1 by at most 2, we assume that A always queries H_1 with $\vec{\mathbf{vk}}$ and $\vec{\mathbf{vk}}^*$.

We construct a probabilistic polynomial-time inversion algorithm I using Lemma 2.2. Concretely, we define an algorithm B which simulates the experiment $\text{Exp}_{1T,A}^{\text{TimeMS}}$ to A and enforces some additional checks. Then, on input (Ind, X) , the algorithm I invokes the forking algorithm F_B for B , guaranteed by Lemma 2.2, to obtain two correlated forgeries which can be used to compute $x' \in \mathcal{X}_{\text{Ind}}$ and $0 \neq d \in \mathbb{K}_{\text{Ind}}$ such that $F_{\text{Ind}}(x') = X^d$.

The algorithm $B(\text{pp}, X, \vec{h})$

On input $\text{pp} = \text{Ind}$, $X = F_{\text{Ind}}(x)$, and $\vec{h} = (h^{(1)}, \dots, h^{(q_{H_1})})$, the algorithm B is defined as follows:

1. Sample internal randomness

$$\rho = \left(i_0^*, i_1^*, k^*, \sigma, \left\{ c^{(i)} \right\}_{i \in [q_{H_0}]}, \left\{ (a_1^{(i)}, \dots, a_n^{(i)}) \right\}_{i \in [q_{H_1}]} \right),$$

where:

- $i_0^*, i_1^* \leftarrow [q_{H_0}]$. Looking ahead, i_0^* and i_1^* are B 's guesses for the indices of A 's H_0 -queries (m, aggvk) and (m^*, aggvk^*) , respectively.
 - $k^* \leftarrow [n]$. Looking ahead, k^* is B 's guess for the coordinate in which A will position vk within $\vec{\text{vk}}^*$.
 - $\sigma \leftarrow \mathcal{X}_{\text{pp}}$. Looking ahead, σ is B 's response to A 's signing query.
 - $c^{(i)} \leftarrow \mathbb{K}_{\text{Ind}}$ for every $i \in [q_{H_0}]$. Looking ahead, these values are B 's responses to A 's H_0 -queries.
 - $(a_1^{(j)}, \dots, a_n^{(j)}) \leftarrow \mathbb{K}_{\text{Ind}}^n$ for every $j \in [q_{H_1}]$. Looking ahead, these are B 's responses to A 's H_1 -queries (except for the k^* -th coordinate $a_{k^*}^{(j)}$ that will be replaced).
2. Set $R = F_{\text{Ind}}(\sigma) \cdot X^{-c^{(i_0^*)}}$ (assuming that i_0^* is guessed correctly, this defines $c^{(i_0^*)} = H_0(m, \text{aggvk})$), and for every $j \in [q_{H_1}]$ set $a_{k^*}^{(j)} = h^{(j)}$.
3. Invoke A on input (pp, vk) , where $\text{vk} = (X, R)$, and respond to A 's oracle queries as follows:
- When A issues a signing query return σ .
 - For every $i \in [q_{H_0}]$, when A issues their i th H_0 -query return $c^{(i)}$.
 - For every $j \in [q_{H_1}]$, when A issues their j th H_1 -query return $(a_1^{(j)}, \dots, a_n^{(j)})$.
4. Given A 's output $(\vec{\text{vk}}^*, m^*, \text{agg}\sigma^*)$, if all of the following conditions are satisfied then output $(j^*, \text{agg}\sigma^*)$, where $j^* \in [q_{H_1}]$ is the index of the H_1 -query $\vec{\text{vk}}^*$, and otherwise output \perp :
- (a) Verify $(\text{pp}, \vec{\text{vk}}^*, m^*, \text{agg}\sigma^*) = 1$, $\text{vk} \in \vec{\text{vk}}^*$, and $(\vec{\text{vk}}^*, m^*) \neq (\vec{\text{vk}}, m)$ (i.e., A won the simulated experiment).
 - (b) A 's i_0^* -th H_0 -query was (m, aggvk) (i.e., B guessed i_0^* correctly).
 - (c) A 's i_1^* -th H_0 -query was (m^*, aggvk^*) (i.e., B guessed i_1^* correctly).
 - (d) $\vec{\text{vk}}^* = (\text{vk}_1^*, \dots, \text{vk}_n^*)$ where $\text{vk}_{k^*}^* = \text{vk}$ (i.e., B guessed k^* correctly).

Claim 4.2. For every $\lambda \in \mathbb{N}$ it holds that

$$\Pr \left[B(\text{pp}, F_{\text{Ind}}(x), \vec{h}) \neq \perp \right] \geq \frac{1}{(q_{H_0})^2 \cdot n} \cdot \text{Adv}_{1T,A}^{\text{TimeMS}}(\lambda),$$

where $\text{pp} = \text{Ind} \leftarrow F.\text{Setup}(1^\lambda)$, $x \leftarrow \mathcal{X}_{\text{Ind}}$ and $\vec{h} \leftarrow \mathbb{K}_{\text{Ind}}^{q_{H_1}}$.

Proof of Claim 4.2. Let win denote the event in which A wins the simulated experiment, and let hit denote the event in which the index i_0^* was guessed correctly by B . The choices of i_1^* and k^* are completely independent of A 's execution, and therefore can be sampled by B at the end. Hence,

$$\begin{aligned} \Pr \left[B \left(\text{pp}, F_{\text{Ind}}(x), \vec{h} \right) \neq \perp \right] &= \frac{1}{q_{\text{H}_0} \cdot n} \cdot \Pr [\text{win} | \text{hit}] \cdot \Pr [\text{hit}] \\ &= \frac{1}{(q_{\text{H}_0})^2 \cdot n} \cdot \Pr [\text{win} | \text{hit}] \\ &= \frac{1}{(q_{\text{H}_0})^2 \cdot n} \cdot \text{Adv}_{\text{IT}, A}^{\text{TimeMS}}(\lambda), \end{aligned} \quad (4.1)$$

where Eq. (4.1) also follows from the fact that B outputs \perp with probability 1 conditioned on $\overline{\text{hit}}$. Since conditioned on hit , B perfectly simulates the one-time unforgeability experiment to A , the claim follows. \blacksquare

Claim 4.2 directly enables us to lower bound the probability that $F_B(\text{Ind}, X) \neq \perp$ via the forking lemma (see Lemma 2.2):

Corollary 4.3. *For every $\lambda \in \mathbb{N}$ it holds that*

$$\Pr [F_B(\text{Ind}, F_{\text{Ind}}(x)) \neq \perp] \geq \frac{\text{Adv}_{\text{IT}, A}^{\text{TimeMS}}(\lambda)}{(q_{\text{H}_0})^2 \cdot n} \cdot \left(\frac{\text{Adv}_{\text{IT}, A}^{\text{TimeMS}}(\lambda)}{(q_{\text{H}_0})^2 \cdot q_{\text{H}_1} \cdot n} - \frac{1}{K(\lambda)} \right),$$

where $\text{Ind} \leftarrow \text{F.Setup}(1^\lambda)$ and $x \leftarrow \mathcal{X}_{\text{Ind}}$.

Having described the algorithms B and F_B , and analyzed their success probability, we can now describe the inversion algorithm I . On input (Ind, X) , the algorithm I invokes the algorithm F_B provided by the forking lemma (see Section 2.2) on the input (pp, X) , where $\text{pp} = \text{Ind}$. If the algorithm F_B returns \perp then I fails (e.g., I outputs some fixed values $x' \in \mathcal{X}_\lambda$ and $d \in \mathbb{K}_{\text{Ind}}$). Otherwise, if the algorithm F_B returns a triplet $(j^*, \text{agg}\sigma^*, \text{agg}\sigma^{**})$, then denote by ρ the random string sampled by F_B for running B , where

$$\rho = \left(i_0^*, i_1^*, k^*, \sigma, \left\{ c^{(i)} \right\}_{i \in [q_{\text{H}_0}]}, \left\{ \left(a_1^{(j)}, \dots, a_n^{(j)} \right) \right\}_{j \in [q_{\text{H}_1}]} \right),$$

and denote by $h^{(1)}, \dots, h^{(q_{\text{H}_1})}$ the sequence of oracle responses sampled by F_B for B 's first execution, and by $\hat{h}^{(j^*)}, \dots, \hat{h}^{(q_{\text{H}_1})}$ the additional oracle responses sampled by F_B for B 's second execution. That is, using this notation, the algorithm F_B first invoked B on the input $(\text{pp}, X, h^{(1)}, \dots, h^{(q_{\text{H}_1})})$ to obtain an output $(\vec{\text{vk}}^*, m^*, \text{agg}\sigma^*)$ from A and an output $(j^*, \text{agg}\sigma^*)$ from B , and then invoked B on the input $(\text{pp}, X, h^{(1)}, \dots, h^{(j^*-1)}, \hat{h}^{(j^*)}, \dots, \hat{h}^{(q_{\text{H}_1})})$ to obtain an output $(\vec{\text{vk}}^{**}, m^{**}, \text{agg}\sigma^{**})$ from A and an output $(j^*, \text{agg}\sigma^{**})$ from B . The inversion algorithm I then computes and outputs

$$\begin{aligned} x' &= \text{agg}\sigma^* - \text{agg}\sigma^{**} - \left(h^{(j^*)} - \hat{h}^{(j^*)} \right) \cdot \sigma \in \mathcal{X}_{\text{Ind}} \\ d &= \left(c^{(i_1^*)} - c^{(i_0^*)} \right) \cdot \left(h^{(j^*)} - \hat{h}^{(j^*)} \right) \in \mathbb{K}_{\text{Ind}}. \end{aligned}$$

We now show that as long as the algorithm F_B does not return \perp , then $F_{\text{Ind}}(x') = X^d$ and thus the inversion algorithm I is successful. Note that both invocations of the algorithm B by the algorithm F_B use the random string ρ , and therefore B uses the same values $i_0^*, i_1^*, k^*, \sigma, \{c^{(i)}\}_{i \in [q_{\text{H}_0}]}$ and $\left\{ \left(a_1^{(j)}, \dots, a_n^{(j)} \right) \right\}_{j \in [q_{\text{H}_1}]}$ for these two executions. In particular it holds that:

1. In both executions of B the vector of verification keys \vec{vk}^* relative to which A produces the forgery was provided as input to A 's j^* -th H_1 -query (this is the query to which A is rewinded by B). Therefore, since the executions are identical up to the response to the j^* -th H_1 -query, and since in both executions the same index k^* is used for guessing the position of vk^* and vk^{**} in \vec{vk}^* , it holds that $vk^* = vk^{**}$. We let

$$vk^* = vk^{**} = ((X_1, R_1), \dots, (X_n, R_n)).$$

2. In both executions of B the verification key $vk = (X, R)$ was positioned by A in coordinate $k^* \in [n]$ of vk^* and vk^{**} . Therefore, $(X_{k^*}, R_{k^*}) = (X, R)$.
3. In both executions of B , the value returned as the response to A 's i_1^* -th query to H_0 was $c^{(i_1^*)}$. Therefore, in the first execution $c^{(i_1^*)} = H_0(m^*, \text{agg}vk^*)$ and in the second execution $c^{(i_1^*)} = H_0(m^{**}, \text{agg}vk^{**})$.

Thus, in B 's first execution, we have that

$$\begin{aligned} H_0(m^*, \text{agg}vk^*) &= c^{(i_1^*)} \\ H_1(\vec{vk}^*) &= (a_1^{(j^*)}, \dots, a_{k^*-1}^{(j^*)}, h^{(j^*)}, a_{k^*+1}^{(j^*)}, a_n^{(j^*)}), \end{aligned}$$

and given that A produced a valid forgery it holds that

$$F_{\text{Ind}}(\text{agg}\sigma^*) = \left(X_{k^*}^{h^{(j^*)}} \cdot \prod_{i \in [n] \setminus \{k^*\}} X_i^{a_i^{(j^*)}} \right)^{c^{(i_1^*)}} \cdot \left(R_{k^*}^{h^{(j^*)}} \cdot \prod_{i \in [n] \setminus \{k^*\}} R_i^{a_i^{(j^*)}} \right). \quad (4.2)$$

Similarly, in B 's second execution, we have that

$$\begin{aligned} H_0(m^{**}, \text{agg}vk^{**}) &= c^{(i_1^*)} \\ H_1(\vec{vk}^{**}) &= (a_1^{(j^*)}, \dots, a_{k^*-1}^{(j^*)}, \hat{h}^{(j^*)}, a_{k^*+1}^{(j^*)}, a_n^{(j^*)}), \end{aligned}$$

and given that A produced a valid forgery it holds that

$$F_{\text{Ind}}(\text{agg}\sigma^{**}) = \left(X_{k^*}^{\hat{h}^{(j^*)}} \cdot \prod_{i \in [n] \setminus \{k^*\}} X_i^{a_i^{(j^*)}} \right)^{c^{(i_1^*)}} \cdot \left(R_{k^*}^{\hat{h}^{(j^*)}} \cdot \prod_{i \in [n] \setminus \{k^*\}} R_i^{a_i^{(j^*)}} \right). \quad (4.3)$$

Therefore, using the facts that $X_{k^*} = X$ and $R_{k^*} = R = F_{\text{Ind}}(\sigma) \cdot X^{-c^{(i_0^*)}}$, Equations (4.2) and (4.3) imply

$$\begin{aligned} F_{\text{Ind}}(\text{agg}\sigma^* - \text{agg}\sigma^{**}) &= X_{k^*}^{c^{(i_1^*)} \cdot (h^{(j^*)} - \hat{h}^{(j^*)})} \cdot R_{k^*}^{h^{(j^*)} - \hat{h}^{(j^*)}} \\ &= X^{c^{(i_1^*)} \cdot (h^{(j^*)} - \hat{h}^{(j^*)})} \cdot \left(F_{\text{Ind}}(\sigma) \cdot X^{-c^{(i_0^*)}} \right)^{h^{(j^*)} - \hat{h}^{(j^*)}} \\ &= X^{(c^{(i_1^*)} - c^{(i_0^*)}) \cdot (h^{(j^*)} - \hat{h}^{(j^*)})} \cdot F_{\text{Ind}}\left(\left(h^{(j^*)} - \hat{h}^{(j^*)}\right) \cdot \sigma\right) \end{aligned}$$

which, for the above setting of x' and d , implies that

$$\begin{aligned} F_{\text{Ind}}(x') &= F_{\text{Ind}}\left(\text{agg}\sigma^* - \text{agg}\sigma^{**} - \left(h^{(j^*)} - \hat{h}^{(j^*)}\right) \cdot \sigma\right) \\ &= X^{(c^{(i_1^*)} - c^{(i_0^*)}) \cdot (h^{(j^*)} - \hat{h}^{(j^*)})} \\ &= X^d. \end{aligned}$$

The following claim concludes the proof by bounding the probability that $d = 0_{\mathbb{K}_{\text{Ind}}}$. This shows, in particular, that the scheme's key-aggregation procedure is collision resistant as a function of the vector of verification keys that it aggregates.

Claim 4.4. *For every $\lambda \in \mathbb{N}$ it holds that*

$$\Pr \left[\begin{array}{l} I(\text{Ind}, X) = (x', d) \\ \text{s.t. } d = 0_{\mathbb{K}_{\text{Ind}}} \end{array} \wedge F_B(\text{pp}, X) \neq \perp \right] \leq \frac{(q_{\text{H}_0})^2 + (q_{\text{H}_1})^2}{K(\lambda)},$$

where $\text{pp} = \text{Ind} \leftarrow \text{F.Setup}(1^\lambda)$, $x \leftarrow \mathcal{X}_{\text{Ind}}$, $X = \text{F}_{\text{Ind}}(x)$ and the probability is additionally taken over the internal randomness of the algorithms I and B .

Proof of Claim 4.4. Conditioned on the event $F_B(\text{pp}, X) \neq \perp$, the algorithm I outputs $d = (c^{(i_1^*)} - c^{(i_0^*)}) \cdot (h^{(j^*)} - \hat{h}^{(j^*)}) \in \mathbb{K}_{\text{Ind}}$, where $c^{(i_1^*)}, c^{(i_0^*)}, h^{(j^*)}, \hat{h}^{(j^*)} \in \mathbb{K}_{\text{Ind}}$ and $\mathbb{K}_{\text{Ind}} = \mathbb{Z}_q$ for some prime $q = q(\text{Ind}) \geq K(\lambda)$. Recall, in addition, that according to the forking lemma, if $h^{(j^*)} = \hat{h}^{(j^*)}$ then F_B outputs \perp . Therefore, conditioned on the event $F_B(\text{pp}, X) \neq \perp$, it holds that $h^{(j^*)} \neq \hat{h}^{(j^*)}$, and we are left to consider the event $c^{(i_1^*)} = c^{(i_0^*)}$. Note that these two values are defined already for the first invocation of B by F_B , and conditioned on $B(\text{pp}, X, \vec{h}) \neq \perp$ (which follows from $F_B(\text{pp}, X) \neq \perp$) it holds that

$$\begin{aligned} c^{(i_0^*)} &= \text{H}_0(m, \text{aggvk}) \\ c^{(i_1^*)} &= \text{H}_0(m^*, \text{aggvk}^*). \end{aligned}$$

Now, there are two disjoint events to consider, denoted \mathcal{E} and $\bar{\mathcal{E}}$:

Event \mathcal{E} : $m \neq m^*$ or $\text{aggvk} \neq \text{aggvk}^*$. Here $(m, \text{aggvk}) \neq (m^*, \text{aggvk}^*)$, and therefore $c^{(i_0^*)} = c^{(i_1^*)}$ implies that the algorithm A found a non-trivial collision for the random oracle H_0 during B 's first invocation. Such a collision may be found with probability at most $(q_{\text{H}_0})^2/|\mathbb{K}_{\text{Ind}}| \leq (q_{\text{H}_0})^2/K(\lambda)$, and therefore

$$\Pr \left[\begin{array}{l} I(\text{Ind}, X) = (x', d) \\ \text{s.t. } d = 0_{\mathbb{K}_{\text{Ind}}} \end{array} \wedge F_B(\text{pp}, X) \neq \perp \wedge \mathcal{E} \right] \leq \frac{(q_{\text{H}_0})^2}{K(\lambda)}. \quad (4.4)$$

Event $\bar{\mathcal{E}}$: $m = m^*$ and $\text{aggvk} = \text{aggvk}^*$. Note that by the definition of B , if $F_B(\text{pp}, X) \neq \perp$ then $(m, \vec{\text{vk}}) \neq (m^*, \vec{\text{vk}}^*)$ and therefore the event $\bar{\mathcal{E}}$ implies that the algorithm A produced $\vec{\text{vk}}$ and $\vec{\text{vk}}^*$ such that $\vec{\text{vk}} \neq \vec{\text{vk}}^*$ and $\text{aggvk} = \text{aggvk}^*$. Among the q_{H_1} queries issued by A to the random oracle H_1 , there are at most $(q_{\text{H}_0})^2$ potential choices for such $\vec{\text{vk}} \neq \vec{\text{vk}}^*$, and we now show that for each such choice it holds that $\text{aggvk} = \text{aggvk}^*$ with probability at most $1/K(\lambda)$.

Specifically, let

$$\vec{\text{vk}} = ((X_1, R_1), \dots, (X_n, R_n))$$

and

$$\vec{\text{vk}}^* = ((X_1^*, R_1^*), \dots, (X_n^*, R_n^*)),$$

and assume without loss of generality that A issues the H_1 -query $\vec{\text{vk}}$ before the H_1 -query $\vec{\text{vk}}^*$. Then, if $(X_1, \dots, X_n) \neq (X_1^*, \dots, X_n^*)$, then by the time A issues that H_1 query $\vec{\text{vk}}^*$ the value $\text{agg}X = \prod_{i=1}^n X_i^{a_i}$ is already fixed, where $(a_1, \dots, a_n) = \text{H}_1(\vec{\text{vk}})$. Therefore, the probability over the choice of $(a_1^*, \dots, a_n^*) = \text{H}_1(\vec{\text{vk}}^*)$ that $\text{agg}X^* = \prod_{i=1}^n (X_i^*)^{a_i^*} = \text{agg}X$ is at most

$1/|\mathbb{K}_{\text{Ind}}| \leq 1/K(\lambda)$. An identical argument holds if $(R_1, \dots, R_n) \neq (R_1^*, \dots, R_n^*)$, and overall we obtain

$$\Pr \left[\begin{array}{l} I(\text{Ind}, X) = (x', d) \\ \text{s.t. } d = 0_{\mathbb{K}_{\text{Ind}}} \end{array} \wedge F_B(\text{pp}, X) \neq \perp \wedge \bar{\mathcal{E}} \right] \leq \frac{(q_{H_1})^2}{K(\lambda)}. \quad (4.5)$$

The claim now follows by combining Equations (4.4) and (4.5). \blacksquare

Equipped with Corollary 4.3 and Claim 4.4, we have

$$\begin{aligned} \Pr \left[\begin{array}{l} I(\text{Ind}, X) = (x', d) \text{ s.t.} \\ \mathbf{F}_{\text{Ind}}(x') \neq X^d \text{ or } d = 0_{\mathbb{K}_{\text{Ind}}} \end{array} \right] \\ \leq \Pr [F_B(\text{pp}, X) = \perp] \\ + \Pr \left[\begin{array}{l} I(\text{Ind}, X) = (x', d) \\ \text{s.t. } d = 0_{\mathbb{K}_{\text{Ind}}} \wedge F_B(\text{pp}, X) \neq \perp \end{array} \right] \\ \leq 1 - \frac{\text{Adv}_{1\text{T},A}^{\text{TimeMS}}(\lambda)}{(q_{H_0})^2 \cdot n} \cdot \left(\frac{\text{Adv}_{1\text{T},A}^{\text{TimeMS}}(\lambda)}{(q_{H_0})^2 \cdot q_{H_1} \cdot n} - \frac{1}{K(\lambda)} \right) \\ + \frac{(q_{H_0})^2 + (q_{H_1})^2}{K(\lambda)}. \end{aligned}$$

Thus, for any $\lambda \in \mathbb{N}$ it holds that

$$\begin{aligned} \text{Adv}_{F,I}^{\text{hom-ow}}(\lambda) &= \Pr \left[\begin{array}{l} I(\text{Ind}, X) = (x', d) \text{ s.t.} \\ \mathbf{F}_{\text{Ind}}(x') = X^d \text{ and } d \neq 0_{\mathbb{K}_{\text{Ind}}} \end{array} \right] \\ &\geq \left(\frac{\text{Adv}_{1\text{T},A}^{\text{TimeMS}}(\lambda)}{(q_{H_0})^2 \cdot q_{H_1} \cdot n} - \frac{1}{|K(\lambda)|} \right)^2 - \frac{(q_{H_0})^2 + (q_{H_1})^2}{K(\lambda)}, \end{aligned}$$

and therefore

$$\text{Adv}_{1\text{T},A}^{\text{TimeMS}}(\lambda) \leq (q_{H_0})^2 \cdot q_{H_1} \cdot n \cdot \left(\sqrt{\text{Adv}_{F,I}^{\text{hom-ow}}(\lambda) + \frac{(q_{H_0})^2 + (q_{H_1})^2}{K(\lambda)}} + \frac{1}{K(\lambda)} \right).$$

This settles the proof of Theorem 4.1. \blacksquare

5 From One-Time to Single-Set Multi-Signatures

In this section we show that any one-time multi-signature scheme with non-interactive signing can be transformed into a multi-signature scheme that satisfies our notion of single-set unforgeability with a two-round signing protocol. Our construction relies on the following building blocks:

- A multi-signature scheme $1\text{T} = (1\text{T.Setup}, 1\text{T.KG}, 1\text{T.KAgg}, 1\text{T.Sign}, 1\text{T.SAgg}, 1\text{T.Verify})$ that is one-time unforgeable with non-interactive signing over a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$. For simplicity we assume that $\mathcal{M}_\lambda = \{0, 1\}^\ell$ for some polynomial $\ell = \ell(\lambda)$, and that ℓ is sufficiently large for enabling the scheme to sign, for example, the concatenation of two verification keys produced by its key-generation algorithm (otherwise, we can additionally rely on a collision-resistant hash function in the standard “hash-then-sign” manner). In addition, since the scheme has non-interactive signing, we assume without loss of generality that its signing algorithm is deterministic.

- A pseudorandom function $\text{PRF} = (\text{PRF.KG}, \text{PRF.Eval})$, where for each $\lambda \in \mathbb{N}$ and for each key k produced by $\text{PRF.KG}(1^\lambda)$ it holds that $\text{PRF.Eval}(k, \cdot) : \{0, 1\}^{\leq \ell(\lambda)} \rightarrow \{0, 1\}^{\ell'(\lambda)}$, and $\ell'(\lambda)$ is the length of the internal random string sampled by the key-generation algorithm 1T.KG .
- A collision-resistant hash function $\text{CRH} = (\text{CRH.KG}, \text{CRH.Eval})$. As noted below, our usage of a collision-resistant hash function is in fact not required, and is done for providing a more direct proof of security.
- A non-interactive statistically-binding commitment scheme $\text{COM} = (\text{COM.Setup}, \text{COM.Commit}, \text{COM.Verify})$.
- A non-interactive simulation-sound zero-knowledge proof system $\text{ZK} = (\text{ZK.Setup}, \text{ZK.P}, \text{ZK.V}, \text{ZK.Sim}_1, \text{ZK.Sim}_2)$ for the language $\mathcal{L} = \{\mathcal{L}_\lambda\}_{\lambda \in \mathbb{N}}$, where

$$\mathcal{L}_\lambda = \left\{ \left(\begin{array}{l} \text{pp}_{\text{1T}}, \text{crs}_{\text{com}}, \\ \text{com}, x, \text{vk} \end{array} \right) : \begin{array}{l} \exists (\text{decom}, k_{\text{PRF}}, \text{sk}) \text{ s.t.} \\ \text{COM.Verify}(\text{crs}_{\text{com}}, \text{com}, \text{decom}, k_{\text{PRF}}) = 1 \text{ and} \\ (\text{sk}, \text{vk}) = \text{1T.KG}(\text{pp}_{\text{1T}}; \text{PRF.Eval}(k_{\text{PRF}}, x)) \end{array} \right\}.$$

Note that even if the security of the scheme 1T is proved in the random-oracle model, then as long as the key-generation algorithm 1T.KG does not access the random oracle (as with our scheme in Section 4), then \mathcal{L} is an NP-language (assuming, of course, that COM is a standard-model commitment scheme, and that PRF is a standard-model pseudorandom function [Nao91, Gol01]). In this case, the single-set security of our construction is proved in the random-oracle model.

For presenting our scheme, we denote by ε the empty string, we denote by $m|_i$ the leftmost i bits of a binary string m (where $m|_0 = \varepsilon$), and we denote by $m|_i b$ the binary string obtained by concatenating the strings $m|_i$ and b .

The scheme $\Pi = (\text{Setup}, \text{KG}, \text{KAgg}, \text{Sign}, \text{SAgg}, \text{Verify})$

Setup(1^λ). On input 1^λ the setup algorithm computes

$$\begin{aligned} \text{pp}_{\text{1T}} &\leftarrow \text{1T.Setup}(1^\lambda) \\ \text{crs}_{\text{ZK}} &\leftarrow \text{ZK.Setup}(1^\lambda) \\ \text{crs}_{\text{COM}} &\leftarrow \text{COM.Setup}(1^\lambda) \\ \text{k}_{\text{CRH}} &\leftarrow \text{CRH.KG}(1^\lambda), \end{aligned}$$

and returns $\text{pp} = (\text{pp}_{\text{1T}}, \text{crs}_{\text{ZK}}, \text{crs}_{\text{COM}}, \text{k}_{\text{CRH}})$.

KG(pp). On input pp as above, the key-generation algorithm computes

$$\begin{aligned} (\text{sk}_\varepsilon, \text{vk}_\varepsilon) &\leftarrow \text{1T.KG}(\text{pp}_{\text{1T}}) \\ \text{k}_{\text{PRF}} &\leftarrow \text{PRF.KG}(1^\lambda) \\ (\text{com}, \text{decom}) &\leftarrow \text{COM.Commit}(\text{crs}_{\text{COM}}, \text{k}_{\text{PRF}}), \end{aligned}$$

and returns $\text{sk} = (\text{sk}_\varepsilon, \text{vk}_\varepsilon, \text{k}_{\text{PRF}}, \text{com}, \text{decom})$ and $\text{vk} = (\text{vk}_\varepsilon, \text{com})$.

KAgg($\text{pp}, \vec{\text{vk}}$). On input pp as above and $\vec{\text{vk}} = \left((\text{vk}_\varepsilon^{(1)}, \text{com}^{(1)}), \dots, (\text{vk}_\varepsilon^{(n)}, \text{com}^{(n)}) \right)$ the key-aggregation

algorithm computes

$$\begin{aligned}\text{aggvk}_\varepsilon &= \text{1T.KAgg}\left(\text{pp}_{\text{1T}}, \left(\text{vk}_\varepsilon^{(1)}, \dots, \text{vk}_\varepsilon^{(n)}\right)\right) \\ \text{tag}_{\vec{\text{vk}}} &= \text{CRH.Eval}\left(\text{k}_{\text{CRH}}, \vec{\text{vk}}\right),\end{aligned}$$

and returns $\text{aggvk} = (\text{aggvk}_\varepsilon, \text{tag}_{\vec{\text{vk}}})$.

Sign $(\text{pp}, \text{sk}, \vec{\text{vk}}, m)$. On input pp as above, $\text{sk} = (\text{sk}_\varepsilon, \text{vk}_\varepsilon, \text{k}, \text{com}, \text{decom})$, $\vec{\text{vk}}$ and $m \in \{0, 1\}^\ell$, the signing algorithm proceeds as follows:

1. Let $\vec{\text{vk}} = \left(\left(\text{vk}_\varepsilon^{(1)}, \text{com}^{(1)}\right), \dots, \left(\text{vk}_\varepsilon^{(n)}, \text{com}^{(n)}\right)\right)$. If there is no index $i \in [n]$ for which $(\text{vk}_\varepsilon, \text{com}) = (\text{vk}_\varepsilon^{(i)}, \text{com}^{(i)})$ or there is more than one such index, then abort. Otherwise, denote by $k^* \in [n]$ the unique such index.
2. For any $j \in \{0, \dots, \ell - 1\}$ and $b \in \{0, 1\}$ compute

$$\begin{aligned}\left(\text{sk}_{m|_j b}^{(k^*)}, \text{vk}_{m|_j b}^{(k^*)}\right) &= \text{1T.KG}\left(\text{pp}_{\text{1T}}; \text{PRF.Eval}(\text{k}, m|_j b)\right) \\ \pi_{m|_j b}^{(k^*)} &= \text{ZK.P}\left(\text{crs}_{\text{ZK}}, \left(\text{pp}_{\text{1T}}, \text{crs}_{\text{com}}, \text{com}, m|_j b, \text{vk}_{m|_j b}^{(k^*)}\right), \left(\text{decom}, \text{k}, \text{sk}_{m|_j b}^{(k^*)}\right)\right)\end{aligned}$$

and send $\left\{\left(\text{vk}_{m|_j b}^{(k^*)}, \pi_{m|_j b}^{(k^*)}\right)\right\}_{j \in \{0, \dots, \ell - 1\}, b \in \{0, 1\}}$ to all other parties.

3. Upon receiving $\left\{\left(\text{vk}_{m|_j b}^{(i)}, \pi_{m|_j b}^{(i)}\right)\right\}_{i \in [n] \setminus \{k^*\}, j \in \{0, \dots, \ell - 1\}, b \in \{0, 1\}}$ from all other parties, if there exists $(i, j, b) \in ([n] \setminus \{k^*\}) \times \{0, \dots, \ell - 1\} \times \{0, 1\}$ for which $\text{vk}_{m|_j b}^{(i)} = \text{vk}_{m|_j b}^{(k^*)}$ or

$$\text{ZK.V}\left(\text{crs}_{\text{ZK}}, \left(\text{pp}_{\text{1T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, m|_j b, \text{vk}_{m|_j b}^{(i)}\right), \pi_{m|_j b}^{(i)}\right) = 0$$

then abort. Otherwise, compute

$$\begin{aligned}\vec{\text{vk}}_{m|_j b} &= \left(\text{vk}_{m|_j b}^{(1)}, \dots, \text{vk}_{m|_j b}^{(n)}\right) \text{ for all } j \in \{0, \dots, \ell - 1\} \text{ and } b \in \{0, 1\} \\ \text{aggvk}_{m|_j b} &= \text{1T.KAgg}\left(\text{pp}_{\text{1T}}, \vec{\text{vk}}_{m|_j b}\right) \text{ for all } j \in \{0, \dots, \ell - 1\} \text{ and } b \in \{0, 1\} \\ \sigma_{m|_j}^{(k^*)} &= \text{1T.Sign}\left(\text{pp}_{\text{1T}}, \text{sk}_{m|_j}^{(k^*)}, \vec{\text{vk}}_{m|_j}, \left(\text{aggvk}_{m|_j 0}, \text{aggvk}_{m|_j 1}\right)\right) \text{ for all } j \in \{0, \dots, \ell - 1\} \\ \text{tag}_{\vec{\text{vk}}} &= \text{CRH.Eval}\left(\text{k}_{\text{CRH}}, \vec{\text{vk}}\right) \\ \sigma_m^{(k^*)} &= \text{1T.Sign}\left(\text{pp}_{\text{1T}}, \text{sk}_m^{(k^*)}, \vec{\text{vk}}_m, (m, \text{tag}_{\vec{\text{vk}}})\right)\end{aligned}$$

and output

$$\sigma^{(k^*)} = \left(\left\{\left(\sigma_{m|_j}^{(k^*)}, \text{vk}_{m|_j 0}^{(k^*)}, \text{vk}_{m|_j 1}^{(k^*)}\right)\right\}_{j \in \{0, \dots, \ell - 1\}}, \sigma_m^{(k^*)}\right).$$

SAgg $(\text{pp}, \vec{\text{vk}}, \vec{\sigma})$. On input pp as above, $\vec{\text{vk}}$ and $\vec{\sigma} = (\sigma^{(1)}, \dots, \sigma^{(n)})$ where

$$\sigma^{(i)} = \left(\left\{\left(\sigma_{m|_j}^{(i)}, \text{vk}_{m|_j 0}^{(i)}, \text{vk}_{m|_j 1}^{(i)}\right)\right\}_{j \in \{0, \dots, \ell - 1\}}, \sigma_m^{(i)}\right)$$

for every $i \in [n]$, the signature-aggregation algorithm computes

$$\begin{aligned} \text{aggvk}_{m|j b} &= \text{1T.KAagg} \left(\text{pp}_{\text{1T}}, \left(\text{vk}_{m|j b}^{(1)}, \dots, \text{vk}_{m|j b}^{(n)} \right) \right) \text{ for all } j \in \{0, \dots, \ell - 1\} \text{ and } b \in \{0, 1\} \\ \text{agg}\sigma_{m|j} &= \text{1T.SAagg} \left(\text{pp}_{\text{1T}}, \vec{\text{vk}}_{m|j}, \left(\sigma_{m|j}^{(1)}, \dots, \sigma_{m|j}^{(n)} \right) \right) \text{ for all } j \in \{0, \dots, \ell\}, \end{aligned}$$

and outputs

$$\text{agg}\sigma = \left(\left\{ \left(\text{agg}\sigma_{m|j}, \text{aggvk}_{m|j 0}, \text{aggvk}_{m|j 1} \right) \right\}_{j \in \{0, \dots, \ell - 1\}}, \text{agg}\sigma_m \right).$$

Verify($\text{pp}, \text{aggvk}, m, \text{agg}\sigma$). On input pp as above, $\text{aggvk} = (\text{aggvk}_\varepsilon, \text{tag}_{\vec{\text{vk}}})$, $m \in \{0, 1\}^\ell$ and $\text{agg}\sigma$, where

$$\text{agg}\sigma = \left(\left\{ \left(\text{agg}\sigma_{m|j}, \text{aggvk}_{m|j 0}, \text{aggvk}_{m|j 1} \right) \right\}_{j \in \{0, \dots, \ell - 1\}}, \text{agg}\sigma_m \right),$$

the verification algorithm outputs 1 if and only if the following two requirements are satisfied:

1. $\text{1T.Verify} \left(\text{pp}_{\text{1T}}, \text{aggvk}_{m|j}, \left(\text{aggvk}_{m|j 0}, \text{aggvk}_{m|j 1} \right), \text{agg}\sigma_{m|j} \right) = 1$ for every $j \in \{0, \dots, \ell - 1\}$.
2. $\text{1T.Verify} \left(\text{pp}_{\text{1T}}, \text{aggvk}_m, (m, \text{tag}_{\vec{\text{vk}}}), \text{agg}\sigma_m \right) = 1$.

The following theorem captures the security of the scheme Π based on that on its underlying building blocks. Here, we note that our usage of a collision-resistant hash function for computing the tags $\text{tag}_{\vec{\text{vk}}} = \text{CRH.Eval}(\text{k}_{\text{CRH}}, \vec{\text{vk}})$ is in fact not required, and is done for providing a more direct proof of security. Specifically, we could have instead used $\text{tag}_{\vec{\text{vk}}} = \text{aggvk}_\varepsilon$, and rely on the fact that the key aggregation of any one-time multi-signature scheme is collision resistant (as effectively shown by the proof of Claim 4.4 for our one-time scheme).

Theorem 5.1. *Let $\epsilon_{\text{binding}} = \epsilon_{\text{binding}}(\lambda)$, $q_{\text{Sign}} = q_{\text{Sign}}(\lambda)$ and $\ell = \ell(\lambda)$ be functions of the security parameter $\lambda \in \mathbb{N}$, and let A be a probabilistic polynomial-time algorithm that issues q_{Sign} signing queries for ℓ -bit messages. Then, assuming that COM is $\epsilon_{\text{binding}}$ -statistically binding, there exist probabilistic polynomial-time algorithms B_1, \dots, B_6 such that for every $\lambda \in \mathbb{N}$ it holds that*

$$\begin{aligned} \text{Adv}_{\Pi, A}^{\text{SSetMS}}(\lambda) &\leq \epsilon_{\text{binding}}(\lambda) + \text{Adv}_{\text{ZK}, B_1}^{\text{zk}}(\lambda) + \text{Adv}_{\text{COM}, B_2}^{\text{hiding}}(\lambda) + \text{Adv}_{\text{PRF}, B_3}^{\text{prf}}(\lambda) \\ &\quad + \text{Adv}_{\text{ZK}, B_4}^{\text{ss}}(\lambda) + \text{Adv}_{\text{CRH}, B_5}^{\text{crh}}(\lambda) + (2 \cdot \ell \cdot q_{\text{Sign}} + 1) \cdot \text{Adv}_{\text{1T}, B_6}^{\text{1TimeMS}}(\lambda). \end{aligned}$$

Proof of Theorem 5.1. Let A be a probabilistic polynomial-time adversary that participates in the experiment $\text{Exp}_{\Pi, A}^{\text{SSetMS}}(\lambda)$ by issuing $q_{\text{Sign}} = q_{\text{Sign}}(\lambda)$ signing queries. Note that the signing oracle, responding to these queries, may invoke the key-generation algorithm 1T.KAagg at most $2 \cdot \ell \cdot q_{\text{Sign}}$ times during the experiment $\text{Exp}_{\Pi, A}^{\text{SSetMS}}$ (recall that ℓ denotes bit-length of messages). In addition, at the beginning of the experiment $\text{Exp}_{\Pi, A}^{\text{SSetMS}}$, the key-generation algorithm 1T.KAagg is invoked once by Π 's key-generation algorithm KG . Thus, throughout the experiment $\text{Exp}_{\Pi, A}^{\text{SSetMS}}$, the key-generation algorithm 1T.KAagg is invoked at most $L = 2 \cdot \ell \cdot q_{\text{Sign}} + 1$ times.

In what follows we describe a sequence of hybrid experiments $\text{Exp}_0, \dots, \text{Exp}_5$, starting with the experiment $\text{Exp}_0 = \text{Exp}_{\Pi, A}^{\text{SSetMS}}$. Based on the security of the building blocks used in the construction of the scheme Π , we bound the differences between the adversary's success probabilities in each consecutive pair of hybrid experiments, and then show that any adversary participating in experiment Exp_5 may be used to attack the underlying scheme 1T .

Experiment Exp_0 . This is the original experiment $\text{Exp}_{\Pi, A}^{\text{SSetMS}}$. Recall that at the beginning of the experiment, the setup algorithm Setup and the key-generation algorithm KG of the scheme Π

are invoked to produce $\text{pp} = (\text{pp}_{1\text{T}}, \text{crs}_{\text{ZK}}, \text{crs}_{\text{COM}}, \text{k}_{\text{CRH}})$ and $\text{vk} = (\text{vk}_\varepsilon, \text{com})$ which are provided to A , whereas the signing key $\text{sk} = (\text{sk}_\varepsilon, \text{vk}_\varepsilon, \text{k}_{\text{PRF}}, \text{com}, \text{decom})$ is not provided to A .

Experiment Exp_1 . This experiment is obtained from Exp_0 via the following modifications. First, instead of producing the common-reference string crs_{ZK} that is included in the public parameters using the setup algorithm $\text{crs}_{\text{ZK}} \leftarrow \text{ZK.Setup}(1^\lambda)$, it is produced by the simulated setup algorithm $(\text{crs}_{\text{ZK}}, \tau_{\text{ZK}}) \leftarrow \text{Sim}_1(1^\lambda)$. Second, when the experiment produces one-time signing and verification keys as

$$\left(\text{sk}_{m|_j b}^{(k^*)}, \text{vk}_{m|_j b}^{(k^*)} \right) = \text{1T.KG}(\text{pp}_{1\text{T}}; \text{PRF.Eval}(\text{k}_{\text{PRF}}, m|_j b))$$

for some $k^* \in [n]$ and for all $j \in \{0, \dots, j-1\}$ and $b \in \{0, 1\}$, then instead of computing

$$\pi_{m|_j b}^{(k^*)} = \text{ZK.P} \left(\text{crs}_{\text{ZK}}, \left(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}, m|_j b, \text{vk}_{m|_j b}^{(k^*)} \right), \left(\text{decom}, \text{k}_{\text{PRF}}, \text{sk}_{m|_j b}^{(k^*)} \right) \right)$$

it uses the simulated prover for computing

$$\pi_{m|_j b}^{(k^*)} = \text{ZK.Sim}_2 \left(\tau_{\text{ZK}}, \left(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}, m|_j b, \text{vk}_{m|_j b}^{(k^*)} \right) \right).$$

That is, in this experiment it still holds that $(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}, m|_j b, \text{vk}_{m|_j b}^{(k^*)}) \in \mathcal{L}$, but the proofs are generated using the trapdoor τ_{ZK} instead of the witness $(\text{decom}, \text{k}_{\text{PRF}}, \text{sk}_{m|_j b}^{(k^*)})$.

A direct reduction to the adaptive zero-knowledge property of the proof system ZK (see Definition 2.6) yields the following claim:

Claim 5.2. *There exists a probabilistic polynomial-time algorithm B_1 such that for every $\lambda \in \mathbb{N}$ it holds that*

$$|\Pr[\text{Exp}_0(\lambda) = 1] - \Pr[\text{Exp}_1(\lambda) = 1]| \leq \text{Adv}_{\text{ZK}, B_1}^{\text{zk}}(\lambda).$$

Experiment Exp_2 . This experiment is obtained from Exp_1 by modifying the distribution of the commitment com that is included in the verification key $\text{vk} = (\text{vk}_\varepsilon, \text{com})$: Instead of computing $(\text{com}, \text{decom}) \leftarrow \text{COM.Commit}(\text{crs}_{\text{COM}}, \text{k}_{\text{PRF}})$, the experiment computes $(\text{com}, \text{decom}) \leftarrow \text{COM.Commit}(\text{crs}_{\text{COM}}, 0^{|\text{k}_{\text{PRF}}|})$. Note that the key k_{PRF} is still sampled and used for deriving randomness for generating the one-time signing and verification keys.

A direct reduction to the computational hiding property of the commitment scheme COM (see Definition 2.5) yields the following claim:

Claim 5.3. *There exists a probabilistic polynomial-time algorithm B_2 such that for every $\lambda \in \mathbb{N}$ it holds that*

$$|\Pr[\text{Exp}_1(\lambda) = 1] - \Pr[\text{Exp}_2(\lambda) = 1]| \leq \text{Adv}_{\text{COM}, B_2}^{\text{hiding}}(\lambda).$$

Experiment Exp_3 . This experiment is obtained from Exp_2 by producing the one-time signing and verification keys using true randomness instead of using the pseudorandom function. That is, instead of computing

$$\left(\text{sk}_{m|_j b}^{(k^*)}, \text{vk}_{m|_j b}^{(k^*)} \right) = \text{1T.KG}(\text{pp}_{1\text{T}}; \text{PRF.Eval}(\text{k}_{\text{PRF}}, m|_j b))$$

for some $k^* \in [n]$ and for all $j \in \{0, \dots, \ell - 1\}$ and $b \in \{0, 1\}$, the experiment samples $r_{m|j,b}^{(k^*)} \leftarrow \{0, 1\}^*$ of the appropriate length, and computes

$$\left(\text{sk}_{m|j,b}^{(k^*)}, \text{vk}_{m|j,b}^{(k^*)} \right) = \text{1T.KG} \left(\text{pp}_{\text{1T}}; r_{m|j,b}^{(k^*)} \right).$$

The experiment stores all of the one-time signing and verification keys that it produces, and whenever such keys are used more than once (i.e., on signing queries corresponding to intersecting paths from the root ε to different leaves), the corresponding one-time signing and verification keys are reused (or, equivalently, produced again using the same random strings).

A direct reduction to the security of the pseudorandom function PRF (see Definition 2.3) yields the following claim:

Claim 5.4. *There exists a probabilistic polynomial-time algorithm B_3 such that for every $\lambda \in \mathbb{N}$ it holds that*

$$|\Pr [\text{Exp}_2(\lambda) = 1] - \Pr [\text{Exp}_3(\lambda) = 1]| \leq \text{Adv}_{\text{PRF}, B_3}^{\text{prf}}(\lambda).$$

Experiment Exp_4 . This experiment is obtained from Exp_3 by modifying its output in some cases, as follows. If there exists a one-time signing key produced by the experiment Exp_3 that is used for signing more than one message then Exp_4 outputs 0, and otherwise it outputs the output of Exp_3 .

The following claim upper bounds the difference between the adversary's success probabilities in the experiments Exp_3 and Exp_4 based on the simulation soundness property of the proof system ZK (see Definition 2.6) and on the statistical binding property of the commitment scheme COM (see Definition 2.5):

Claim 5.5. *Assuming that COM is $\epsilon_{\text{binding}}$ -statistically binding, there exists a probabilistic polynomial-time algorithm B_4 such that for every $\lambda \in \mathbb{N}$ it holds that*

$$|\Pr [\text{Exp}_3(\lambda) = 1] - \Pr [\text{Exp}_4(\lambda) = 1]| \leq 2 \cdot \text{Adv}_{\text{ZK}, B_4}^{\text{ss}}(\lambda) + \epsilon_{\text{binding}}(\lambda).$$

Proof of Claim 5.5. Recall that the adversary A is restricted to issuing signing queries with respect to a single vector of verification keys. We denote this vector by

$$\vec{\text{vk}} = \left(\left(\text{vk}_{\varepsilon}^{(1)}, \text{com}^{(1)} \right), \dots, \left(\text{vk}_{\varepsilon}^{(n)}, \text{com}^{(n)} \right) \right),$$

and we denote by $k^* \in [n]$ the unique integer for which it holds that $(\text{vk}_{\varepsilon}, \text{com}) = \left(\text{vk}_{\varepsilon}^{(k^*)}, \text{com}^{(k^*)} \right)$, where $\text{vk} = (\text{vk}_{\varepsilon}, \text{com})$ is the honest verification key provided as input to the adversary (note that without loss of generality such an integer exists and it is unique since otherwise the adversary obtains no responses to signing queries). Additionally, denote by $\left\{ \left(\text{sk}_{\alpha}^{(k^*)}, \text{vk}_{\alpha}^{(k^*)} \right) \right\}_{\alpha \in \{\varepsilon\} \cup \{0, 1\}^{\leq \ell}}$ the subset of one-time signing and verification keys used by experiment Exp_3 for responding to the A 's signing queries (note that not all of these keys have actually been generated, but rather a subset that corresponds to the paths from the root ε to the leaves corresponding to the signed messages). Denote by ReUsedKey the event in which on these pairs is used for signing more than one message in experiment Exp_3 , then by the definition of experiment Exp_4 it holds that

$$|\Pr [\text{Exp}_3(\lambda) = 1] - \Pr [\text{Exp}_4(\lambda) = 1]| \leq \Pr [\text{ReUsedKey}].$$

Consider now an execution of Exp_3 in which some signing key $\text{sk}_\alpha^{(k^*)}$ was used for signing more than one message. Note that a signing key $\text{sk}_\alpha^{(k^*)}$ for $\alpha \in \{0, 1\}^\ell$ is always used for signing exactly one message (which is the message $(\alpha, \text{tag}_{\vec{vk}})$). Therefore, it must be that $\alpha \in \{\varepsilon\} \cup \{0, 1\}^{<\ell}$, and thus $\text{sk}_\alpha^{(k^*)}$ was used for signing the concatenation of two aggregated verification keys. Moreover, since all signing queries are issued with respect to the same “root” keys $\text{vk}_\varepsilon^{(1)}, \dots, \text{vk}_\varepsilon^{(n)}$, then there must exist a signing key $\text{sk}_\alpha^{(k^*)}$ for $\alpha \in \{0, 1\}^\ell$ that was used for signing more than one message with respect to the same vector of verification keys $\vec{vk}_\alpha = (\text{vk}_\alpha^{(1)}, \dots, \text{vk}_\alpha^{(n)})$. We denote the two signing operations as

$$\begin{aligned}\sigma_\alpha^{(k^*)} &= \text{1T.Sign}\left(\text{pp}_{1\text{T}}, \text{sk}_\alpha^{(k^*)}, \vec{vk}_\alpha, (\text{aggvk}_{\alpha 0}, \text{aggvk}_{\alpha 1})\right) \\ \hat{\sigma}_\alpha^{(k^*)} &= \text{1T.Sign}\left(\text{pp}_{1\text{T}}, \text{sk}_\alpha^{(k^*)}, \vec{vk}_\alpha, (\text{agg}\hat{\text{v}}k_{\alpha 0}, \text{agg}\hat{\text{v}}k_{\alpha 1})\right)\end{aligned}$$

where $(\text{aggvk}_{\alpha 0}, \text{aggvk}_{\alpha 1}) \neq (\text{agg}\hat{\text{v}}k_{\alpha 0}, \text{agg}\hat{\text{v}}k_{\alpha 1})$. This means that for some $b \in \{0, 1\}$, the adversary provided two inputs $(\text{vk}_{\alpha b}^{(i)}, \pi_{\alpha b}^{(i)})$ and $(\hat{\text{v}}k_{\alpha b}^{(i)}, \hat{\pi}_{\alpha b}^{(i)})$ such that $\text{vk}_{\alpha b}^{(i)} \neq \hat{\text{v}}k_{\alpha b}^{(i)}$, and these were used for computing the aggregated verification keys $\text{aggvk}_{\alpha b} \neq \text{agg}\hat{\text{v}}k_{\alpha b}$ after verifying the proofs $\pi_{\alpha b}^{(i)}$ and $\hat{\pi}_{\alpha b}^{(i)}$.

However, except with probability $\epsilon_{\text{binding}}$, the common-references string crs_{com} is such that $\text{com}^{(i)}$ uniquely determines a key $\text{k}_{\text{PRF}}^{(i)}$ for the pseudorandom function, and therefore at most one of the keys $\text{vk}_{\alpha b}^{(i)}$ and $\hat{\text{v}}k_{\alpha b}^{(i)}$ was actually produced using the randomness $\text{PRF.Eval}(\text{k}_{\text{PRF}}^{(i)}, \alpha b)$. That is, although both proofs pass verification, it must hold that either

$$\left(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \text{vk}_{\alpha b}^{(i)}\right) \notin \mathcal{L}$$

or

$$\left(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \hat{\text{v}}k_{\alpha b}^{(i)}\right) \notin \mathcal{L}.$$

This leads to an algorithm B_4 that participates in the simulation-soundness experiment $\text{Exp}_{\text{ZK}, B_4}^{\text{SS}}(\lambda)$ by simulating the experiment Exp_3 to A , where crs_{ZK} is generated by Sim_1 and provided as input to B_4 , and where B_4 generates proofs using oracle access to Sim_2 (recall Definition 2.6). Whenever B_4 observes $(\text{vk}_{\alpha b}^{(i)}, \pi_{\alpha b}^{(i)})$ and $(\hat{\text{v}}k_{\alpha b}^{(i)}, \hat{\pi}_{\alpha b}^{(i)})$ such that $\text{vk}_{\alpha b}^{(i)} \neq \hat{\text{v}}k_{\alpha b}^{(i)}$ and both pass verification as above, then B_4 outputs

$$\left(\left(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \text{vk}_{\alpha b}^{(i)}\right), \pi_{\alpha b}^{(i)}\right)$$

with probability 1/2 and

$$\left(\left(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \hat{\text{v}}k_{\alpha b}^{(i)}\right), \hat{\pi}_{\alpha b}^{(i)}\right)$$

with probability 1/2. We observe that the proofs $\pi_{\alpha b}^{(i)}$ and $\hat{\pi}_{\alpha b}^{(i)}$ could not have been previously provided by the oracle Sim_2 to the algorithm B_4 as responses to the queries $(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \text{vk}_{\alpha b}^{(i)})$ or $(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \hat{\text{v}}k_{\alpha b}^{(i)})$, respectively (if such queries were issued). Specifically, for the internal nodes $\alpha 0$ and $\alpha 1$, the algorithm B_4 queries the oracle Sim_2 only with the two instances $(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(k^*)}, \alpha b, \text{vk}_{\alpha b}^{(k^*)})$ for $b = 0$ and $b = 1$. However, as instructed in Steps 1 and 3 of the algorithm Sign , the algorithm B_4 aborts the signing protocol if either $\text{vk}_{\alpha b}^{(k^*)} = \text{vk}_{\alpha b}^{(i)}$ or $\text{vk}_{\alpha b}^{(k^*)} = \hat{\text{v}}k_{\alpha b}^{(i)}$. Therefore, the algorithm B_4 does not query the oracle Sim_2 with either $(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \text{vk}_{\alpha b}^{(i)})$

or $(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}^{(i)}, \alpha b, \hat{\text{vk}}_{\alpha b}^{(i)})$. This means that for the algorithm B_4 it holds that

$$\begin{aligned} \text{Adv}_{\text{ZK}, B_4}^{\text{SS}}(\lambda) &\geq \frac{1}{2} \cdot \left(\Pr[\text{ReUsedKey}] - \epsilon_{\text{binding}}(\lambda) \right) \\ &\geq \frac{1}{2} \cdot \left(|\Pr[\text{Exp}_3(\lambda) = 1] - \Pr[\text{Exp}_4(\lambda) = 1]| - \epsilon_{\text{binding}}(\lambda) \right) \end{aligned}$$

as required. \blacksquare

Experiment Exp_5 . This experiment is obtained from Exp_4 by modifying its output in some cases, as follows. Denote by $\vec{\text{vk}}$ the vector of one-time verification keys with respect to which A issues signing queries, and denote by $\vec{\text{vk}}$ the vector of verification keys included in A 's output. If $\vec{\text{vk}} \neq \vec{\text{vk}}$ and $\text{CRH.Eval}(k_{\text{CRH}}, \vec{\text{vk}}) = \text{CRH.Eval}(k_{\text{CRH}}, \vec{\text{vk}})$ then Exp_5 outputs 0, and otherwise it outputs the output of Exp_4 .

A direct reduction to the collision resistance property of the function family CRH (see Definition 2.4) yields the following claim:

Claim 5.6. *There exists a probabilistic polynomial-time algorithm B_5 such that for every $\lambda \in \mathbb{N}$ it holds that*

$$|\Pr[\text{Exp}_4(\lambda) = 1] - \Pr[\text{Exp}_5(\lambda) = 1]| \leq \text{Adv}_{\text{CRH}, B_5}^{\text{crh}}(\lambda).$$

The following claim upper bounds the success probability of the adversary A in the experiment Exp_5 based on the security of the one-time multi-signature scheme 1T :

Claim 5.7. *There exists a probabilistic polynomial-time algorithm B_6 such that for every $\lambda \in \mathbb{N}$ it holds that*

$$\Pr[\text{Exp}_5(\lambda) = 1] \leq L \cdot \text{Adv}_{1\text{T}, B_6}^{1\text{TimeMS}}(\lambda).$$

Proof of Claim 5.7. Let B_6 be the algorithm that on input $\text{crs}_{1\text{T}}$ and $\text{vk}_{1\text{T}}$ participates in the experiment $\text{Exp}_{1\text{T}, B_6}^{1\text{TimeMS}}$ by emulating the experiment Exp_5 to the algorithm A as follows. First, the algorithm B_6 computes $(\text{crs}_{\text{ZK}}, \tau_{\text{ZK}}) \leftarrow \text{Sim}_1(1^\lambda)$, $\text{crs}_{\text{COM}} \leftarrow \text{COM.Setup}(1^\lambda)$, $k_{\text{CRH}} \leftarrow \text{CRH.KG}(1^\lambda)$, and $(\text{com}, \text{decom}) \leftarrow \text{COM.Commit}(\text{crs}_{\text{COM}}, 0^{|\text{k}|})$. Then, recall that in the experiment Exp_5 , the algorithm B_6 generates at most $L = 2 \cdot \ell \cdot q_{\text{Sign}} + 1$ one-time signing and verification keys. The algorithm B_6 samples $v \leftarrow [L]$, determining the index of the one-time verification key for which $\text{vk}_{1\text{T}}$ will be used instead of sampling a fresh key pair. In particular, if $v = 1$ then B_6 sets $\text{vk}_\varepsilon = \text{vk}_{1\text{T}}$, and otherwise B_6 samples $(\text{sk}_\varepsilon, \text{vk}_\varepsilon) \leftarrow 1\text{T.KG}(\text{pp}_{1\text{T}})$. Next, B_6 invokes the algorithm A on the public parameters $\text{pp} = (\text{pp}_{1\text{T}}, \text{crs}_{\text{ZK}}, \text{crs}_{\text{COM}}, k_{\text{CRH}})$ and the verification key $\text{vk} = (\text{vk}_\varepsilon, \text{com})$, and responds to A 's signing queries as follows.

When A issues a signing query of the form $(\vec{\text{vk}}, m)$, where

$$\vec{\text{vk}} = \left((\text{vk}_\varepsilon^{(1)}, \text{com}^{(1)}), \dots, (\text{vk}_\varepsilon^{(n)}, \text{com}^{(n)}) \right),$$

the algorithm B_6 first verifies that there is a unique $k^* \in [n]$ for which $(\text{vk}_\varepsilon, \text{com}) = (\text{vk}_\varepsilon^{(i)}, \text{com}^{(i)})$, and otherwise responds with \perp . Then, for any $j \in \{0, \dots, \ell - 1\}$ and $b \in \{0, 1\}$, the algorithm B_6 samples $r_{m|_j b}^{(k^*)} \leftarrow \{0, 1\}^*$ of the appropriate length, and computes

$$\begin{aligned} (\text{sk}_{m|_j b}^{(k^*)}, \text{vk}_{m|_j b}^{(k^*)}) &= 1\text{T.KG}(\text{pp}_{1\text{T}}; r_{m|_j b}^{(k^*)}) \\ \pi_{m|_j b}^{(k^*)} &= \text{ZK.Sim}_2\left(\tau_{\text{ZK}}, \left(\text{pp}_{1\text{T}}, \text{crs}_{\text{com}}, \text{com}, m|_j b, \text{vk}_{m|_j b}^{(k^*)}\right)\right), \end{aligned}$$

with the exception of setting $\mathbf{vk}_{m|jb}^{(k^*)} = \mathbf{vk}_{1T}$ for the v -th key $\mathbf{vk}_{m|jb}^{(k^*)}$ that is produced as part of the emulation whenever $v > 1$ (recall that the case $v = 1$ was already handled differently by setting $\mathbf{vk}_\varepsilon = \mathbf{vk}_{1T}$). The algorithm B_6 responds to the signing query with $\left\{ \left(\mathbf{vk}_{m|jb}^{(k^*)}, \pi_{m|jb}^{(k^*)} \right) \right\}_{j \in \{0, \dots, \ell-1\}, b \in \{0,1\}}$. In addition, when A issues a signing query of the form

$$\left\{ \left(\mathbf{vk}_{m|jb}^{(i)}, \pi_{m|jb}^{(i)} \right) \right\}_{i \in [n] \setminus \{k^*\}, j \in \{0, \dots, \ell-1\}, b \in \{0,1\}},$$

corresponding to an earlier signing query of the form $(\vec{\mathbf{vk}}, m)$, the algorithm B_6 first verifies that for all $(i, j, b) \in ([n] \setminus \{k^*\}) \times \{0, \dots, \ell-1\} \times b \in \{0,1\}$ it holds that

$$\text{ZK.V} \left(\text{crs}_{\text{ZK}}, \left(\text{pp}_{1T}, \text{crs}_{\text{com}}, \text{com}^{(i)}, m|_jb, \mathbf{vk}_{m|jb}^{(i)}, \pi_{m|jb}^{(i)} \right) \right) = 1$$

and $\mathbf{vk}_{m|jb}^{(i)} \neq \mathbf{vk}_{m|jb}^{(k^*)}$, and otherwise responds with \perp . Then, the algorithm B_6 computes

$$\begin{aligned} \vec{\mathbf{vk}}_{m|jb} &= \left(\mathbf{vk}_{m|jb}^{(1)}, \dots, \mathbf{vk}_{m|jb}^{(n)} \right) \text{ for all } j \in \{0, \dots, \ell-1\} \text{ and } b \in \{0,1\} \\ \text{aggvk}_{m|jb} &= \text{1T.KAagg} \left(\text{pp}_{1T}, \vec{\mathbf{vk}}_{m|jb} \right) \text{ for all } j \in \{0, \dots, \ell-1\} \text{ and } b \in \{0,1\} \\ \sigma_{m|j}^{(k^*)} &= \text{1T.Sign} \left(\text{pp}_{1T}, \text{sk}_{m|j}^{(k^*)}, \vec{\mathbf{vk}}_{m|j}, \left(\text{aggvk}_{m|j0}, \text{aggvk}_{m|j1} \right) \right) \text{ for all } j \in \{0, \dots, \ell-1\} \\ \sigma_m^{(k^*)} &= \text{1T.Sign} \left(\text{pp}_{1T}, \text{sk}_m^{(k^*)}, \vec{\mathbf{vk}}_m, m \right), \end{aligned}$$

with the exception of accessing the signing oracle $\text{1T.Sign}(\text{pp}_{1T}, \text{sk}_{1T}, \cdot, \cdot)$ for obtaining the signature $\sigma_{m|j}^{(k^*)}$ or $\sigma_m^{(k^*)}$ corresponding to the v -th one-time verification key that is produced as part of the emulation (recall that the v -th one-time verification key was set to \mathbf{vk}_{1T} and thus its corresponding signing key sk_{1T} is not known to B_6). The algorithm B_6 locally stores all signatures produced throughout the emulation, so that no signature has to be recomputed, and no query to the oracle $\text{1T.Sign}(\text{pp}_{1T}, \text{sk}_{1T}, \cdot, \cdot)$ has to be repeated (as specified above, we assume without loss of generality that the signing algorithm 1T.Sign is deterministic). The algorithm B_6 responds to the signing query with

$$\sigma^{(k^*)} = \left(\left\{ \left(\sigma_{m|j}^{(k^*)}, \mathbf{vk}_{m|j0}^{(k^*)}, \mathbf{vk}_{m|j1}^{(k^*)} \right) \right\}_{j \in \{0, \dots, \ell-1\}}, \sigma_m^{(k^*)} \right).$$

At some point, the algorithm A produces an output $(\vec{\mathbf{vk}}, \hat{m}, \text{agg}\sigma)$ and halts. We let

$$\begin{aligned} \vec{\mathbf{vk}} &= \left(\left(\hat{\mathbf{vk}}_\varepsilon^{(1)}, \hat{\text{com}}^{(1)} \right), \dots, \left(\hat{\mathbf{vk}}_\varepsilon^{(n)}, \hat{\text{com}}^{(n)} \right) \right) \\ \text{tag}_{\vec{\mathbf{vk}}} &= \text{CRH.Eval} \left(\text{k}_{\text{CRH}}, \vec{\mathbf{vk}} \right) \\ \text{aggvk}_\varepsilon &= \text{1T.KAagg} \left(\text{pp}_{1T}, \left(\hat{\mathbf{vk}}_\varepsilon^{(1)}, \dots, \hat{\mathbf{vk}}_\varepsilon^{(n)} \right) \right) \\ \text{agg}\sigma &= \left(\left\{ \left(\text{agg}\sigma_{\hat{m}|j}, \text{aggvk}_{\hat{m}|j0}, \text{aggvk}_{\hat{m}|j1} \right) \right\}_{j \in \{0, \dots, \ell-1\}}, \text{agg}\sigma_{\hat{m}} \right). \end{aligned}$$

In addition, as above, we denote by $\vec{\mathbf{vk}} = \left(\left(\mathbf{vk}_\varepsilon^{(1)}, \text{com}^{(1)} \right), \dots, \left(\mathbf{vk}_\varepsilon^{(n)}, \text{com}^{(n)} \right) \right)$ the vector of verification keys relative to which A issued signing queries (recall that there is only one such vector). For determining its output, the algorithm B_6 first checks whether or not the event $\{\text{Exp}_5(\lambda) = 1\}$

occurs, when the event is defined within the probability space induced by the emulation. If the event $\{\text{Exp}_5(\lambda) = 1\}$ does not occur, then B_6 outputs \perp . Otherwise, note that if the event $\{\text{Exp}_5(\lambda) = 1\}$ does occur, then since A issued all signing queries with respect to the same vector of verification keys \vec{vk} , the fact that no signing key was used to sign more than one message implies that A 's signing queries and their corresponding responses uniquely define the values

$$\begin{aligned}\vec{vk}_{m|jb} &= \left(vk_{m|jb}^{(1)}, \dots, vk_{m|jb}^{(n)} \right) \text{ for all } j \in \{0, \dots, \ell - 1\} \text{ and } b \in \{0, 1\} \\ \text{aggvk}_{m|jb} &= \text{1T.KAgg} \left(\text{pp}_{1\text{T}}, \vec{vk}_{m|jb} \right) \text{ for all } j \in \{0, \dots, \ell - 1\} \text{ and } b \in \{0, 1\}\end{aligned}$$

for each message m for which a signing query was issued. Equipped with this observation, if the event $\{\text{Exp}_5(\lambda) = 1\}$ occurs then the algorithm B_6 determines its output by distinguishing between the following three disjoint cases:

Case I: If $\left(vk_\varepsilon^{(1)}, \dots, vk_\varepsilon^{(n)} \right) \neq \left(\hat{vk}_\varepsilon^{(1)}, \dots, \hat{vk}_\varepsilon^{(n)} \right)$, then B_6 outputs

$$\begin{aligned}\vec{vk}^* &= \left(\hat{vk}_\varepsilon^{(1)}, \dots, \hat{vk}_\varepsilon^{(n)} \right) \\ m^* &= \left(\text{aggvk}_0, \text{aggvk}_1 \right) \\ \text{agg}\sigma^* &= \text{agg}\sigma_\varepsilon\end{aligned}$$

Case II: Else, if there exists $j \in \{0, \dots, \ell - 1\}$ for which $\text{aggvk}_{\hat{m}|j} = \text{agg}\hat{vk}_{\hat{m}|j}$ but it holds that $\left(\text{aggvk}_{\hat{m}|j0}, \text{aggvk}_{\hat{m}|j1} \right) \neq \left(\text{agg}\hat{vk}_{\hat{m}|j0}, \text{agg}\hat{vk}_{\hat{m}|j1} \right)$ (including the case where $\text{aggvk}_{\hat{m}|j0}$ or $\text{aggvk}_{\hat{m}|j1}$ were not at all generated during B_6 's responses to A 's signing queries), then B_6 outputs

$$\begin{aligned}\vec{vk}^* &= \left(vk_{\hat{m}|j}^{(1)}, \dots, vk_{\hat{m}|j}^{(n)} \right) \\ m^* &= \left(\text{agg}\hat{vk}_{\hat{m}|j0}, \text{agg}\hat{vk}_{\hat{m}|j1} \right) \\ \text{agg}\sigma^* &= \text{agg}\sigma_{\hat{m}|j}\end{aligned}$$

for the minimal such $j \in \{0, \dots, \ell - 1\}$.

Case III: Else, B_6 outputs

$$\begin{aligned}\vec{vk}^* &= \left(vk_{\hat{m}}^{(1)}, \dots, vk_{\hat{m}}^{(n)} \right) \\ m^* &= \left(\hat{m}, \text{tag}_{\vec{vk}} \right) \\ \text{agg}\sigma^* &= \text{agg}\sigma_{\hat{m}}.\end{aligned}$$

This concludes the description of the algorithm B_6 , and now we turn to analyzing its success probability in the experiment $\text{Exp}_{1\text{T}, B_6}^{\text{1TimeMS}}$. The algorithm B_6 perfectly emulates the experiment Exp_5 to A , and therefore

$$\begin{aligned}\text{Adv}_{1\text{T}, B_6}^{\text{1TimeMS}}(\lambda) &= \Pr \left[\text{Exp}_{1\text{T}, B_6}^{\text{1TimeMS}}(\lambda) = 1 \right] \\ &= \Pr \left[\text{Exp}_5(\lambda) = 1 \right] \cdot \Pr \left[\text{Exp}_{1\text{T}, B_6}^{\text{1TimeMS}}(\lambda) = 1 \mid \text{Exp}_5(\lambda) = 1 \right],\end{aligned}\tag{5.1}$$

where, as noted above, the event $\{\text{Exp}_5(\lambda) = 1\}$ is defined within the probability space induced by the emulation. Note that, conditioned on the event $\{\text{Exp}_5(\lambda) = 1\}$, no one-time signing key used throughout the emulation was used to sign more than one message, and in addition:

1. $\text{vk} = (\text{vk}_\varepsilon, \text{com}) \in \vec{\text{vk}}$ and the algorithm A did not issue a signing query $(\vec{\text{vk}}, \hat{m})$.
2. $1\text{T.Verify}(\text{pp}_{1\text{T}}, \text{aggvk}_{m|j}, (\text{aggvk}_{m|j0}, \text{aggvk}_{m|j1}), \text{agg}\sigma_{m|j}) = 1$ for every $j \in \{0, \dots, \ell - 1\}$.
3. $1\text{T.Verify}(\text{pp}_{1\text{T}}, \text{aggvk}_{\hat{m}}, (\hat{m}, \text{tag}_{\vec{\text{vk}}}), \text{agg}\sigma_{\hat{m}}) = 1$.

We now show that, conditioned on the event $\{\text{Exp}_5(\lambda) = 1\}$, then in each of Case I, Case II and Case II it holds that $\text{Exp}_{1\text{T}, B_6}^{1\text{TimeMS}}(\lambda) = 1$ with probability at least $1/L$ over the choice of $v \leftarrow [L]$ (which is independent of A 's view):

Case I: In this case, if $v = 1$ (an event that occurs with probability $1/L$), then $\text{vk}_\varepsilon = \text{vk}_{1\text{T}}$ (recall that $\text{vk}_{1\text{T}}$ is the one-time verification key provided to B_6 at the beginning of the experiment $\text{Exp}_{1\text{T}, B_6}^{1\text{TimeMS}}$). We now verify that the three conditions required for $\text{Exp}_{1\text{T}, B_6}^{1\text{TimeMS}}(\lambda) = 1$ are satisfied:

- Since $(\text{vk}_\varepsilon, \text{com}) \in \vec{\text{vk}} = \left((\hat{\text{vk}}_\varepsilon^{(1)}, \text{com}^{(1)}), \dots, (\hat{\text{vk}}_\varepsilon^{(n)}, \text{com}^{(n)}) \right)$ and $\text{vk}_\varepsilon = \text{vk}_{1\text{T}}$, then $\text{vk}_{1\text{T}} \in (\hat{\text{vk}}_\varepsilon^{(1)}, \dots, \hat{\text{vk}}_\varepsilon^{(n)}) = \vec{\text{vk}}^*$.
- Since $1\text{T.Verify}(\text{pp}_{1\text{T}}, \text{aggvk}_\varepsilon, (\text{aggvk}_0, \text{aggvk}_1), \text{agg}\sigma_\varepsilon) = 1$, it holds that

$$\begin{aligned} & 1\text{T.Verify}(\text{pp}_{1\text{T}}, \text{KAgg}(\text{pp}_{1\text{T}}, \vec{\text{vk}}^*), m^*, \text{agg}\sigma^*) \\ &= 1\text{T.Verify}(\text{pp}_{1\text{T}}, \text{aggvk}_\varepsilon, (\text{aggvk}_0, \text{aggvk}_1), \text{agg}\sigma_\varepsilon) \\ &= 1 \end{aligned}$$

- Since B_6 queried the signing oracle $1\text{T.Sign}(\text{pp}_{1\text{T}}, \text{sk}_{1\text{T}}, \cdot, \cdot)$ exactly once, and the input to this query was $\left((\text{vk}_\varepsilon^{(1)}, \dots, \text{vk}_\varepsilon^{(n)}), (\text{aggvk}_0, \text{aggvk}_1) \right)$, and since in Case I we have

$$(\text{vk}_\varepsilon^{(1)}, \dots, \text{vk}_\varepsilon^{(n)}) \neq (\hat{\text{vk}}_\varepsilon^{(1)}, \dots, \hat{\text{vk}}_\varepsilon^{(n)}),$$

this means that B_6 did not query the signing oracle with

$$(\vec{\text{vk}}^*, m^*) = \left((\hat{\text{vk}}_\varepsilon^{(1)}, \dots, \hat{\text{vk}}_\varepsilon^{(n)}), (\text{aggvk}_0, \text{aggvk}_1) \right).$$

Case II: Let $j \in \{0, \dots, \ell - 1\}$ be the minimal integer for which $\text{aggvk}_{\hat{m}|j} = \text{aggvk}_{\hat{m}|j}$ but it holds that $(\text{aggvk}_{\hat{m}|j0}, \text{aggvk}_{\hat{m}|j1}) \neq (\text{aggvk}_{\hat{m}|j0}, \text{aggvk}_{\hat{m}|j1})$ (or that $\text{aggvk}_{\hat{m}|j0}$ or $\text{aggvk}_{\hat{m}|j1}$ were not at all generated during B_6 's responses to A 's signing queries), and let $\vec{\text{vk}}_{\hat{m}|j} = (\text{vk}_{\hat{m}|j}^{(1)}, \dots, \text{vk}_{\hat{m}|j}^{(n)})$. In this case, if $v \in [L]$ equals the index of the verification key $\text{vk}_{\hat{m}|j}^{(k^*)}$ (an event that occurs with probability $1/L$), then $\text{vk}_{\hat{m}|j}^{(k^*)} = \text{vk}_{1\text{T}}$. We now verify that the three conditions required for $\text{Exp}_{1\text{T}, B_6}^{1\text{TimeMS}}(\lambda) = 1$ are satisfied:

- Since $\text{vk}_{\hat{m}|j}^{(k^*)} \in (\text{vk}_{\hat{m}|j}^{(1)}, \dots, \text{vk}_{\hat{m}|j}^{(n)}) = \vec{\text{vk}}^*$ and $\text{vk}_{\hat{m}|j}^{(k^*)} = \text{vk}_{1\text{T}}$, then $\text{vk}_{1\text{T}} \in \vec{\text{vk}}^*$.

- Since $1\text{T.Verify}\left(\text{pp}_{1\text{T}}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j}, \left(\text{agg}\hat{\text{vk}}_{\hat{m}|_j 0}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j 1}\right), \text{agg}\hat{\sigma}_{\hat{m}|_j}\right) = 1$ and $\text{agg}\hat{\text{vk}}_{\hat{m}|_j} = \text{agg}\hat{\text{vk}}_{\hat{m}|_j}$, it holds that

$$\begin{aligned} 1\text{T.Verify}\left(\text{pp}_{1\text{T}}, \text{KAgg}\left(\text{pp}_{1\text{T}}, \vec{\text{vk}}^*\right), m^*, \text{agg}\sigma^*\right) \\ &= 1\text{T.Verify}\left(\text{pp}_{1\text{T}}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j}, \left(\text{agg}\hat{\text{vk}}_{\hat{m}|_j 0}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j 1}\right), \text{agg}\hat{\sigma}_{\hat{m}|_j}\right) \\ &= 1 \end{aligned}$$

- Since B_6 queried the signing oracle $1\text{T.Sign}(\text{pp}_{1\text{T}}, \text{sk}_{1\text{T}}, \cdot, \cdot)$ at most once, and the input to this query (if existed) was

$$\left(\left(\text{vk}_{\hat{m}|_j}^{(1)}, \dots, \text{vk}_{\hat{m}|_j}^{(n)}\right), \left(\text{agg}\hat{\text{vk}}_{\hat{m}|_j 0}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j 1}\right)\right),$$

and since in Case II we have

$$\left(\text{agg}\hat{\text{vk}}_{\hat{m}|_j 0}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j 1}\right) \neq \left(\text{agg}\hat{\text{vk}}_{\hat{m}|_j 0}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j 1}\right),$$

this means that B_6 did not query the signing oracle with

$$\left(\vec{\text{vk}}^*, m^*\right) = \left(\left(\text{vk}_{\hat{m}|_j}^{(1)}, \dots, \text{vk}_{\hat{m}|_j}^{(n)}\right), \left(\text{agg}\hat{\text{vk}}_{\hat{m}|_j 0}, \text{agg}\hat{\text{vk}}_{\hat{m}|_j 1}\right)\right).$$

Case III: In this case, if $v \in [L]$ equals the index of the verification key $\text{vk}_{\hat{m}}^{(k^*)}$ (an event that occurs with probability $1/L$), then $\text{vk}_{\hat{m}}^{(k^*)} = \text{vk}_{1\text{T}}$. We now verify that the three conditions required for $\text{Exp}_{1\text{T}, B_6}^{\text{TimeMS}}(\lambda) = 1$ are satisfied:

- Since $\text{vk}_{\hat{m}}^{(k^*)} \in \left(\text{vk}_{\hat{m}}^{(1)}, \dots, \text{vk}_{\hat{m}}^{(n)}\right) = \vec{\text{vk}}^*$ and $\text{vk}_{\hat{m}}^{(k^*)} = \text{vk}_{1\text{T}}$, then $\text{vk}_{1\text{T}} \in \vec{\text{vk}}^*$.
- Since $1\text{T.Verify}\left(\text{pp}_{1\text{T}}, \text{agg}\hat{\text{vk}}_{\hat{m}}, \left(\hat{m}, \text{tag}_{\vec{\text{vk}}}^{\hat{m}}\right), \text{agg}\hat{\sigma}_{\hat{m}}\right) = 1$, and since in Case III we have $\text{KAgg}\left(\text{pp}_{1\text{T}}, \vec{\text{vk}}^*\right) = \text{agg}\hat{\text{vk}}_{\hat{m}} = \text{agg}\hat{\text{vk}}_{\hat{m}}$, it holds that

$$\begin{aligned} 1\text{T.Verify}\left(\text{pp}_{1\text{T}}, \text{KAgg}\left(\text{pp}_{1\text{T}}, \vec{\text{vk}}^*\right), m^*, \text{agg}\sigma^*\right) \\ &= 1\text{T.Verify}\left(\text{pp}_{1\text{T}}, \text{agg}\hat{\text{vk}}_{\hat{m}}, \left(\hat{m}, \text{tag}_{\vec{\text{vk}}}^{\hat{m}}\right), \text{agg}\hat{\sigma}_{\hat{m}}\right) \\ &= 1 \end{aligned}$$

- The algorithm A did not issue the signing query $\left(\vec{\text{vk}}, \hat{m}\right)$, and therefore if $\vec{\text{vk}} = \vec{\text{vk}}$ this means that A did not issue any signing query with the message \hat{m} . Therefore, if $\vec{\text{vk}} = \vec{\text{vk}}$, then B_6 did not query the oracle $1\text{T.Sign}(\text{pp}_{1\text{T}}, \text{sk}_{1\text{T}}, \cdot, \cdot)$ even once.

In addition, if $\vec{\text{vk}} \neq \vec{\text{vk}}$, it may be that A issued the signing query $(\vec{\text{vk}}, \hat{m})$, and then B_6 queried the oracle $1\text{T.Sign}(\text{pp}_{1\text{T}}, \text{sk}_{1\text{T}}, \cdot, \cdot)$ exactly once. However, the input to this query was $\left(\left(\text{vk}_{\hat{m}}^{(1)}, \dots, \text{vk}_{\hat{m}}^{(n)}\right), \left(\hat{m}, \text{tag}_{\vec{\text{vk}}}^{\hat{m}}\right)\right)$, and since $\text{tag}_{\vec{\text{vk}}}^{\hat{m}} \neq \text{tag}_{\vec{\text{vk}}}^{\hat{m}}$ this means that B_6 did not query the signing oracle with

$$\left(\vec{\text{vk}}^*, m^*\right) = \left(\left(\text{vk}_{\hat{m}}^{(1)}, \dots, \text{vk}_{\hat{m}}^{(n)}\right), \left(\hat{m}, \text{tag}_{\vec{\text{vk}}}^{\hat{m}}\right)\right).$$

Using the fact that Case I, Case II and Case III correspond to disjoint events, we obtain

$$\Pr \left[\text{Exp}_{1T, B_6}^{\text{1TimeMS}}(\lambda) = 1 \mid \text{Exp}_5(\lambda) = 1 \right] \geq \frac{1}{L},$$

and then Equation (5.1) yields

$$\text{Adv}_{1T, B_6}^{\text{1TimeMS}}(\lambda) \geq \Pr [\text{Exp}_5(\lambda) = 1] \cdot \frac{1}{L}.$$

■

Putting together Claims 5.2–5.7, there exist probabilistic polynomial-time algorithms B_1, \dots, B_6 such that for every $\lambda \in \mathbb{N}$ it holds that

$$\begin{aligned} \text{Adv}_{\text{II}, A}^{\text{SSetMS}}(\lambda) &= \Pr [\text{Exp}_0(\lambda) = 1] \\ &\leq |\Pr [\text{Exp}_0(\lambda) = 1] - \Pr [\text{Exp}_1(\lambda) = 1]| \\ &\quad + |\Pr [\text{Exp}_1(\lambda) = 1] - \Pr [\text{Exp}_2(\lambda) = 1]| \\ &\quad + |\Pr [\text{Exp}_2(\lambda) = 1] - \Pr [\text{Exp}_3(\lambda) = 1]| \\ &\quad + |\Pr [\text{Exp}_3(\lambda) = 1] - \Pr [\text{Exp}_4(\lambda) = 1]| \\ &\quad + |\Pr [\text{Exp}_4(\lambda) = 1] - \Pr [\text{Exp}_5(\lambda) = 1]| \\ &\quad + \Pr [\text{Exp}_5(\lambda) = 1] \\ &\leq \epsilon_{\text{binding}}(\lambda) + \text{Adv}_{\text{ZK}, B_1}^{\text{zk}}(\lambda) + \text{Adv}_{\text{COM}, B_2}^{\text{hiding}}(\lambda) + \text{Adv}_{\text{PRF}, B_3}^{\text{prf}}(\lambda) \\ &\quad + \text{Adv}_{\text{ZK}, B_4}^{\text{ss}}(\lambda) + \text{Adv}_{\text{CRH}, B_5}^{\text{crh}}(\lambda) + L \cdot \text{Adv}_{1T, B_6}^{\text{1TimeMS}}(\lambda), \end{aligned}$$

where $L = 2 \cdot \ell \cdot q_{\text{sign}} + 1$. This settles the proof of Theorem 5.1.

■

References

- [AAB⁺02] M. Abdalla, J. H. An, M. Bellare, and C. Nampreppe. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In *Advances in Cryptology – EUROCRYPT ’02*, pages 418–433, 2002.
- [AB21] H. K. Alper and J. Burdges. Two-round trip Schnorr multi-signatures via delinearized witnesses. In *Advances in Cryptology – CRYPTO ’21*, pages 157–188, 2021.
- [AE18] J.-P. Aumasson and G. Endignoux. Improving stateless hash-based signatures. In *Cryptographers’ Track at the RSA Conference*, pages 219–242. Springer, 2018.
- [AHK20] T. Agrikola, D. Hofheinz, and J. Kastner. On instantiating the algebraic group model from falsifiable assumptions. In *Advances in Cryptology – EUROCRYPT ’20*, pages 96–126, 2020.
- [Arg22] Argent. Part I: WTF is account abstraction. Available at <https://www.argent.xyz/blog/wtf-is-account-abstraction/>, 2022.
- [BBB⁺18] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *IEEE Symposium on Security and Privacy*, pages 315–334, 2018.

- [BCI⁺13] N. Bitansky, A. Chiesa, Y. Ishai, R. Ostrovsky, and O. Paneth. Succinct non-interactive arguments via linear interactive proofs. In *Proceedings of the 10th Theory of Cryptography Conference*, pages 315–333, 2013.
- [BCS16] E. Ben-Sasson, A. Chiesa, and N. Spooner. Interactive oracle proofs. In *Proceedings of the 14th Theory of Cryptography Conference*, pages 31–60, 2016.
- [BD21] M. Bellare and W. Dai. Chain reductions for multi-signatures and the HBMS scheme. In *Advances in Cryptology – ASIACRYPT ’21*, pages 650–678, 2021.
- [BDN18] D. Boneh, M. Drijvers, and G. Neven. Compact multi-signatures for smaller blockchains. In *Advances in Cryptology – ASIACRYPT ’18*, pages 435–464, 2018.
- [BFL20] B. Bauer, G. Fuchsbauer, and J. Loss. A classification of computational assumptions in the algebraic group model. In *Advances in Cryptology – CRYPTO ’20*, pages 121–151, 2020.
- [BFM88] M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 103–112, 1988.
- [BGO⁺07] A. Boldyreva, C. Gentry, A. O’Neill, and D. H. Yum. Ordered multisignatures and identity-based sequential aggregate signatures, with applications to secure routing. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 276–285, 2007.
- [BHH⁺15] D. J. Bernstein, D. Hopwood, A. Hülsing, T. Lange, R. Niederhagen, L. Pappachristodoulou, M. Schneider, P. Schwabe, and Z. Wilcox-O’Hearn. SPHINCS: Practical stateless hash-based signatures. In *Advances in Cryptology – EUROCRYPT ’15*, pages 368–397, 2015.
- [BHK⁺19] D. J. Bernstein, A. Hülsing, S. Kölbl, R. Niederhagen, J. Rijneveld, and P. Schwabe. The SPHINCS⁺ signature framework. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, pages 2129–2146. ACM, 2019.
- [BK20] D. Boneh and S. Kim. One-time and interactive aggregate signatures from lattices. Available at https://crypto.stanford.edu/~skim13/agg_ots.pdf, 2020.
- [BKK⁺15] O. Blazy, S. A. Kakvi, E. Kiltz, , and J. Pan. Tightly-secure signatures from chameleon hash functions. In *Public Key Cryptography – PKC ’15*, pages 256–279, 2015.
- [BN06] M. Bellare and G. Neven. Multi-signatures in the plain public-key model and a general forking lemma. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 390–399, 2006.
- [Bol03] A. Boldyreva. Threshold signatures, multisignatures and blind signatures based on the Gap-Diffie-Hellman-group signature scheme. In *Public Key Cryptography – PKC ’03*, pages 31–46, 2003.
- [BR93] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 62–73, 1993.

- [BS08] M. Bellare and S. Shoup. Two-tier signatures from the Fiat-Shamir transform, with applications to strongly unforgeable and one-time signatures. *IET Information Security*, 2(2):47–63, 2008.
- [BSBH⁺18] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev. Scalable, transparent, and post-quantum secure computational integrity. Cryptology ePrint Archive, Paper 2018/046, 2018.
- [BSM⁺91] M. Blum, A. D. Santis, S. Micali, and G. Persiano. Non-interactive zero-knowledge. *SIAM Journal on Computing*, 20(6):1084–1118, 1991.
- [BTT22] C. Boschini, A. Takahashi, and M. Tibouchi. MuSig-L: Lattice-based multi-signature with single-round online phase. In *Advances in Cryptology – CRYPTO ’22*, pages 276–305, 2022.
- [BWG⁺21] V. Buterin, Y. Weiss, K. Gazso, N. Patel, D. Tirosh, S. Nacson, and T. Hess. EIP-4337: Account abstraction using Alt mempool [DRAFT]. Ethereum Improvement Proposals, no. 4337 (available at <https://eips.ethereum.org/EIPS/eip-4337>), 2021.
- [BY96] M. Bellare and M. Yung. Certifying permutations: Noninteractive zero-knowledge based on any trapdoor permutation. *Journal of Cryptology*, 9(3):149–166, 1996.
- [CD98] R. Cramer and I. Damgård. Zero-knowledge proofs for finite field arithmetic; or: Can zero-knowledge be for free? In *Advances in Cryptology – CRYPTO ’98*, pages 424–441, 1998.
- [CFG⁺15] D. Catalano, D. Fiore, R. Gennaro, and K. Vamvourellis. Algebraic (trapdoor) one-way functions: Constructions and applications. *Theoretical Computer Science*, 592:143–165, 2015.
- [CHM⁺20] A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In *Advances in Cryptology – EUROCRYPT ’20*, pages 738–768, 2020.
- [CL18] R. Canetti and A. Lichtenberg. Certifying trapdoor permutations, revisited. In *Proceedings of the 16th Theory of Cryptography Conference*, pages 476–506, 2018.
- [DEF⁺19] M. Drijvers, K. Edalatnejad, B. Ford, E. Kiltz, J. Loss, G. Neven, and I. Stepanovs. On the security of two-round multi-signatures. In *IEEE Symposium on Security and Privacy*, pages 1084–1101, 2019.
- [DOT⁺22] I. Damgård, C. Orlandi, A. Takahashi, and M. Tibouchi. Two-round n -out-of- n and multi-signatures and trapdoor commitment from lattices. *Journal of Cryptology*, 35(2):14, 2022.
- [FH21] M. Fukumitsu and S. Hasegawa. A tightly secure DDH-based multisignature with public-key aggregation. *International Journal of Networking and Computing*, 11(2):319–337, 2021.
- [FKL18] G. Fuchsbauer, E. Kiltz, and J. Loss. The algebraic group model and its applications. In *Advances in Cryptology – CRYPTO ’18*, pages 33–62, 2018.

- [FLS90] U. Feige, D. Lapidot, and A. Shamir. Multiple non-interactive zero knowledge proofs based on a single random string. In *Proceedings of the 31st Annual IEEE Symposium on Foundations of Computer Science*, pages 308–317, 1990.
- [FPS20] G. Fuchsbauer, A. Plouviez, and Y. Seurin. Blind Schnorr signatures and signed ElGamal encryption in the algebraic group model. In *Advances in Cryptology – EUROCRYPT ’20*, pages 63–95, 2020.
- [FS86] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology – CRYPTO ’86*, pages 186–194, 1986.
- [FSZ22] N. Fleischhacker, M. Simkin, and Z. Zhang. Squirrel: Efficient synchronized multi-signatures from lattices. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1109–1123, 2022.
- [GLS⁺04] R. Gennaro, D. Leigh, R. Sundaram, and W. S. Yezauris. Batching Schnorr identification scheme with applications to privacy-preserving authorization and low-bandwidth communication devices. In *Advances in Cryptology – ASIACRYPT ’04*, pages 276–292, 2004.
- [GLS⁺23] A. Golovnev, J. Lee, S. Setty, J. Thaler, and R. S. Wahby. Brakedown: Linear-time and field-agnostic SNARKs for R1CS. In *Advances in Cryptology – CRYPTO ’23*, pages 193–226, 2023.
- [Gol01] O. Goldreich. *Foundations of Cryptography – Volume 1: Basic Techniques*. Cambridge University Press, 2001.
- [Gol11] O. Goldreich. In a world of $P=BPP$. In *Studies in Complexity and Cryptography*, pages 191–232. Springer, 2011.
- [GOS06] J. Groth, R. Ostrovsky, and A. Sahai. Non-interactive zaps and new techniques for NIZK. In *Advances in Cryptology – CRYPTO ’06*, pages 97–111, 2006.
- [GR13] O. Goldreich and R. D. Rothblum. Enhancements of trapdoor permutations. *Journal of Cryptology*, 26(3):484–512, 2013.
- [Gro06] J. Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In *Advances in Cryptology – ASIACRYPT ’06*, pages 444–459, 2006.
- [Gro10] J. Groth. Short non-interactive zero-knowledge proofs. In *Advances in Cryptology – ASIACRYPT ’10*, pages 341–358, 2010.
- [Gro16] J. Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology – EUROCRYPT ’16*, pages 305–326, 2016.
- [GS08] J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Advances in Cryptology – EUROCRYPT ’08*, pages 415–432, 2008.
- [GW11] C. Gentry and D. Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 99–108, 2011.

- [GWC19] A. Gabizon, Z. J. Williamson, and O. Ciobotaru. PLONK: permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. *Cryptology ePrint Archive*, Paper 2019/953, 2019.
- [HK22] A. Hülsing and M. A. Kudinov. Recovering the tight security proof of SPHINCS⁺. *Cryptology ePrint Archive*, Paper 2022/346, 2022.
- [IN83] K. Itakura and K. Nakamura. A public-key cryptosystem suitable for digital multisignatures. *NEC Research & Development*, 71:1–8, 1983.
- [IR89] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 44–61, 1989.
- [KHR⁺22] M. A. Kudinov, A. Hülsing, E. Ronen, and E. Yogev. SPHINCS+C: Compressing SPHINCS+ with (almost) no cost. *Cryptology ePrint Archive*, Paper 2022/778, 2022.
- [KL21] J. Katz and Y. Lindell. *Introduction to Modern Cryptography* (3rd Edition). CRC Press, 2021.
- [KMP16] E. Kiltz, D. Masny, and J. Pan. Optimal security proofs for signatures from identification schemes. In *Advances in Cryptology – CRYPTO ’16*, pages 33–61, 2016.
- [Lam79] L. Lamport. Constructing digital signatures from a one way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, 1979.
- [LHL94] C. Li, T. Hwang, and N. Lee. Threshold-multisignature schemes where suspected forgery implies traceability of adversarial shareholders. In *Advances in Cryptology – EURO-CRYPT ’94*, pages 194–204, 1994.
- [Lin06] Y. Lindell. A simpler construction of CCA2-secure public-key encryption under general assumptions. *Journal of Cryptology*, 19(3):359–377, 2006.
- [LK23] K. Lee and H. Kim. Two-round multi-signatures from Okamoto signatures. *Mathematics*, 11(14), 2023.
- [LOS⁺06] S. Lu, R. Ostrovsky, A. Sahai, H. Shacham, and B. Waters. Sequential aggregate signatures and multisignatures without random oracles. In *Advances in Cryptology – EUROCRYPT ’06*, pages 465–485, 2006.
- [LYG19] D. Le, G. Yang, and A. A. Ghorbani. A new multisignature scheme with public key aggregation for blockchain. In *Proceedings of the 17th International Conference on Privacy, Security and Trust*, pages 1–7, 2019.
- [Mic00] S. Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000.
- [MOR01] S. Micali, K. Ohta, and L. Reyzin. Accountable-subgroup multisignatures: extended abstract. In *Proceedings of the 8th ACM Conference on Computer and Communications Security (CCS)*, pages 245–254, 2001.
- [MPS⁺19] G. Maxwell, A. Poelstra, Y. Seurin, and P. Wuille. Simple Schnorr multi-signatures with applications to Bitcoin. *Designs, Codes and Cryptography*, 87(9):2139–2164, 2019.

- [MTT19] T. Mizuide, A. Takayasu, and T. Takagi. Tight reductions for Diffie-Hellman variants in the algebraic group model. In *Topics in Cryptology – CT-RSA ’19*, pages 169–188, 2019.
- [Nak09] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://www.bitcoin.org/bitcoin.pdf>, 2009.
- [Nao91] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991.
- [NRS⁺20] J. Nick, T. Ruffing, Y. Seurin, and P. Wuille. MuSig-DN: Schnorr multi-signatures with verifiably deterministic nonces. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1717–1731, 2020.
- [NRS21] J. Nick, T. Ruffing, and Y. Seurin. MuSig2: Simple two-round Schnorr multi-signatures. In *Advances in Cryptology – CRYPTO ’21*, pages 189–221, 2021.
- [NY89] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 33–43, 1989.
- [OO91] K. Ohta and T. Okamoto. A digital multisignature scheme based on the Fiat-Shamir scheme. In *Advances in Cryptology – ASIACRYPT ’91*, pages 139–148, 1991.
- [PS00] D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13:361–396, 2000.
- [PW23] J. Pan and B. Wagner. Chopsticks: Fork-free two-round multi-signatures from non-interactive assumptions. In *Advances in Cryptology – EUROCRYPT ’23*, pages 597–627, 2023.
- [RS20] L. Rotem and G. Segev. Algebraic distinguishers: From discrete logarithms to decisional Uber assumptions. In *Proceedings of the 18th Theory of Cryptography Conference*, pages 366–389, 2020.
- [RY07] T. Ristenpart and S. Yilek. The power of proofs-of-possession: Securing multiparty signatures against rogue-key attacks. In M. Naor, editor, *Advances in Cryptology – EUROCRYPT ’07*, pages 228–245, 2007.
- [Sah99] A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 543–553, 1999.
- [Sch91] C. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.
- [SCO⁺01] A. D. Santis, G. D. Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In *Advances in Cryptology – CRYPTO ’01*, pages 566–598, 2001.
- [Sho97] V. Shoup. Lower bounds for discrete logarithms and related problems. In *Advances in Cryptology – EUROCRYPT ’97*, pages 256–266, 1997.

- [Sim98] D. R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Advances in Cryptology – EUROCRYPT ’98*, pages 334–345, 1998.
- [Sta23] StarkWare. Account abstraction: Improving security and user experience for mainstream crypto adoption. Available at <https://medium.com/starkware/account-abstraction-improving-security-and-user-experience-for-mainstream-crypto-adoption-eb57cb09023>, 2023.
- [TSS⁺23] K. Takemure, Y. Sakai, B. Santoso, G. Hanaoka, and K. Ohta. More efficient two-round multi-signature scheme with provably secure parameters. *Cryptology ePrint Archive*, Paper 2023/155, 2023.
- [TZ23] S. Tessaro and C. Zhu. Threshold and multi-signature schemes from linear hash functions. In *Advances in Cryptology – EUROCRYPT ’23*, pages 628–658, 2023.
- [XZS22] T. Xie, Y. Zhang, and D. Song. Orion: Zero knowledge proof with linear prover time. In *Advances in Cryptology – CRYPTO ’22*, pages 299–328, 2022.

A Extension: A t -Time Multi-Signature Scheme

In this section we show that our one-time multi-signature scheme, described in Section 4, naturally extends to a t -time multi-signature scheme without introducing any additional assumptions.

The t -time scheme’s verification and aggregated verification keys consist of $t + 1$ group elements, whereas signatures still consist of a single group element. At a high level, the extension relies on ideas developed originally for batch identification schemes, such as the batch Schnorr identification scheme constructed by Gennaro, Leigh, Sundaram and Yerazunis [GLS⁺04].

For extending our one-time scheme, the key-generation algorithm invoked by each signer samples $x^{(1)}, \dots, x^{(t)}, r \leftarrow \mathcal{X}_{\text{Ind}}$, computes $X^{(1)} = \text{F}_{\text{Ind}}(x^{(1)}), \dots, X^{(t)} = \text{F}_{\text{Ind}}(x^{(t)}), R = \text{F}_{\text{Ind}}(r)$, and then outputs

$$(\text{sk}, \text{vk}) = \left((x^{(1)}, \dots, x^{(t)}, r), (X^{(1)}, \dots, X^{(t)}, R) \right).$$

For aggregating n verification keys, a vector $\vec{\text{vk}} = (\text{vk}_1, \dots, \text{vk}_n)$ of verification keys, where $\text{vk}_i = (X_i^{(1)}, \dots, X_i^{(t)}, R_i)$ for every $i \in [n]$, is aggregated by computing

$$\begin{aligned} (a_1, \dots, a_n) &= \text{H}_1(\vec{\text{vk}}) \in \mathbb{K}_{\text{Ind}}^n \\ \text{agg}X^{(j)} &= \prod_{i=1}^n (X_i^{(j)})^{a_i} \in \mathcal{Y}_{\text{Ind}} \text{ for every } j \in [t] \\ \text{agg}R &= \prod_{i=1}^n R_i^{a_i} \in \mathcal{Y}_{\text{Ind}}, \end{aligned}$$

and outputting the aggregated verification key

$$\text{aggvk} = (\text{agg}X^{(1)}, \dots, \text{agg}X^{(t)}, \text{agg}R).$$

For signing a message m with respect to an aggregated verification key aggvk , each signer first computes $\beta = \text{H}_0(m, \text{aggvk})$ and then uses their signing key $\text{sk}_i = (x_i^{(1)}, \dots, x_i^{(t)}, r_i)$ for computing

a signature $\sigma_i = r_i + \sum_{j=1}^t \beta^j \cdot x_i^{(j)}$. The signatures $\sigma_1, \dots, \sigma_n$ are then aggregated exactly as in the one-time scheme 1Γ by computing $\text{agg}\sigma = \sum_{i=1}^n a_i \cdot \sigma_i$, and the verification algorithm verifies that

$$F_{\text{Ind}}(\text{agg}\sigma) = \text{agg}R \cdot \prod_{j=1}^t \left(\text{agg}X^{(j)} \right)^{\beta^j}.$$

In terms of correctness, letting $\beta = H_0(m, \text{aggvk})$, it holds that

$$\begin{aligned} F_{\text{Ind}}(\text{agg}\sigma) &= F_{\text{Ind}}\left(\sum_{i=1}^n a_i \cdot \sigma_i\right) \\ &= F_{\text{Ind}}\left(\sum_{i=1}^n a_i \cdot r_i + \sum_{j=1}^t \beta^j \cdot \sum_{i=1}^n a_i \cdot x_i^{(j)}\right) \\ &= \left(\prod_{i=1}^n F_{\text{Ind}}(r_i)^{a_i}\right) \cdot \prod_{j=1}^t \left(\prod_{i=1}^n F_{\text{Ind}}(x_i)^{a_i}\right)^{\beta^j} \\ &= \text{agg}R \cdot \prod_{j=1}^t \left(\text{agg}X^{(j)}\right)^{\beta^j}, \end{aligned}$$

and thus the scheme provides perfect correctness.

In terms of security, the extended analysis is obtained from the proof of Theorem 4.1 by following the exact same structure, while enabling the algorithm B to respond to A 's t signing queries. This is done via the following two main modifications. First, the algorithm B guesses the indices of A 's H_0 -queries $\{(m_i, \text{aggvk}_i)\}_{i \in [t]}$ and (m^*, aggvk^*) , where m_1, \dots, m_t are the messages for which A issues signing queries, and $\text{aggvk}_1, \dots, \text{aggvk}_t$ are the corresponding aggregated verification keys. This leads to a security loss of $q_{H_0}^{t+1}$, instead of $q_{H_0}^2$ as in the one-time scheme (similar to the loss in the batch protocol of Gennaro et al. [GLS⁺04]), which is still polynomial for any constant $t \geq 1$. Second, given $X = F_{\text{Ind}}(x)$, the algorithm B needs to define the honest verification key $\text{vk} = (X^{(1)}, \dots, X^{(t)}, R)$ that will be provided to the algorithm A . This is done by sampling responses $\beta_1, \dots, \beta_t \leftarrow \mathbb{K}_{\text{Ind}}$ that will be provided by B to A 's H_0 -queries that correspond to the t signing queries, as well as signatures $\sigma_1, \dots, \sigma_t \leftarrow \mathcal{X}_{\text{pp}}$ that will be provided by B as responses to A 's signing queries. Letting $X^{(1)} = F_{\text{Ind}}(x^{(1)}), \dots, X^{(t)} = F_{\text{Ind}}(x^{(t)})$, and $R = F_{\text{Ind}}(r)$ for unknowns $x^{(1)}, \dots, x^{(t)}$ and r , the following linear equations must be satisfied:

$$\begin{pmatrix} \sigma_1 \\ \vdots \\ \sigma_t \end{pmatrix} = \begin{pmatrix} 1 & \beta_1 & \dots & \beta_1^t \\ \vdots & & \ddots & \\ 1 & \beta_t & \dots & \beta_t^t \end{pmatrix} \cdot \begin{pmatrix} r \\ x^{(1)} \\ \vdots \\ x^{(t)} \end{pmatrix}$$

Note that whenever the values β_1, \dots, β_t are distinct, then the above matrix has full rank. Therefore, for any fixing of one of the $x^{(i)}$'s to unknown value x , there is a unique (and linear) solution for the remaining $x^{(j)}$'s and r . Relying on this fact, the algorithm B uniformly samples an index $i \leftarrow [t]$, sets $X^{(i)} = X$, and can then compute the values of the remaining $X^{(j)}$'s and R in a linearly homomorphic manner given X , $\sigma_1, \dots, \sigma_t$ and β_1, \dots, β_t (note that the case $t = 1$ corresponds to the proof of Theorem 4.1).

B From One-Time Multi-Signatures to Collision-Resistant Hashing

In this section we show that a collision-resistant hash function can be constructed in a black-box manner from any one-time multi-signature scheme satisfying a natural property (which is satisfied by our one-time multi-signature scheme in Section 4). The property asks that the scheme’s key-generation algorithm defines an injective mapping of its randomness to its verification keys. We refer to this property as “randomness-injective key generation”, and an interesting open problem is to study the extent to which this property is essential in this context. Given that one-way functions cannot be used in a black-box manner to construct a collision-resistant hash function [Sim98], this rules out various natural constructions of one-time multi-signature schemes based on one-way functions.

Note that, in addition to randomness-injective key generation, one must also always assume that the scheme’s key-aggregation algorithm produces rather short keys. Otherwise, any one-time single-signer scheme (which can be constructed based on any one-way function [Lam79]) can be used to construct a one-time multi-signer scheme in a black-box manner by concatenating individual verification keys for producing an aggregated verification key.

Our result shows that, unlike the case of standard (i.e., single-user) one-time signatures, a one-time multi-signature scheme with short keys (and randomness-injective key generation) cannot be constructed based on any one-way function in a black-box manner [IR89]. This justifies, in some sense, the fact that our one-time multi-signature scheme is based on an assumption that seems stronger than the assumption that one-way functions exist. Specifically, our one-time multi-signature scheme is based on any ring-homomorphic one-way function, which is a significantly more structured primitive when compared to plain one-way functions.

In what follows we describe the construction and prove its collision resistance based on the security of the underlying one-time multi-signature scheme. Let $1T = (1T.Setup, 1T.KG, 1T.KAgg, 1T.Sign, 1T.SAgg, 1T.Verify)$ be a one-time multi-signature scheme over the message space $\mathcal{M} = \{0, 1\}$. We denote by $\ell = \ell(\lambda)$ the bit-length of the randomness provided to the key-generation algorithm $1T.KG$.

The function family $CRH = (KG, Eval)$

KG(1^λ). On input 1^λ the key-generation algorithm computes

$$\begin{aligned} pp &\leftarrow 1T.Setup(1^\lambda) \\ (sk_0, vk_0) &\leftarrow 1T.KG(1^\lambda) \end{aligned}$$

and outputs $k = (pp, vk_0)$.

Eval(k, x_1, \dots, x_n). On input $k = (pp, vk_0)$ and $x = (x_1, \dots, x_n) \in (\{0, 1\}^\ell)^n$ the evaluation algorithm computes

$$\begin{aligned} (sk_i, vk_i) &\leftarrow 1T.KG(pp; x_i) \text{ for all } i \in [n] \\ y &= 1T.KAgg(pp, vk_0, vk_1, \dots, vk_n) \end{aligned}$$

and outputs y .

In terms of compression, note that if there exists any $n = n(\lambda)$ for which the scheme’s key-aggregation algorithm aggregates $n + 1$ keys into an aggregated key of length less than $n \cdot \ell$ bits, then the function family CRH is compressing. This holds, in particular if the length of aggregated keys is independent of the number of signers, as in our one-time multi-signature scheme from Section 4.

The following theorem captures the collision resistance of the function family CRH based on the security of the one-time multi-signature scheme $1T$. As discussed above, the theorem assumes that,

for the key-generation algorithm 1T.KG , different random strings lead to different verification keys. Formally, we assume that for any $\lambda \in \mathbb{N}$, for any pp generated by $1\text{T.Setup}(1^\lambda)$ and for any $x \neq x' \in \{0, 1\}^{\ell(\lambda)}$ it holds that $\text{vk} \neq \text{vk}'$, where $(\text{sk}, \text{vk}) = 1\text{T.KG}(\text{pp}; x)$ and $(\text{sk}', \text{vk}') = 1\text{T.KG}(\text{pp}; x')$. We refer to this property as “randomness-injective key generation”, and note once again that it is satisfied in particular by our one-time multi-signature scheme from Section 4.

Theorem B.1. *Assuming that the scheme 1T has randomness-injective key generation, then for any probabilistic polynomial-time algorithm A there exists a probabilistic polynomial-time algorithm B such that for every $\lambda \in \mathbb{N}$ it holds that*

$$\text{Adv}_{\text{CRH},A}^{\text{crh}}(\lambda) \leq \text{Adv}_{1\text{T},B}^{1\text{TimeMS}}(\lambda).$$

Proof of Theorem B.1. Let A be a probabilistic polynomial-time algorithm that participates in the collision-resistance experiment $\text{Exp}_{\text{CRH},A}^{\text{crh}}$, and consider the algorithm B that participates in the experiment $\text{Exp}_{1\text{T},B}^{1\text{TimeMS}}$ as follows. On input (pp, vk_0) , the algorithm B invokes the algorithm A on input $\mathbf{k} = (\text{pp}, \text{vk}_0)$ to obtain $x = (x_1, \dots, x_n)$ and $x' = (x'_1, \dots, x'_n)$. If $x = x'$ or $\text{CRH.Eval}(\mathbf{k}, x) = \text{CRH.Eval}(\mathbf{k}, x')$ then B outputs \perp and halts. Otherwise, B computes

$$\begin{aligned} (\text{sk}_i, \text{vk}_i) &\leftarrow 1\text{T.KG}(\text{pp}; x_i) \text{ for all } i \in [n] \\ \vec{\text{vk}} &= (\text{vk}_0, \text{vk}_1, \dots, \text{vk}_n) \\ \sigma_i &= 1\text{T.Sign}(\text{pp}, \text{sk}_i, \vec{\text{vk}}, 0) \text{ for all } i \in [n] \\ (\text{sk}'_i, \text{vk}'_i) &\leftarrow 1\text{T.KG}(\text{pp}; x'_i) \text{ for all } i \in [n] \\ \vec{\text{vk}}' &= (\text{vk}_0, \text{vk}'_1, \dots, \text{vk}'_n), \end{aligned}$$

and queries the signing oracle with $(\vec{\text{vk}}, 0)$ to obtain a signature σ_0 . Then, the algorithm B outputs

$$\begin{aligned} \vec{\text{vk}}^* &= \vec{\text{vk}}' \\ m^* &= 0 \\ \text{agg}\sigma^* &= 1\text{T.SAgg}(\text{pp}, (\sigma_0, \sigma_1, \dots, \sigma_n)). \end{aligned}$$

For analyzing the success probability of the algorithm B , note that $\text{CRH.Eval}(\mathbf{k}, x) = \text{CRH.Eval}(\mathbf{k}, x')$ implies $1\text{T.KAgg}(\text{pp}, \vec{\text{vk}}) = 1\text{T.KAgg}(\text{pp}, \vec{\text{vk}}')$, and therefore the correctness of the scheme 1T implies that

$$\begin{aligned} &1\text{T.Verify}(\text{pp}, 1\text{T.KAgg}(\text{pp}, \vec{\text{vk}}^*), m^*, \text{agg}\sigma^*) \\ &= 1\text{T.Verify}(\text{pp}, 1\text{T.KAgg}(\text{pp}, \vec{\text{vk}}'), 0, 1\text{T.SAgg}(\text{pp}, (\sigma_0, \sigma_1, \dots, \sigma_n))) \\ &= 1\text{T.Verify}(\text{pp}, 1\text{T.KAgg}(\text{pp}, \vec{\text{vk}}), 0, 1\text{T.SAgg}(\text{pp}, (\sigma_0, \sigma_1, \dots, \sigma_n))) \\ &= 1. \end{aligned}$$

In addition, the description of B guarantees that $\text{vk}_0 \in \vec{\text{vk}}^*$, and the assumption that 1T has randomness-injective key generation together with the fact that $x \neq x'$ guarantee that $\vec{\text{vk}} \neq \vec{\text{vk}}'$, and therefore B did not query the signing oracle with $(\vec{\text{vk}}^*, m^*)$. Therefore, it holds that

$$\text{Adv}_{1\text{T},B}^{1\text{TimeMS}}(\lambda) \geq \text{Adv}_{\text{CRH},A}^{\text{crh}}(\lambda).$$

■