

History-Free Sequential Aggregation of Hash-and-Sign Signatures

Alessio Meneghetti¹ and Edoardo Signorini^{2,3}

¹ University of Trento, Trento, Italy

² Telsy, Turin, Italy

³ Politecnico di Torino, Turin, Italy

Abstract. A sequential aggregate signature (SAS) scheme allows multiple users to sequentially combine their respective signatures in order to reduce communication costs. Historically, early proposals required the use of trapdoor permutation (e.g., RSA). In recent years, a number of attempts have been made to extend SAS schemes to post-quantum assumptions. Many post-quantum signatures have been proposed in the hash-and-sign paradigm, which requires the use of trapdoor functions and appears to be an ideal candidate for sequential aggregation attempts. However, the hardness in achieving post-quantum one-way permutations makes it difficult to obtain similarly general constructions. Direct attempts at generalizing permutation-based schemes have been proposed, but they either lack formal security or require additional properties on the trapdoor function, which are typically not available for multivariate or code-based functions. In this paper, we propose a (partial-signature) history-free SAS within the probabilistic hash-and-sign with retry paradigm, generalizing existing techniques to generic trapdoor functions. We prove the security of our scheme in the random oracle model and we instantiate our construction with three post-quantum schemes, comparing their compression capabilities. Finally, we discuss how direct extensions of permutation-based SAS schemes are not possible without additional properties, showing the lack of security of two existing multivariate schemes.

Keywords: sequential aggregate signature · post-quantum cryptography · hash-and-sign

1 Introduction

An Aggregate Signature (AS) scheme allows n users to combine their individual signatures on separate messages to produce a single, directly verifiable aggregate signature. This approach aims to achieve shorter signature lengths compared to trivial concatenation of individual signatures. Aggregate signatures have interesting applications in various network scenarios with high communication costs, such as PKI certificate chains or secure routing protocols authentication. The concept of aggregate signatures was initially introduced in a seminal paper by Boneh, Gentry, Lynn, and Shacham [11]. They proposed a method that allows

any participant to aggregate signatures from distinct users using a public aggregation algorithm. Although this general aggregation approach is efficient and valuable, it is limited to the use of bilinear pairings. A restricted variant of aggregate signatures, known as Sequential Aggregated Signature (SAS), was introduced by Lysyanskaya, Micali, Reyzin, and Shacham [31]. In SAS schemes, signatures are aggregated in a specific sequence, starting from the so-far aggregated signature and possibly from the public key and message information of previous users. The sequential structure is still beneficial in many applications. Numerous works have been pursued in this direction, proposing constructions based on trapdoor permutations [31,32,13,23] or the use of bilinear pairings [30,4,21]. Additionally, it is possible to consider further aggregation variants such as the synchronous model [25,1], in which signers cannot produce more than one signature within a certain time interval, or the interactive model [5], in which signers participate in an interactive multi-round protocol to produce a signature.

Aggregate Signatures from Post-Quantum Assumptions In recent years, there has been a growing interest in post-quantum signatures, leading researchers to explore AS schemes in this field. Numerous Lattice-based schemes have been proposed, with generic solutions based on non-interactive arguments [18,2], results in the synchronous model [22] and sequential aggregations both in the Fiat-Shamir [12] and Hash-and-Sign paradigms [19,35]. Focusing on the latter, both [19,35] have extended previous trapdoor permutation-based approaches [32,23], but their security relies on the collision-resistance property of lattice trapdoor Preimage Sampleable Functions (PSF) [24]. These additional properties are not available for generic trapdoor functions employed, for instance, in multivariate-quadratic-based (MQ-based) or code-based signature schemes. Currently, SAS schemes based on these assumptions are very limited, with only two existing MQ-based schemes [20,15], which follow the construction of [31]. Unfortunately, both [20,15] lack formal security and there are instances of the underlying function for which they are insecure, as outlined below.

Our Contribution In this work, we address the extension of permutation-based SAS schemes to generic trapdoor functions, making them applicable to a broader range of post-quantum signatures. In Section 3, we present a partial-signature history-free SAS scheme based on generic trapdoor functions. In a history-free SAS, introduced in [13], signers receive only the so-far aggregated signature without requiring previous users' public keys and messages. The partial-signature variant, initially presented in [16], reduces the amount of information the signer needs to receive from the previous one, but requires a final (public) aggregation step.

Our approach can be seen as a generalization of the work of Brogle, Goldberg, and Reyzin [13] for trapdoor permutations, adapting the encoding technique of [32,19] to include trapdoor functions beyond permutations. The main novelty of our work is the extension of the probabilistic hash-and-sign with retry paradigm, which is common in post-quantum signature constructions, to sequential aggregations.

gate signature schemes. As a result, we are able to reduce the security of our scheme to the one-wayness of the trapdoor function and to an additional notion of Preimage Sampling (PS) indistinguishability. While this further notion may appear restrictive in the choice of trapdoor functions, it turns out to be quite natural in security proofs of post-quantum hash-and-sign schemes, as recently shown in [28].

In Section 4, we apply our scheme to MQ-based signature schemes, specifically UOV [27] and MAYO [8], and the code-based scheme Wave [17]. For each scheme, we evaluate its compression capabilities and review its PS security so that it can be covered in our security proof.

Finally, in Section 5, we argue that the simpler approaches of [31,32] are not viable for generic trapdoor functions. As evidence, we show how two existing MQ-based aggregate signature schemes [20,15] are universally forgeable when instantiated with UOV and discuss their lack of provable security.

2 Notation and Preliminaries

For $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$. For a finite set X , we write $|X|$ for the cardinality of X and $\text{len}(X)$ for the bit size of an element in X . By $x \leftarrow_s X$, we denote the sample of the element x from $U(X)$, the uniform distribution over X . For an algorithm A , we write $x \leftarrow A(y)$ to denote the assignment of x to the output of A on input y . We denote with \mathcal{A} a (probabilistic) polynomial-time adversary. For an adversary \mathcal{A} and a function F , we write $x \leftarrow \mathcal{A}^{\text{OF}}$ the assignment of x of the output of \mathcal{A} with oracle access to F . For two bit strings $x, y \in \{0, 1\}^*$, we denote by $x \parallel y$ the bit string obtained by their concatenation. We write \mathbb{F}_q for a finite field of q elements. We denote by $\mathbb{F}_q^{m \times n}$ the set of matrices over \mathbb{F}_q with m rows and n columns. $\mathbf{I}_{n \times n}$ is the identity matrix of size n . $\mathbf{0}_{m \times n}$ is the $m \times n$ zero matrix and $\mathbf{0}_n$ is the zero vector in \mathbb{F}_q^n . In the remainder of this section, we introduce standard definitions and notions related to digital signature schemes based on trapdoor functions.

2.1 Trapdoor Functions

Definition 1. A trapdoor function (TDF) T is a tuple of four algorithms $(\text{TrapGen}, F, \mathsf{I}, \text{SampDom})$:

- $\text{TrapGen}(1^\lambda)$: takes as input a security parameter 1^λ and generates an efficiently computable function $F: \mathcal{X} \rightarrow \mathcal{Y}$ and a trapdoor I that allow to invert F .
- $F(x)$: takes as input $x \in \mathcal{X}$ and outputs $F(x) \in \mathcal{Y}$.
- $\mathsf{I}(y)$: takes as input $y \in \mathcal{Y}$ and outputs $x \in \mathcal{X}$ such that $F(x) = y$ or it fails by returning \perp .
- $\text{SampDom}(F)$ takes as input a function $F: \mathcal{X} \rightarrow \mathcal{Y}$ and outputs $x \in \mathcal{X}$.

We define the following notion of one-wayness (OW) for a trapdoor function.

Definition 2. Let $T = (\text{TrapGen}, F, I, \text{SampDom})$ be a TDF and let \mathcal{A} be an adversary. We define the advantage of \mathcal{A} playing the OW game against T as

$$\text{Adv}_T^{\text{OW}}(\mathcal{A}) = \Pr \left[F(x) = y \mid \begin{array}{l} (F, I) \leftarrow \text{TrapGen}(1^\lambda) \\ y \leftarrow \mathcal{Y} \\ x \leftarrow \mathcal{A}(F, y) \end{array} \right]$$

Sometimes, the notion of one-wayness requires that the challenge is obtained as $F(x)$ with x sampled following some distribution on \mathcal{X} . The use of uniformly random challenges appears necessary to prove the security of hash-and-sign schemes if the image via F of random input is not uniformly distributed in the codomain.

Definition 3 (Trapdoor Permutation (TDP)). A TDF $T = (\text{TrapGen}, F, I, \text{SampDom})$ is said to be a TDP if F and I are permutations.

2.2 Digital Signatures

A digital signature scheme Sig is a tuple of three algorithms ($\text{KGen}, \text{Sign}, \text{Vrfy}$):

- $\text{KGen}(1^\lambda)$: takes as input a security parameter 1^λ and generates a key pair (pk, sk) .
- $\text{Sign}(\text{sk}, m)$: takes as input a signing key sk and a message m and returns a signature σ .
- $\text{Vrfy}(\text{pk}, m, \sigma)$: takes as input a verification key pk , a message m and a signature σ and returns 1 for acceptance or 0 for rejection.

We define the standard notion of existential unforgeability against chosen-message attack (EUF-CMA).

Definition 4 (EUF-CMA security). Let \mathcal{O} be a random oracle, let $\text{Sig} = (\text{KGen}, \text{Sign}, \text{Vrfy})$ be a signature scheme, let \mathcal{A} be an adversary. We define the advantage of \mathcal{A} playing the EUF-CMA game against Sig in the random oracle model as:

$$\text{Adv}_{\text{Sig}}^{\text{EUF-CMA}}(\mathcal{A}) = \Pr \left[\begin{array}{l} \text{Vrfy}(\text{pk}, m, \sigma) = 1 \\ \text{OSign}(\text{sk}, \cdot) \text{ not queried on } m \end{array} \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{KGen}(1^\lambda) \\ (m, \sigma) \leftarrow \mathcal{A}^{\mathcal{O}, \text{OSign}(\text{sk}, \cdot)}(\text{pk}) \end{array} \right]$$

2.3 Hash-and-Sign Schemes

The (probabilistic) hash-and-sign (HaS) paradigm is a standard approach to building digital signature schemes in the random oracle model from a trapdoor function T and a hash function $H: \{0, 1\}^* \rightarrow \mathcal{Y}$. To sign a message m , a signer with secret key $\text{sk} = I$ applies the hash function, modeled as a random oracle, to the message $y \leftarrow H(m)$ and computes its inverse $x \leftarrow I(y)$ through the secret trapdoor. In some scenarios, the HaS paradigm requires using a random string r , which acts as a salt for the hash function, i.e., $y \leftarrow H(m \parallel r)$. The resulting

Algorithm 1: Hash-and-sign with retry	
KGen (1^λ): 1: $(F, l) \leftarrow \text{TrapGen}(1^\lambda)$ 2: return (F, l) Vrfy ($F, m, (r, x)$): 1: return $F(x) = H(r, m)$	Sign (l, m): 1: repeat 2: $r \leftarrow_s \{0, 1\}^\lambda$ 3: $x \leftarrow l(H(r, m))$ 4: until $x \neq \perp$ 5: return (r, x)

Game 1: PS_b	
1: $(F, l) \leftarrow \text{TrapGen}(1^\lambda)$ 2: $b^* \leftarrow \mathcal{A}^{\text{Sample}_b}(F)$ 3: return $b^* \in \{0, 1\}$ Sample ₁ : 1: $r_i \leftarrow_s \{0, 1\}^\lambda$ 2: $x_i \leftarrow \text{SampDom}(F)$ 3: return (r_i, x_i)	Sample ₀ : 1: repeat 2: $r_i \leftarrow_s \{0, 1\}^\lambda$ 3: $y_i \leftarrow_s \mathcal{Y}$ 4: $x_i \leftarrow l(y_i)$ 5: until $x_i \neq \perp$ 6: return (r_i, x_i)

signature is the couple $\sigma = (x, r)$. A verifier uses the corresponding public key $\text{pk} = F$ to verify whether $F(x) = H(m \parallel r)$.

When T is a trapdoor permutation, this construction is known as Full Domain Hash, and the EUF-CMA security of the signature scheme can be proved from the one-wayness assumption of T [6]. For generic TDF, a black-box security proof is not known, and custom reductions are needed for different constructions. This becomes particularly significant in the post-quantum setting, where no constructions of one-way permutations are known. In order to achieve a secure signature scheme, it is possible to consider trapdoor functions with additional properties. For instance, Preimage Sampleable Functions (PSF) [24], which can be constructed from lattices [33], or Average Trapdoor PSF (ATPSF) [14], which can be constructed from code-based assumptions [17]. More generally, a slightly different paradigm, known as probabilistic hash-and-sign with retry (Algorithm 1), is used to prove the EUF-CMA security. With this approach, a random string r is sampled until a preimage for $H(m \parallel r)$ is found. The security is based on the one-wayness of the trapdoor function and on the additional condition that the output of the signing algorithm (r, x) is indistinguishable from a couple (r', x') with $r' \leftarrow_s \{0, 1\}^\lambda$ and $x' \leftarrow_s \text{SampDom}(F)$. Ad-hoc versions of this paradigm are commonly employed for MQ-based signatures, and have been utilized to prove the security of Unbalanced Oil and Vinegar (UOV), Hidden-Field Equation (HFE) [34], and MAYO [8] signature schemes.

In this work, we build a partial-signature history-free sequential aggregate signature scheme HaS-HF-SAS within the probabilistic hash-and-sign with retry paradigm. The security of our scheme requires the indistinguishability condition on preimages, which we formalize by adopting the following notion from [28].

Definition 5 (Preimage Sampling [28]). Let $\mathsf{T} = (\text{TrapGen}, \mathsf{F}, \mathsf{l}, \text{SampDom})$ be a TDF, let \mathcal{A} be an adversary. We define the advantage of \mathcal{A} playing the PS game (Game 1) against T as:

$$\text{Adv}_{\mathsf{T}}^{\text{PS}}(\mathcal{A}) = |\Pr[\text{PS}_0(\mathcal{A}) = 1] - \Pr[\text{PS}_1(\mathcal{A}) = 1]|$$

Note that the OW notion of Definition 2 includes trapdoor functions for which the probability that a pre-image exists via F for a random element $y \in \mathcal{Y}$ is negligible. Such functions could not be used as a building block for hash-and-sign schemes, as even knowledge of the secret trapdoor would not allow the computation of a pre-image. Nevertheless, a non-negligible failure probability could impact the tightness of a security reduction from OW. In the security proof of our scheme, we explicitly consider this possibility by bounding the number of queries made to the random oracle during a signature attempt.

2.4 History-Free Sequential Aggregate Signature

History-Free Sequential Aggregate Signatures (HF-SAS) were first introduced in [13,21] as a variant of the original Full-History construction of [31] that does not require knowledge of previous messages and public keys in the aggregation step.

Definition 6 (HF-SAS). A History-Free Sequential Aggregate Signature is a tuple of three algorithms $(\text{KGen}, \text{AggSign}, \text{AggVrfy})$:

- $\text{KGen}(1^\lambda)$: takes as input a security parameter 1^λ and generates a key pair (pk, sk) .
- $\text{AggSign}(\text{sk}_i, m_i, \Sigma_{i-1})$: takes as input the secret key sk_i and the message m_i of the i th user and the previous aggregate signature Σ_{i-1} . Returns an aggregate signature Σ_i .
- $\text{AggVrfy}(L_n, \Sigma_n)$: takes as input the full history $L_n = (\text{pk}_1, m_1), \dots, (\text{pk}_n, m_n)$ of public key, message pairs and an aggregate signature Σ_n . Returns 1 if Σ_n is a valid aggregate signature and 0 otherwise.

Every signer has a key pair $(\text{pk}_i, \text{sk}_i) \leftarrow \text{KGen}(1^\lambda)$. The signature aggregation process is done iteratively: the first signer with keys $(\text{pk}_1, \text{sk}_1)$ generates a signature Σ_1 for message m_1 with $\Sigma_1 \leftarrow \text{AggSign}(\text{sk}_1, m_1, \varepsilon)$, where ε represents the empty string to indicate that this is the first signature in the sequence. The i th signer with keys $(\text{pk}_i, \text{sk}_i)$ receives an aggregate signature Σ_{i-1} from the $(i-1)$ th signer and aggregate his signature on message m_i to obtain the aggregate signature $\Sigma_i \leftarrow \text{AggSign}(\text{sk}_i, m_i, \Sigma_{i-1})$. Note that the aggregation algorithm AggSign does not require the public keys and messages from the previous signers. Finally, the verifier can check the validity of the aggregate signature by running $\text{AggVrfy}(L_n, \Sigma_n)$.

SAS schemes were originally introduced by [31] for generic trapdoor permutation with the FDH approach. The history-free variant of [13] still requires trapdoor permutation, while [21] relies on bilinear pairing. The main intuition behind the aggregation process in TDP schemes is to “embed” the previous




Fig. 1. High level description of SAS scheme from [31].

Game 2: strong PS-HF-UF-CMA_S	
<ol style="list-style-type: none"> 1: $(pk^*, sk^*) \leftarrow KGen(1^\lambda)$ 2: $\mathcal{Q} \leftarrow \emptyset$ 3: $(L_n, \bar{\Sigma}_n) \leftarrow \mathcal{A}^{O, OAggSign}(pk^*)$ 4: $(pk_1, m_1), \dots, (pk_n, m_n) \leftarrow L_n$ 5: if $\nexists i^* : (pk_{i^*} = pk^* \wedge (m_{i^*}, s_{i^*}) \notin \mathcal{Q})$ then 6: return \perp 7: return $AggVrfy(L_n, \bar{\Sigma}_n)$ 	$OAggSign(m, \varrho)$: <ol style="list-style-type: none"> 1: $(\varrho', \varsigma') \leftarrow AggSign(sk^*, m, \varrho)$ 2: $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \varsigma')\}$ 3: return ϱ', ς'

aggregate signature into the new message to be signed. This ensures that the aggregate signature can be retrieved during the verification process, as depicted in Figure 1.

The main challenge in extending previous schemes to trapdoor functions that are not permutations lies in their lack of injectivity. This issue was addressed in [19] within the context of lattice-based signatures by employing an encoding technique derived from [32]. Subsequently, this idea was applied to MQ-based schemes instantiated with HFEv- [20] and UOV [15]. The proposed solution is to use a suitable encoding function, which splits the signature into two components. The first component can be injected into the codomain of the trapdoor function and subsequently made part of the computation of the aggregate signature, similar to the approach used in [31]. The second component is transmitted to the next signer and becomes part of the final aggregate signature. During the verification phase, this component is used to recover the partial aggregate signature through a corresponding decoding function. Notice that this method has a drawback in terms of the efficiency of the SAS scheme. In fact, storing part of the signature of each user without further aggregation causes a linear growth of the aggregate signature in the number of users. Therefore, it is currently unknown how to achieve sequential aggregate signatures of constant size in the post-quantum setting, where there are no one-way TDPs.

In the following, we define a slight modification of HF-SAS, as formalized in [12]. In this variant, the aggregation step requires only partial knowledge about the so-far aggregated signature. This description better captures the intuition behind the use of the encoding function and is better suited to our proposed

scheme. During each aggregation step, the signer produces a partial signature information, which will be sent to the next signer, along with a complementary component. At the end of the aggregation sequence, an additional `Combine` step is performed, potentially by a third party. This step combines all the complementary information and the last signature of the sequence, resulting in the complete aggregated signature.

Definition 7 (PS-HF-SAS). *A Partial-Signature History-Free Sequential Aggregate Signature is a tuple of four algorithms (KGen, AggSign, AggVrfy, Combine):*

- KGen and AggVrfy as described in Definition 6.
- AggSign(sk_i, m_i, ρ_{i-1}): takes as input the secret key sk_i and the message m_i of the i th user and a partial description ρ_{i-1} of the previous aggregate signature Σ_{i-1} . Computes an updated aggregate signature Σ_i and returns a partial description ρ_i and some complementary information ς_i .
- Combine($\varsigma_1, \dots, \varsigma_{n-1}, \Sigma_n$): takes as input the complementary information ς_i of the first $n-1$ signatures and the full description of the last signature Σ_n . Returns the complete description of the aggregate signature $\bar{\Sigma}_n$.

Below, we show the definition of partial-signature history-free unforgeability under adaptive chosen message (PS-HF-UF-CMA). In this model, the forger controls all signers’ private keys except for at least one honest signer. The forger can choose the keys of the rogue signers and adaptively query an aggregate signature oracle. Finally, to win the experiment, the forger must produce a valid, non-trivial aggregate signature involving the public key of the honest signer. A stronger notion is also considered in [12], where the adversary may produce forgery on messages already queried to the signing oracle, provided that the complementary part of the corresponding signature is distinct from the oracle’s response. We denote this variant as `strong PS-HF-UF-CMA`.

Definition 8 (PS-HF-UF-CMA Security). *Let \mathcal{O} be a random oracle, let $S = (\text{KGen}, \text{AggSign}, \text{AggVrfy}, \text{Combine})$ be a PS-HF-SAS scheme, let \mathcal{A} be an adversary. We define the advantage of \mathcal{A} playing the `strong PS-HF-UF-CMA` game (Game 2) against S as follows:*

$$\text{Adv}_S^{\text{PS-HF-UF-CMA}}(\mathcal{A}) = \Pr[\text{PS-HF-UF-CMA}_S(\mathcal{A}) = 1].$$

3 Sequential Aggregation of Hash-and-Sign Signatures

We present a PS-HF-SAS scheme following the probabilistic hash-and-sign with retry paradigm. The intuition behind our scheme closely follows the one of [32,13]. We use a two-step hash procedure: first, the signature x_{i-1} of the previous signer and the message m_i are contracted to a short value h_i , and then expanded to the codomain of the trapdoor function. The value h_i can be aggregated and made available to the verifier, who can expand it before knowing x_{i-1} . In this way, Neven [32] showed that the verifier does not need to check




Fig. 2. High level description of our HaS-HF-SAS scheme. The dashed arrows in the output represent the complementary part of the signature.

the validity of the signers' public keys. When calculating the first hash, we also require the signer to concatenate a random salt r_i , which will later be part of the aggregate signature. In [13], the salt is introduced to prevent a chosen message attack in the history-free setting. In our scheme, the use of the salt descends from the probabilistic hash-and-sign paradigm and provides a solution to overcome the technical challenges of using trapdoor functions that are not permutations. As discussed later in Section 5, attempting to remove the random salt would make the construction insecure even in the full-history setting.

A high-level description of the scheme HaS-HF-SAS is shown in Figure 2 while a detailed description is given in Algorithm 2. The properties of the underlying trapdoor function are as described in Section 2.1.

3.1 Security Proof

In the following, we prove the strong PS-HF-UF-CMA security of Algorithm 2.

Theorem 1. *Let T be a TDF. Let \mathcal{A} be a strong PS-HF-UF-CMA adversary against the HaS-HF-SAS scheme on T in the random oracle model, which runs in time t and makes q_S signing queries, q_H queries to the random oracle H and q_G queries to the random oracle G . Then, there exist a OW adversary \mathcal{B} against T that runs in time $t + \mathcal{O}((q_H + q_S + 1) \cdot \text{poly}(\text{len}(\mathcal{X}), \text{len}(\mathcal{Y})))$, and a PS adversary \mathcal{D} against T issuing q_S sampling queries that runs in time $t + \mathcal{O}(q_S \cdot \text{poly}(\text{len}(\mathcal{X}), \text{len}(\mathcal{Y})))$, such that*

$$\begin{aligned} \text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A}) &\leq (\psi q_H) \cdot \text{Adv}_{\mathsf{T}}^{\text{OW}}(\mathcal{B}) + \text{Adv}_{\mathsf{T}}^{\text{PS}}(\mathcal{D}) + \frac{(q_S + q_H)(q'_S + q_H + q_G)}{2^{2\lambda}} \\ &\quad + \frac{q_S(q'_S + q_H)}{2^\lambda} + \frac{\psi q_H^2}{2|\mathcal{Y}|} + \frac{(\psi q_H)^{\psi+1} |\mathcal{X}|}{(\psi + 1)! \cdot |\mathcal{Y}|^{\psi+1}}, \end{aligned}$$

where $\psi \geq \lceil \text{len}(\mathcal{X}) / \text{len}(\mathcal{Y}) \rceil$, and q'_S is a bound on the total number of queries to H in all the signing queries.

Algorithm 2: HaS-HF-SAS

Let $h_0 = \varepsilon, x_0 = \varepsilon$. The random oracles are $\mathbf{H}: \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$ and $\mathbf{G}: \{0, 1\}^{2\lambda} \rightarrow \mathcal{Y}$. The encoding function is $\text{enc}: \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}'$ and the corresponding decoding function is $\text{dec}: \mathcal{Y} \times \mathcal{X}' \rightarrow \mathcal{X}$ such that $\text{dec}(\text{enc}(x)) = x$. ϱ_i and ς_i are the partial description and the complementary information of the aggregate signature Σ_i , respectively.

KGen(1^λ):

- 1: $(\mathbf{F}, \mathbf{l}) \leftarrow \text{TrapGen}(1^\lambda)$
- 2: **return** $\text{pk} \leftarrow \mathbf{F}, \text{sk} \leftarrow (\mathbf{F}, \mathbf{l})$

AggSign(($\mathbf{F}_i, \mathbf{l}_i$), m_i, ϱ_{i-1}):

- 1: $(h_{i-1}, x_{i-1}) \leftarrow \varrho_{i-1}$
- 2: $(\alpha_{i-1}, \beta_{i-1}) \leftarrow \text{enc}(x_{i-1})$
- 3: **repeat**
- 4: $r_i \leftarrow \mathfrak{s} \{0, 1\}^\lambda$
- 5: $\eta_i \leftarrow \mathbf{H}(\mathbf{F}_i, m_i, r_i, x_{i-1})$
- 6: $h_i \leftarrow h_{i-1} \oplus \eta_i$
- 7: $g_i \leftarrow \mathbf{G}(h_i)$
- 8: $y_i \leftarrow g_i \oplus \alpha_{i-1}$
- 9: $x_i \leftarrow \mathbf{l}_i(y_i)$
- 10: **until** $x_i \neq \perp$
- 11: $\varrho_i \leftarrow (h_i, x_i)$
- 12: $\varsigma_i \leftarrow (r_i, \beta_{i-1})$
- 13: **return** ϱ_i, ς_i

AggVrfy($L_n, \bar{\Sigma}_n$):

- 1: $(\mathbf{F}_1, m_1), \dots, (\mathbf{F}_n, m_n) \leftarrow L_n$
- 2: $(\vec{r}_n, \vec{\beta}_{n-1}, h_n, x_n) \leftarrow \bar{\Sigma}_n$
- 3: **for** $i \leftarrow n, \dots, 2$ **do**
- 4: $y_i \leftarrow \mathbf{F}_i(x_i)$
- 5: $g_i \leftarrow \mathbf{G}(h_i)$
- 6: $\alpha_{i-1} \leftarrow g_i \oplus y_i$
- 7: $x_{i-1} \leftarrow \text{dec}(\alpha_{i-1}, \beta_{i-1})$
- 8: $\eta_i \leftarrow \mathbf{H}(\mathbf{F}_i, m_i, r_i, x_{i-1})$
- 9: $h_{i-1} \leftarrow h_i \oplus \eta_i$
- 10: **return** $h_1 = \mathbf{H}(\mathbf{F}_1, r_1, m_1, \varepsilon) \wedge \mathbf{F}_1(x_1) = \mathbf{G}(h_1)$

Combine($\varsigma_1, \dots, \varsigma_{n-1}, \Sigma_n$):

- 1: $(r_i, \beta_{i-1}) \leftarrow \varsigma_i$
- 2: $(r_n, \beta_{n-1}, h_n, x_n) \leftarrow \Sigma_n$
- 3: $\vec{r}_n \leftarrow (r_1, \dots, r_n)$
- 4: $\vec{\beta}_{n-1} \leftarrow (\beta_1, \dots, \beta_{n-1})$
- 5: **return** $\bar{\Sigma}_n \leftarrow (\vec{r}_n, \vec{\beta}_{n-1}, h_n, x_n)$

Proof (sketch). We sketch the high-level idea of the proof; full details can be found in Appendix A.1. The complete reduction is described in Algorithm 3. We prove the reduction by showing that the strong PS-HF-UF-CMA game can be simulated by the OW adversary \mathcal{B} . First, we modify the PS-HF-UF-CMA game such that in OAggSign the salt r is chosen uniformly at random in $\{0, 1\}^\lambda$ and the preimage is generated by $x \leftarrow \text{SampDom}(\mathbf{F}^*)$ instead of iterating until $\mathbf{l}^*(y) \neq \perp$. The PS adversary \mathcal{D} can simulate the two games by either playing PS_0 or PS_1 and the advantage in distinguishing the two games can therefore be estimated with $\text{Adv}_T^{\text{PS}}(\mathcal{D})$. Once the preimages are produced by $x \leftarrow \text{SampDom}(\mathbf{F}^*)$ without retry, we can adapt the techniques for trapdoor permutations of [13] to complete the reduction. In particular, we will use a labeled tree HTree whose nodes will be populated by some of the queries to the random oracle \mathbf{H} . The HTree is initialized with a root node with a single value $h_0 = \varepsilon$. Each subsequent node N_i is added following a query to the random oracle \mathbf{H} with input $Q_i = (\mathbf{F}_i, m_i, r_i, x_{i-1})$ and will store the following values:

- a reference to its parent node N_{i-1} ;
- the query Q_i to the random oracle \mathbf{H} ;
- the hash response to the query $\eta_i \leftarrow \mathbf{H}(Q_i)$;

- the hash state $h_i \leftarrow h_{i-1} \oplus \eta_i$, where h_{i-1} is the hash state stored in the parent node N_{i-1} ;
- an additional value $y_i \leftarrow \mathsf{G}(h_i) \oplus \alpha_{i-1}$ (where α_{i-1} is computed from $\mathsf{enc}(x_{i-1})$) that will be used to establish if future nodes can be added as children of N_i .

A node N_i can be added as child of a node N_{i-1} if it satisfies the relation $\mathsf{F}_{i-1}(x_{i-1}) = y_{i-1}$, where F_{i-1} and y_{i-1} are stored in N_{i-1} while x_{i-1} is stored in N_i . This relationship establishes that the query Q_i can be properly used by the signer with key F_i to aggregate its signature on message m_i with previous signature x_{i-1} , produced by key F_{i-1} and hash state h_{i-1} , which in turn are stored in N_{i-1} . Whenever a query $Q_i = (\mathsf{F}_i, m_i, r_i, x_{i-1})$, with $x_{i-1} \neq \varepsilon$, satisfies this relation with a node N_{i-1} we say that Q_i can be *tethered* to N_{i-1} . If $x_{i-1} = \varepsilon$, then Q_i can always be tethered to the root of the HTree.

Eventually, when the adversary \mathcal{A} outputs a valid aggregate signature $\bar{\Sigma}_n$ for the history $L_n = (\mathsf{pk}_1, m_1), \dots, (\mathsf{pk}_n, m_n)$, the simulator takes $i^* \in [n]$ such that $\mathsf{pk}_{i^*} = \mathsf{pk}^*$ and $(m_{i^*}, \varsigma_{i^*}) \notin \mathcal{Q}$ (the index i^* is guaranteed to exist when \mathcal{A} is winning). It then recovers x_{i^*} by iterating the procedure of Lines 3 to 9 in `AggVrfy` for $n - i^*$ steps. Then, the simulator checks if x_{i^*} is a preimage of a y_{i^*} in the HTree as a child of the node N_{i^*-1} storing $Q_{i^*-1} = (\mathsf{pk}_{i^*-1}, m_{i^*-1}, r_{i^*-1}, x_{i^*-2})$, which is itself a child of the node N_{i^*-2} , and so on until the node N_1 . If this is not the case, the simulator aborts by raising `badteth`. Otherwise, the value x_{i^*} produced by the forgery will satisfy $\mathsf{F}^*(x_{i^*}) = y_{i^*}$ for some y_{i^*} produced either on Line 19 or on Line 21 of `H` and stored in N_{i^*} . With probability $1/(\psi_{\mathsf{qH}})$, we have that y_{i^*} was produced on Line 21 of `H` and it is equal to y^* . Therefore $\mathsf{F}^*(x_{i^*}) = y^*$ and \mathcal{B} wins his OW game by returning x_{i^*} . \square

4 Instantiation and Evaluation

In this section, we will provide some concrete applications of the HaS-HF-SAS of Section 3 to MQ-based and code-based HaS signature schemes. In particular, we analyze the compression capabilities of the scheme when instantiated with UOV, MAYO, and Wave. More details on the trapdoor functions and the application of Theorem 1 to these schemes are given in Appendix B. Our scheme could also be extended to lattice-based schemes, such as the NIST PQC selected algorithm Falcon [33], and generally with PSF-based signatures [24]. In particular, applying our construction would allow to achieve history-free aggregation over existing schemes⁴. However, we already noted how different design choices become feasible due to the additional properties of trapdoor PSF. Accordingly, a direct application would lead to unnecessary loss of efficiency. An analysis of how the HaS-HF-SAS scheme can be modified with PSF can be found in Appendix C.

The main measure of the efficiency of an aggregate signature scheme is the *compression rate*, i.e., the reduction in the length of the aggregate signature $\bar{\Sigma}_N$ of N users compared to the trivial concatenation of N individual signatures

⁴ An attack on the, previously unique, history-free SAS of [35] was recently proposed in [12].

Algorithm 3: OW \implies PS-HF-UF-CMA

```

 $\mathcal{B}(\mathbf{F}^*, y^*)$ :
1:  $\mathcal{Q} \leftarrow \emptyset; c^* \leftarrow_{\$} [\mathbf{q}_H]; c \leftarrow 0$ 
2:  $(L_n, \bar{\Sigma}_n) \leftarrow \mathcal{A}^{\text{H.G.OAggSign}}(\mathbf{F}^*)$ 
3:  $(\mathbf{F}_1, m_1), \dots, (\mathbf{F}_n, m_n) \leftarrow L_n$ 
4: if  $\text{AggVrfy}(L_n, \Sigma_n) \wedge \exists i^* : (\mathbf{F}_{i^*} = \mathbf{F}^* \wedge m_{i^*} \notin \mathcal{Q})$  then
5:   Recover  $x_{i^*}$  as in  $\text{AggVrfy}$ 
6:    $\text{NList} \leftarrow \text{Lookup}(x_{i^*})$ 
7:   if  $\text{NList} = \perp$  then
8:     raise  $\text{bad}_{\text{teth}}$ 
9:   for  $N_{i^*} \in \text{NList}$  do
10:    Retrieve  $y_{i^*}$  from  $N_{i^*}$ 
11:    if  $y_{i^*} = y^*$  then
12:      return  $x_{i^*}$ 
13:   raise  $\text{bad}_{\text{inv}}$ 

 $\text{OAggSign}(m, \varrho = (h, x))$ :
1:  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{m\}$ 
2:  $(\alpha, \beta) \leftarrow \text{enc}(x)$ 
3:  $r \leftarrow_{\$} \{0, 1\}^\lambda$ 
4: if  $\text{HT}[\mathbf{F}^*, m, r, x] \neq \perp$  then
5:   raise  $\text{bad}_{\text{hcol}}$ 
6:  $\eta \leftarrow_{\$} \{0, 1\}^{2\lambda}$ 
7:  $\text{HT}[\mathbf{F}^*, m, r, x] \leftarrow \eta$ 
8:  $h' \leftarrow h \oplus \eta$ 
9: if  $\text{GT}[h'] \neq \perp$  then
10:  raise  $\text{bad}_{\text{gcol1}}$ 
11:  $x' \leftarrow_{\$} \mathcal{X}$ 
12:  $y' \leftarrow \mathbf{F}^*(x')$ 
13:  $\text{GT}[h'] \leftarrow y' \oplus \alpha$ 
14: return  $(r, \beta), (h', x')$ 

 $\mathbf{G}(h)$ :
1: if  $\text{GT}[h] = \perp$  then
2:   $g \leftarrow_{\$} \mathcal{Y}$ 
3:   $\text{GT}[h] \leftarrow g$ 
4: return  $\text{GT}[h]$ 

 $\mathbf{H}(\mathbf{F}, m, r, x)$ :
1:  $\mathcal{Q} \leftarrow (\mathbf{F}, m, r, x)$ 
2:  $c \leftarrow c + 1$ 
3: if  $\text{HT}[\mathcal{Q}] = \perp$  then
4:   $\eta \leftarrow_{\$} \{0, 1\}^{2\lambda}$ 
5:   $\text{HT}[\mathcal{Q}] \leftarrow \eta$ 
6:   $\text{NList} \leftarrow \text{Lookup}(x)$ 
7:  if  $\mathbf{F} = \mathbf{F}^* \wedge c = c^*$  then
8:     $i^* \leftarrow_{\$} [|\text{NList}|]$ 
9:    for  $i \in [|\text{NList}|]$  do
10:      $N_i \leftarrow \text{NList}[i]$ 
11:      $N'_i \leftarrow$  new node with parent  $N_i$ 
12:     Retrieve  $h_i$  from  $N_i$ 
13:      $h'_i \leftarrow h_i \oplus \eta$ 
14:      $(\alpha, \beta) \leftarrow \text{enc}(x)$ 
15:     if  $\text{GT}[h'_i] \neq \perp$  then
16:       raise  $\text{bad}_{\text{gcol2}}$ 
17:     if  $\mathbf{F} \neq \mathbf{F}^* \vee c \neq c^* \vee i \neq i^*$  then
18:        $g'_i \leftarrow \mathbf{G}(h'_i)$ 
19:        $y'_i \leftarrow g'_i \oplus \alpha$ 
20:     else
21:        $y'_i \leftarrow y_i^*$ 
22:        $\text{GT}[h'_i] \leftarrow y'_i \oplus \alpha$ 
23:     Populate node  $N'_i$  with  $\mathcal{Q}, \eta, h'_i, y'_i$ 
24: return  $\text{HT}[\mathcal{Q}]$ 

 $\text{Lookup}(x)$ :
1: if  $x = \varepsilon$  then
2:  return Root of  $\text{HTree}$ 
3:  $\text{NList} \leftarrow \{N \in \text{HTree} : (\mathbf{F}, y) \in N \wedge \mathbf{F}(x) = y\}$ 
4: if  $|\text{NList}| > \psi$  then
5:  raise  $\text{bad}_{\text{tcol}}$ 
6: else if  $|\text{NList}| = 0$  then
7:  return  $\perp$ 
8: else
9:  return  $\text{NList}$ 

```

σ . The compression rate of N signatures is defined as $\tau(N) = 1 - \frac{|\bar{\Sigma}_N|}{N \cdot |\sigma|}$. An HaS-HF-SAS signature of N users is the output of the **Combine** algorithm on $\varsigma_1, \dots, \varsigma_{N-1}, \Sigma_N$ and is given by $\bar{\Sigma}_N = (\vec{r}_N, \beta_{N-1}, h_N, x)$. An individual signature of a generic HaS scheme as described in Section 2.3 is given by $\sigma = (r, x)$. In the following, we assume that the aggregation scheme is applied to the signature scheme without further possible optimization, so that we have the same size for salts $|r| = \lambda$ and preimages $|x| = \text{len}(\mathcal{X})$, where $\text{len}(X)$ denotes the bit size of an element in X .

The compression rate is thus given by

$$\begin{aligned} \tau(N) &= 1 - \frac{N \cdot \lambda + (N-1) \cdot \text{len}(\mathcal{X}') + 2\lambda + \text{len}(\mathcal{X})}{N \cdot (\lambda + \text{len}(\mathcal{X}))} \\ &= 1 - \frac{N \cdot (\lambda + \text{len}(\mathcal{X}) - \text{len}(\mathcal{Y})) + 2\lambda + \text{len}(\mathcal{Y})}{N \cdot (\lambda + \text{len}(\mathcal{X}))} \quad (1) \\ &= \frac{\text{len}(\mathcal{Y})}{\lambda + \text{len}(\mathcal{X})} - \frac{2\lambda + \text{len}(\mathcal{Y})}{N \cdot (\lambda + \text{len}(\mathcal{X}))}. \end{aligned}$$

Notice that the aggregate signature is smaller than the trivial concatenation whenever $N > \frac{2\lambda}{\text{len}(\mathcal{Y})} + 1$, which for typical parameters is as soon as $N > 2$.

In [20,15], the size of an aggregate signature of N users is $N \cdot (\text{len}(\mathcal{X}) - \text{len}(\mathcal{Y})) + \text{len}(\mathcal{Y})$. Compared to our scheme, we have a small overhead of $(N+2)\lambda$ bytes due to the aggregated hash state and the random salts. However, as we will see in the next section, this increase in signature size is necessary to guarantee the security of the scheme.

UOV (Appendices B.2 and B.3) We consider the parameters proposed in [10] with respect to NIST security levels I, III, and V. For UOV, the domain \mathcal{X} is given by \mathbb{F}_q^n with elements of length $\text{len}(\mathcal{X}) = n \lceil \log_2 q \rceil$. The codomain \mathcal{Y} is \mathbb{F}_q^m with elements of length $\text{len}(\mathcal{Y}) = m \lceil \log_2 q \rceil$. Regardless of the security level, [10] use 128-bit salts. In our comparison, we consider salts of length $\lambda = 128, 192$ and 256 bits, respectively.

The size of a sequential aggregate signature instantiated with UOV is given by $|\bar{\Sigma}_N| = N \cdot (\lambda + (n-m) \lceil \log_2 q \rceil) + 2\lambda + m \lceil \log_2 q \rceil$ and the size of a single signature is given by $|\sigma| = n \lceil \log_2 q \rceil + \lambda$.

Concrete numbers for different security parameters and the number of signers are given in Table 1.

MAYO (Appendix B.4) We consider the parameters proposed in [9] with respect to NIST security levels I, III, and V. For MAYO, the domain \mathcal{X} is given by \mathbb{F}_q^{kn} with elements of length $\text{len}(\mathcal{X}) = kn \lceil \log_2 q \rceil$. The codomain \mathcal{Y} is \mathbb{F}_q^m with elements of length $\text{len}(\mathcal{Y}) = m \lceil \log_2 q \rceil$. In [9], the salt length $|r|$ is slightly longer than the security parameter for consistency with the security proof. In our comparison, we maintain this choice, adjusting the compression rate computation of Equation (1).

The size of a sequential aggregate signature instantiated with MAYO is given by $|\bar{\Sigma}_N| = N \cdot (|r| + (kn - m)\lceil \log_2 q \rceil) + 2\lambda + m\lceil \log_2 q \rceil$ and the size of a single signature is given by $|\sigma| = kn\lceil \log_2 q \rceil + |r|$.

Concrete numbers for different security parameters and the number of signers are given in Table 2.

Wave (Appendix B.5) We consider the parameters proposed in [17] with respect to NIST security level I. Notice that this is not the same scheme later submitted to NIST’s call for additional digital signatures. The submitted scheme incorporates an optimization derived from the Wavelet variant [3], which cannot be used during aggregation and for which only an asymptotically trivial compression rate can be obtained. For Wave, the domain \mathcal{X} is given by $S_{w,n}$, the subset of \mathbb{F}_q^n with vectors of hamming weight w , with elements of length $\text{len}(\mathcal{X}) = \lceil n \log_2 q \rceil$. The codomain \mathcal{Y} is \mathbb{F}_q^{n-k} with elements of length $\text{len}(\mathcal{Y}) = \lceil (n - k) \log_2 q \rceil$.

The size of a sequential aggregate signature instantiated with Wave is given by $|\bar{\Sigma}_N| = N \cdot (\lambda + \lceil k \log_2 q \rceil) + 2\lambda + \lceil (n - k) \log_2 q \rceil$ and the size of a single signature is given by $|\sigma| = \lceil n \log_2 q \rceil + \lambda$.

Concrete numbers for different security parameters and the number of signers are given in Table 3.

4.1 Single-signature optimizations

For proper evaluation of the efficiency of SAS, it is necessary to consider any optimizations of the single signature that cannot be used in aggregation. Ignoring possible optimizations can lead to an unfair comparison and an incorrect calculation of the compression rate.

In [12], Boudgoust and Takahashi observe that in the context of lattice-based signatures built on PSFs (e.g., Falcon) it is possible to reduce the size of signatures considerably by slightly modifying the verification process. However, the same variant is not applicable in the context of aggregate signatures. Similarly, as noted in the previous section, the optimization introduced by the Wavelet variant cannot be applied in the aggregation phase, causing a significant loss of efficiency.

More generally, any compression method applied to a single hash-and-sign signature can be employed in our construction, provided sufficient information exists in the signature to recover the message hash. To elaborate further, consider the generic hash-and-sign scheme outlined in Algorithm 1. Suppose that the Sign algorithm returns the pair $(r, C(x))$, where C is a compression algorithm on the preimage x . If, during the verification process, it is possible to recover $H(r, m)$ from the public key F and $C(x)$, then the same optimization can be effectively employed within the HaS-HF-SAS scheme. In fact, the aggregation process of Algorithm 2 can be tweaked to aggregate part of the compressed preimage $C(x)$ without the need to modify the verification step further. However, the optimized versions of Falcon and Wave(let) do not conform to this description, as their verification process does not enable the recovery of the message hash from the

Table 1. Aggregate signature sizes and compression rates of HaS-HF-SAS scheme on UOV parameters from [10].

Parameter	ov-Ip	ov-Is	ov-III	ov-V
NIST SL	I	I	III	V
(n, m, q)	(112,44,256)	(160,64,16)	(184,72,256)	(244,96,256)
$N \cdot \sigma $ (bytes)	$128 \cdot N$	$96 \cdot N$	$208 \cdot N$	$276 \cdot N$
$ \bar{\Sigma}_N $ (bytes)	$84 \cdot N + 76$	$64 \cdot N + 64$	$136 \cdot N + 120$	$180 \cdot N + 160$
$\tau(5)$	0.23	0.20	0.23	0.23
$\tau(20)$	0.31	0.30	0.32	0.32
$\tau(100)$	0.34	0.33	0.34	0.34
Asym. $\tau(N)$	0.34	0.33	0.35	0.35

Table 2. Aggregate signature sizes and compression rates of HaS-HF-SAS scheme on MAYO parameters from [9].

Parameter	MAYO ₁	MAYO ₂	MAYO ₃	MAYO ₅
NIST SL	I	I	III	V
(n, m, o, k, q)	(66,64,8,9,16)	(78,64,18,4,16)	(99,96,10,11,16)	(133,128,12,12,16)
$ r $ (bytes)	24	24	32	40
$N \cdot \sigma $ (bytes)	$321 \cdot N$	$180 \cdot N$	$577 \cdot N$	$838 \cdot N$
$ \bar{\Sigma}_N $ (bytes)	$289 \cdot N + 64$	$148 \cdot N + 64$	$529 \cdot N + 96$	$774 \cdot N + 128$
$\tau(5)$	0.06	0.11	0.05	0.05
$\tau(20)$	0.09	0.16	0.07	0.07
$\tau(100)$	0.10	0.17	0.08	0.07
Asym. $\tau(N)$	0.10	0.18	0.08	0.08

Table 3. Aggregate signature sizes and compression rates of HaS-HF-SAS scheme on Wave parameters from [17].

Parameter	128g
NIST SL	I
(n, k, w, q)	(8492,5605,7980,3)
$N \cdot \sigma $ (bytes)	$1699 \cdot N$
$ \bar{\Sigma}_N $ (bytes)	$1127 \cdot N + 604$
$\tau(5)$	0.27
$\tau(20)$	0.32
$\tau(100)$	0.33
Asym. $\tau(N)$	0.34

Game 3: FH-UF-CMA_{S'}

1: $(pk^*, sk^*) \leftarrow \text{KGen}(1^\lambda)$ 2: $\mathcal{Q} \leftarrow \emptyset$ 3: $(L_n, \Sigma_n) \leftarrow \mathcal{A}^{\text{O,OAgsign}}(pk^*)$ 4: $(pk_1, m_1), \dots, (pk_n, m_n) \leftarrow L_n$ 5: if $\nexists i^* : (pk_{i^*} = pk^* \wedge (m_{i^*}, L_{i^*}) \notin \mathcal{Q})$ then 6: return \perp 7: return $\text{AggVrfy}(L_n, \Sigma_n)$	OAgsign $(m_i, L_{i-1}, \Sigma_{i-1})$: 1: if $\text{AggVrfy}(L_{i-1}, \Sigma_{i-1}) = 0$ then 2: return \perp 3: $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(L_{i-1}, m_i)\}$ 4: $\Sigma_i \leftarrow \text{AggSign}(sk^*, m_i, L_{i-1}, \Sigma_{i-1})$ 5: return Σ_i
---	--

signature and the public key. Instead, the verification is based on a custom assertion involving the knowledge of F , $C(x)$, and $H(r, m)$.

5 Security of Existing Multivariate SAS Schemes

In this section, we show a universal forgery for the sequential aggregate signature schemes of [20,15] when instantiated with the UOV signature scheme. Both schemes are based on the variant with encoding of [31] and require an alternative definition of SAS with the notion of *full-history*: at each aggregation step, the signer needs the so-far aggregated signature and the complete list of messages and public keys of previous signers. Moreover, knowledge of the full description of the aggregate signature is required, as the signer needs to check its validity before adding its own.

Definition 9 (FH-SAS). *A Full-History Sequential Aggregate Signature is a tuple of three algorithms $(\text{KGen}, \text{AggSign}, \text{AggVrfy})$:*

- $\text{KGen}(1^\lambda)$ as described in Definition 6.
- $\text{AggSign}(sk_i, m_i, L_{i-1}, \Sigma_{i-1})$: takes as input the secret key sk_i and the message m_i of the i th user, a list $L_{i-1} = (pk_1, m_1), \dots, (pk_{i-1}, m_{i-1})$ of public keys, message pairs, and the previous aggregate signature Σ_{i-1} . If $\text{AggVrfy}(L_{i-1}, \Sigma_{i-1}) = 1$, it returns an updated aggregate signature Σ_i .
- $\text{AggVrfy}(L_n, \Sigma_n)$: takes as input the full history $L_n = (pk_1, m_1), \dots, (pk_n, m_n)$ of public key, message pairs, and an aggregate signature Σ_n . Returns 1 if Σ_n is a valid aggregate signature and 0 otherwise.

Below we show the definition of full-history unforgeability under adaptive chosen message (FH-UF-CMA). Compared to the notion of PS-HF-UF-CMA (Definition 8), the signing oracle OAgsign requires sending the list L_{i-1} of public keys and messages along with the aggregate signature Σ_{i-1} and returns the updated signature if and only if Σ_{i-1} is valid.

Definition 10 (FH-UF-CMA Security). *Let \mathcal{O} be a random oracle, let $S' = (\text{KGen}, \text{AggSign}, \text{AggVrfy})$ be a FH-SAS scheme, let \mathcal{A} be an adversary. We define the advantage of \mathcal{A} playing the FH-UF-CMA game (Game 3) against S' as follows:*

$$\text{Adv}_{S'}^{\text{FH-UF-CMA}}(\mathcal{A}) = \Pr[\text{FH-UF-CMA}_{S'}(\mathcal{A}) = 1].$$

Let $\Sigma_0 = (\emptyset, \varepsilon)$.	
KGen (1^λ): 1: $(F, l) \leftarrow \text{TrapGen}(1^\lambda)$ 2: return $\text{pk} \leftarrow F, \text{sk} \leftarrow (F, l)$	AggSign ($((F_i, l_i), m_i, L_{i-1}, \Sigma_{i-1})$): 1: $(\beta_{i-2}, x_{i-1}) \leftarrow \Sigma_{i-1}$ 2: if $\text{AggVrfy}(L_{i-1}, \Sigma_{i-1}) = 0$ then 3: return \perp 4: $L_i \leftarrow L_{i-1} \cup \{(F_i, m_i)\}$ 5: $(\alpha_{i-1}, \beta_{i-1}) \leftarrow \text{enc}(x_i)$ 6: $h_i \leftarrow H(L_i)$ 7: $\vec{x}_i \leftarrow l_i(\alpha_{i-1} \oplus h_i)$ 8: $\beta_{i-1} \leftarrow \beta_{i-2} \cup \{\beta_{i-1}\}$ 9: return $\Sigma_i \leftarrow (\beta_{i-1}, x_i)$
AggVrfy ((L_n, Σ_n)): 1: $(F_1, m_1), \dots, (F_n, m_n) \leftarrow L_n$ 2: $(\beta_{n-1}, x_n) \leftarrow \Sigma_n$ 3: for $i \leftarrow n, \dots, 2$ do 4: $L_i \leftarrow (F_1, m_1), \dots, (F_i, m_i)$ 5: $h_i \leftarrow H(L_i)$ 6: $\alpha_{i-1} \leftarrow F_i(x_i) \oplus h_i$ 7: $x_{i-1} \leftarrow \text{dec}(\alpha_{i-1}, \beta_{i-1})$ 8: return $F_1(x_1) = H(F_1, m_1)$	

5.1 Multivariate FH-SAS

The FH-SAS schemes of [20,15] are instantiated with HFEv- and UOV, respectively, but no explicit use is made of unique features of these trapdoor functions. The description of Algorithm 4 refers to a generic multivariate trapdoor function T (as in Section 2.1) and is based on the construction of [20], which is slightly more general.

In Algorithm 4, the random oracle is $\mathsf{H}: \{0, 1\}^* \rightarrow \mathcal{Y}$. The encoding function is $\text{enc}: \mathcal{X} \rightarrow \mathcal{Y} \times \mathcal{X}'$ that splits an element x_i as $\text{enc}(x_i) = (\alpha_i, \beta_i)$ and the corresponding decoding function is $\text{dec}: \mathcal{Y} \times \mathcal{X}' \rightarrow \mathcal{X}$ such that $\text{dec}(\text{enc}(x)) = x$. To simplify the description we will also use the notation $\alpha(x_i) = \alpha_i$ and $\beta(x_i) = \beta_i$, where $\alpha(\cdot), \beta(\cdot)$ are implicitly defined by enc .

Both [20] and [15] provide a similar claim on the formal security of their sequential aggregate signature scheme. In the following, we are considering a generic multivariate trapdoor function since their choice does not influence the security claim.

Theorem 2 ([20]). *Let T be a multivariate trapdoor function. Let \mathcal{A} be a FH-UF-CMA adversary against the FH-SAS scheme on T in the random oracle model, which makes q_S signing queries and q_H queries to the random oracle. Then, there exist a OW adversary \mathcal{B} against T such that*

$$\text{Adv}_{\text{FH-SAS}_{\mathsf{T}}}^{\text{FH-UF-CMA}}(\mathcal{A}) \leq (q_S q_H + 1) \cdot \text{Adv}_{\mathsf{T}}^{\text{OW}}(\mathcal{B})$$

and the running time of \mathcal{B} is about that of \mathcal{A} .

In [15], the authors omit the proof for their security claim, while in [20] the authors provide a sketch of the proof in which they state that almost all the steps of the security proof follow [31] with only some slight modifications taking into account the use of the encoding function.

In the next section, we show an explicit universal forgery on FH-SAS when instantiated with the UOV signature scheme. Then, in Section 5.3, we provide some insight into why the security proof of [31] cannot be applied to multivariate schemes and, more generally, to signature schemes based on trapdoor functions that are not permutations.

5.2 Description of the Forgery

We recall that in UOV, the trapdoor function is a multivariate quadratic map $\mathcal{P}: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on a secret linear subspace $O \subset \mathbb{F}_q^n$ of dimension m . A more in-depth description can be found in Appendix B.1.

In the following we are assuming that the encoding function $\text{enc}(\mathbf{x})$ can be expressed via an appropriate affine map and, accordingly, $\alpha(\mathbf{x}) = R(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$, where $\mathbf{A} \in \mathbb{F}_q^{m \times n}$, $\mathbf{b} \in \mathbb{F}_q^m$. In [20,15], $\text{enc}(\mathbf{x})$ is always the projection in the first m and the last $n - m$ components⁵ of \mathbf{x} . This is a slight generalization that captures the intuition that there must be a corresponding efficient decoding function.

Lemma 1. *The multivariate FH-SAS scheme of Section 5.1, instantiated with UOV, is not FH-UF-CMA.*

Proof. Let $\text{pk}_i = \mathcal{P}_i$ be the target public key and assume that the forger \mathcal{F} knows a valid aggregate signature $\Sigma_i = (\beta_1, \dots, \beta_{i-1}, \mathbf{x}_i)$ for a honest history $L_i = (\text{pk}_1, m_1), \dots, (\text{pk}_i, m_i)$. This is a typical attack environment, much weaker than the notion of FH-UF-CMA that we introduced in Definition 10. Then, \mathcal{F} select a message m_i^* on which it will produce a forged signature for the target user.

The forger \mathcal{F} computes a forged signature by replacing the $(i - 1)$ th honest signer, as follows:

1. First, \mathcal{F} appropriately generates a UOV key pair $(\text{pk}_{\mathcal{F}}, \text{sk}_{\mathcal{F}}) = (\mathcal{P}_{\mathcal{F}}, O_{\mathcal{F}})$ by randomly choosing an m -dimensional linear subspace $O_{\mathcal{F}} \subset \ker \mathbf{A}$ and use the same procedure of $\text{TrapGen}_{\text{uov}}$ (Algorithm 5) to sample $\mathcal{P}_{\mathcal{F}}$ that vanishes on $O_{\mathcal{F}}$.
2. Then, \mathcal{F} arbitrarily chooses a message $m_{\mathcal{F}}$, computes a corresponding forged history $L^* = L_{i-2} \cup \{(\text{pk}_{\mathcal{F}}, m_{\mathcal{F}}), (\text{pk}_i, m_i)\}$ and computes $\alpha^* \leftarrow \mathcal{P}_i(\mathbf{x}_i) \oplus \mathbf{H}(L^*)$.
3. Finally, \mathcal{F} finds a preimage $\mathbf{x}_{\mathcal{F}}$ under $\mathcal{P}_{\mathcal{F}}$ for $L_{\mathcal{F}} = L_{i-2} \cup \{(\text{pk}_{\mathcal{F}}, m_{\mathcal{F}})\}$ such that

$$\mathcal{P}_{\mathcal{F}}(\mathbf{x}_{\mathcal{F}}) = \alpha_{i-2} \oplus \mathbf{H}(L_{\mathcal{F}}) \quad \text{and} \quad \alpha(\mathbf{x}_{\mathcal{F}}) = \alpha^*. \quad (2)$$

Then $\Sigma^* = (\beta_1, \dots, \beta_{i-2}, \beta(\mathbf{x}_{\mathcal{F}}), \mathbf{x}_i)$ is a valid aggregate signature for the forged history L^* .

⁵ In this case we would have that $\alpha(\mathbf{x}) = \mathbf{A}\mathbf{x}$ with $\mathbf{A} = [\mathbf{I}_m \mid \mathbf{0}_{m, n-m}]$.

Finding a preimage $\mathbf{x}_{\mathcal{F}}$ that satisfies Equation (2) is equivalent to finding partially fixed preimage for $\mathcal{P}_{\mathcal{F}}$ under the map R . In particular, the forger can use the appropriately generated secret key $O_{\mathcal{F}}$ to restrict the preimage search to an appropriate affine subspace and guarantee the condition $R(\mathbf{x}_{\mathcal{F}}) = \alpha^*$. The forger searches for a preimage of $\mathbf{t} = \alpha_{i-2} \oplus \mathbf{H}(L_{\mathcal{F}})$ by using a procedure similar to the Sign procedure described in Algorithm 5: on Line 1, instead of randomly sampling the vectors \mathbf{v} from \mathbb{F}_q^n , samples \mathbf{v} from $\ker R'$, with $R'(\mathbf{x}) = \mathbf{A}\mathbf{x} + (\mathbf{b} - \alpha^*)$. Then when a preimage $\mathbf{x}_{\mathcal{F}} \in \mathbb{F}_q^n$ of \mathbf{t} is found, the forger would have $\mathbf{x}_{\mathcal{F}} \in \ker R'$, since $\mathbf{x}_{\mathcal{F}} = \mathbf{v} + \mathbf{o}$ with $\mathbf{v} \in \ker R'$ and $\mathbf{o} \in \ker \mathbf{A}$. Therefore $\alpha(\mathbf{x}_{\mathcal{F}}) = R(\mathbf{x}_{\mathcal{F}}) = \alpha^*$.

We then show that Σ^* passes the verification correctly for the forged history L^* :

1. The verifier applies the first iteration of `AggVrfy` (Algorithm 4) to recover the previous signature $\mathbf{x}_{\mathcal{F}}$ from \mathbf{x}_i as follows:

$$\alpha(\mathbf{x}_{\mathcal{F}}) \leftarrow \mathcal{P}_i(\mathbf{x}_i) \oplus \mathbf{H}(L^*) = \alpha^*, \quad \mathbf{x}_{\mathcal{F}} \leftarrow \text{dec}(\alpha(\mathbf{x}_{\mathcal{F}}), \beta(\mathbf{x}_{\mathcal{F}}))$$

2. Since $\mathbf{x}_{\mathcal{F}}$ is a preimage of $\alpha_{i-2} \oplus \mathbf{H}(L_{\mathcal{F}})$, the verifier correctly obtains \mathbf{x}_{i-2} proceeding in the iterations of `AggVrfy`:

$$\alpha_{i-2} \leftarrow \mathcal{P}_{\mathcal{F}}(\mathbf{x}_{\mathcal{F}}) \oplus \mathbf{H}(L_{\mathcal{F}}), \quad \mathbf{x}_{i-2} \leftarrow \text{dec}(\alpha_{i-2}, \beta_{i-2}).$$

3. The $(i-2)$ th signer was not tampered and, hence, the intermediate signature $\Sigma_{i-2} = (\beta_1, \dots, \beta_{i-3}, \mathbf{x}_{i-2})$ can be correctly verified with `AggVrfy` on honest history L_{i-2} .

Therefore, the verifier determines the forged signature Σ^* as valid. \square

5.3 Discussion

The previous forging procedure can be directly applied to constructions derived from [31] and instantiated via UOV, such as [20,15]. In particular, we have shown how the existential unforgeability claims of [20,15] are incorrect when the schemes are instantiated with UOV. The attack essentially involves finding a partially fixed preimage, following an adversarial key generation based on the public parameters of the aggregate signature scheme, specifically the encoding function. Although this attack may have applicability beyond UOV, it is not a universal forgery for generic trapdoor functions. However, this result aligns with the analysis of critical issues encountered when attempting to extend the security proof of [31] to generic trapdoor functions, as outlined in the following.

Programming the Random Oracle In [31], the random oracle can be simulated to determine preimages for any permutation $\pi : \mathcal{X} \rightarrow \mathcal{X}$. This is typically achieved by sampling a uniformly random $x \leftarrow_s \mathcal{X}$ and returning $\pi(x)$, which is uniformly distributed in $\mathcal{Y} = \mathcal{X}$. However, in the case of a generic trapdoor function, we cannot assume that the image of \mathbf{F} is uniform. Consequently, to provide an accurate simulation, we must sample and return a uniformly random value in \mathcal{Y} .

Uniqueness of the Aggregate Signature If we relax the previous condition and assume a uniformity property of the trapdoor functions⁶ we may attempt to replicate the process described in [31] to answer the signing oracle. Indeed, on input $Q = (m^*, L_{n-1}, \Sigma_{n-1})$ the simulator can use the knowledge of appropriate preimages for F_i to craft a valid aggregate signature on $L_{n-1} \cup \{(F^*, m^*)\}$. However, we argue that this property alone is not sufficient for a correct simulation, which is instead based on the following fact of TDP-based constructions: for a fixed input $L_n = (F_1, m_1), \dots, (F_n, m_n)$ there exists a unique aggregate signature on L_n . Otherwise, if the aggregate signature is not unique, the simulator would be unable to provide a valid response to the aggregate signature query. In fact, on input Q , the simulator would take the preimage x^* for F^* on $\alpha(x_{n-1}) \oplus H(L_n)$ associated to the random oracle query on input $L_n = L_{n-1} \cup \{(F^*, m^*)\}$. But, without the knowledge that x_{n-1} is equal to the preimage computed by the adversary on input L_{n-1} , the aggregate signature produced by the simulator may not be properly verified, resulting in an invalid response.

Reduction to OW Eventually, the adversary will produce a valid non-trivial aggregate signature Σ_n on input $L_n = (F_1, m_1), \dots, (F_n, m_n)$, where we assume, for simplicity, that $F_n = F^*$ is the target public key. Since the aggregate signature is correct, it follows that $F^*(x_n) = \alpha(x_{n-1}) \oplus H(L_n) = y$. In the context of TDPs, [31] shows that y is equal to the target y^* of the OW game, with probability $(q_H + q_S + 1)^{-1}$.

When we consider generic TDFs, the previously mentioned approach is not valid, and it is necessary to modify the simulation of H by returning a fresh random value for each query. Moreover, we claim that in this setting it is not possible to correctly simulate the response to the oracles in order to obtain a preimage of y^* . In fact, to obtain a valid preimage, the simulator would require $F^*(x_n) = \alpha(x_{n-1}) \oplus H(L_n) = y^*$ and therefore $H(L_n) = y^* \oplus \alpha(x_{n-1})$. It should then have been able to simulate the random oracle to return $y^* \oplus \alpha(x_{n-1})$ when given the input L_n . However, it is not possible to provide this answer, as x_{n-1} is not part of the query input and is not uniquely determined.

Fixing the Forging Vulnerability The main vulnerability of FH-SAS concerns the overall malleability of the aggregate signature. In the original scheme for TPDs, once the input $L_n = (F_1, m_1), \dots, (F_n, m_n)$ is fixed, it was observed that there is a unique aggregate signature on messages m_1, \dots, m_n under public keys F_1, \dots, F_n . Instead, in the extended version, uniqueness is lost due to the probabilistic nature of the inversion process. Consequently, it is always possible to construct two aggregate signatures on the same input, $\Sigma = (\beta_1, \dots, \beta_{i-1}, x_n)$ and $\Sigma' = (\beta_1, \dots, \beta_{i-1}, x'_i)$, which differ only in the aggregation of the last signature. Furthermore, as shown in the forgery presented in Section 5.2, it is possible to have two aggregate signatures on the same input $\Sigma = (\beta_1, \dots, \beta_{i-1}, x_i)$ and $\Sigma' = (\beta_1, \dots, \beta'_{i-1}, x_i)$ which differ only in the intermediate partial encodings. While the loss of uniqueness is unavoidable, it is possible to modify the

⁶ For instance, this is the case for Trapdoor Preimage Sampleable Function [24].

scheme to prevent this additional form of malleability by making partial β encodings part of the random oracle input. We modify the aggregation step of $\text{AggSign}((F_i, l_i), m_i, L_{i-1}, \Sigma_{i-1})$ (Algorithm 4): let $\Sigma_{i-1} = (\beta_1, \dots, \beta_{i-2}, x_{i-1})$ and compute

$$(\alpha_{i-1}, \beta_{i-1}) \leftarrow \text{enc}(x_{i-1}), \quad x_i \leftarrow l_i(\alpha_{i-1} \oplus H(L_i, \vec{\beta}_{i-1})),$$

where $L_i = L_{i-1} \cup \{(F_i, m_i)\}$ and $\vec{\beta}_{i-1} = (\beta_1, \dots, \beta_{i-1})$.

Observe that now, once a new signature has been aggregated, it is no longer possible to modify the previous partial encodings while maintaining the validity of the aggregated signature. That is, if $\Sigma = (\beta_1, \dots, \beta_{i-1}, x_i)$ and $\Sigma' = (\beta_1, \dots, \beta'_{i-1}, x'_i)$ are valid aggregate signatures on the same input with $\beta_{i-1} \neq \beta'_{i-1}$, then $x_i \neq x'_i$. As a result, the forging procedure described in Section 5.2 is no longer applicable, as the adversary now needs to guess the partial encoding $\beta(x_{\mathcal{F}})$ of its own signature. However, in doing so $\beta(x_{\mathcal{F}})$ becomes fixed and α^* is not under the adversary's direct control. Once α^* is computed, the entire signature $x_{\mathcal{F}}$ is fixed, and with high probability, it is not a valid signature.

This minor modification addresses the vulnerability exploited by our attack. However, from a provable security perspective, this construction presents similar problems to the original attempt to generalize [31]. As a result, we are unable to provide a formal proof of security.

6 Conclusions

We proposed a partial-signature history-free sequential aggregate signature within the probabilistic hash-and-sign with retry paradigm, generalizing previous results to generic trapdoor functions. We proved the security of our scheme in the random oracle model, assuming only the one-wayness of the underlying TDF and an additional notion of indistinguishability on preimages. This additional property has been demonstrated for numerous post-quantum TDFs to achieve the security of the related signature schemes. This allowed us to easily instantiate our construction in Section 4 with the UOV, MAYO, and Wave schemes, for which we obtained a compression rate between 5% and 34%. We also pointed out in Section 5 how existing aggregation schemes for multivariate TDFs lack formal security and are insecure for some choices of the underlying function. Therefore, within our knowledge, ours is the first scheme that allows the aggregation of multivariate and code-based HaS signature schemes.

Acknowledgments. The authors would like to thank the anonymous reviewers of CT-RSA 2024 for their valuable feedback and suggestions. This publication was created with the co-financing of the European Union FSE-REACT-EU, PON Research and Innovation 2014-2020 DM1062/2021. The first author is a member of the INdAM Research Group GNSAGA. The second author is a member of CryptO, the group of Cryptography and Number Theory of Politecnico di Torino. The first author acknowledges support from Ripple's University Blockchain Research Initiative.

References

1. Ahn, J.H., Green, M., Hohenberger, S.: Synchronized aggregate signatures: new definitions, constructions and applications. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM CCS 2010. pp. 473–484. ACM Press (Oct 2010). <https://doi.org/10.1145/1866307.1866360>
2. Albrecht, M.R., Cini, V., Lai, R.W.F., Malavolta, G., Thyagarajan, S.A.K.: Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 102–132. Springer, Heidelberg (Aug 2022). https://doi.org/10.1007/978-3-031-15979-4_4
3. Banegas, G., Debris-Alazard, T., Nedeljković, M., Smith, B.: Wavelet: Code-based postquantum signatures with fast verification on microcontrollers. Cryptology ePrint Archive, Report 2021/1432 (2021), <https://eprint.iacr.org/2021/1432>
4. Bellare, M., Namprempre, C., Neven, G.: Unrestricted aggregate signatures. In: Arge, L., Cachin, C., Jurdzinski, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 411–422. Springer, Heidelberg (Jul 2007). https://doi.org/10.1007/978-3-540-73420-8_37
5. Bellare, M., Neven, G.: Multi-signatures in the plain public-key model and a general forking lemma. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006. pp. 390–399. ACM Press (Oct / Nov 2006). <https://doi.org/10.1145/1180405.1180453>
6. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.) ACM CCS 93. pp. 62–73. ACM Press (Nov 1993). <https://doi.org/10.1145/168588.168596>
7. Beullens, W.: Improved cryptanalysis of UOV and Rainbow. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 348–373. Springer, Heidelberg (Oct 2021). https://doi.org/10.1007/978-3-030-77870-5_13
8. Beullens, W.: MAYO: Practical post-quantum signatures from oil-and-vinegar maps. In: AlTawy, R., Hülsing, A. (eds.) SAC 2021. LNCS, vol. 13203, pp. 355–376. Springer, Heidelberg (Sep / Oct 2022). https://doi.org/10.1007/978-3-030-99277-4_17
9. Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: MAYO. Tech. rep., National Institute of Standards and Technology (2023), available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>
10. Beullens, W., Chen, M., Ding, J., Gong, B., Kannwischer, M.J., Patarin, J., Peng, B., Schmidt, D., Shih, C., Tao, C., Yang, B.: UOV — Unbalanced Oil and Vinegar. Tech. rep., National Institute of Standards and Technology (2023), available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>
11. Boneh, D., Gentry, C., Lynn, B., Shacham, H.: Aggregate and verifiably encrypted signatures from bilinear maps. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 416–432. Springer, Heidelberg (May 2003). https://doi.org/10.1007/3-540-39200-9_26
12. Boudgoust, K., Takahashi, A.: Sequential half-aggregation of lattice-based signatures. Cryptology ePrint Archive, Report 2023/159 (2023), <https://eprint.iacr.org/2023/159>
13. Brogle, K., Goldberg, S., Reyzin, L.: Sequential aggregate signatures with lazy verification from trapdoor permutations - (extended abstract). In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 644–662. Springer, Heidelberg (Dec 2012). https://doi.org/10.1007/978-3-642-34961-4_39

14. Chailloux, A., Debris-Alazard, T.: Tight and optimal reductions for signatures based on average trapdoor preimage sampleable functions and applications to code-based signatures. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020, Part II. LNCS, vol. 12111, pp. 453–479. Springer, Heidelberg (May 2020). https://doi.org/10.1007/978-3-030-45388-6_16
15. Chen, J., Ling, J., Ning, J., Peng, Z., Tan, Y.: MQ aggregate signature schemes with exact security based on UOV signature. In: Liu, Z., Yung, M. (eds.) Information Security and Cryptology - 15th International Conference, Inscrypt 2019, Nanjing, China, December 6–8, 2019, Revised Selected Papers. Lecture Notes in Computer Science, vol. 12020, pp. 443–451. Springer (2019). https://doi.org/10.1007/978-3-030-42921-8_26
16. Chen, Y., Zhao, Y.: Half-aggregation of schnorr signatures with tight reductions. In: Atluri, V., Di Pietro, R., Jensen, C.D., Meng, W. (eds.) ESORICS 2022, Part II. LNCS, vol. 13555, pp. 385–404. Springer, Heidelberg (Sep 2022). https://doi.org/10.1007/978-3-031-17146-8_19
17. Debris-Alazard, T., Sendrier, N., Tillich, J.P.: Wave: A new family of trapdoor one-way preimage sampleable functions based on codes. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part I. LNCS, vol. 11921, pp. 21–51. Springer, Heidelberg (Dec 2019). https://doi.org/10.1007/978-3-030-34578-5_2
18. Devadas, L., Goyal, R., Kalai, Y., Vaikuntanathan, V.: Rate-1 non-interactive arguments for batch-NP and applications. In: 63rd FOCS. pp. 1057–1068. IEEE Computer Society Press (Oct / Nov 2022). <https://doi.org/10.1109/FOCS54457.2022.00103>
19. El Bansarkhani, R., Buchmann, J.: Towards lattice based aggregate signatures. In: Pointcheval, D., Vergnaud, D. (eds.) AFRICACRYPT 14. LNCS, vol. 8469, pp. 336–355. Springer, Heidelberg (May 2014). https://doi.org/10.1007/978-3-319-06734-6_21
20. El Bansarkhani, R., Mohamed, M.S.E., Petzoldt, A.: MQSAS - A multivariate sequential aggregate signature scheme. In: Bishop, M., Nascimento, A.C.A. (eds.) ISC 2016. LNCS, vol. 9866, pp. 426–439. Springer, Heidelberg (Sep 2016). https://doi.org/10.1007/978-3-319-45871-7_25
21. Fischlin, M., Lehmann, A., Schröder, D.: History-free sequential aggregate signatures. In: Visconti, I., Prisco, R.D. (eds.) SCN 12. LNCS, vol. 7485, pp. 113–130. Springer, Heidelberg (Sep 2012). https://doi.org/10.1007/978-3-642-32928-9_7
22. Fleischhacker, N., Simkin, M., Zhang, Z.: Squirrel: Efficient synchronized multi-signatures from lattices. In: Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) ACM CCS 2022. pp. 1109–1123. ACM Press (Nov 2022). <https://doi.org/10.1145/3548606.3560655>
23. Gentry, C., O’Neill, A., Reyzin, L.: A unified framework for trapdoor-permutation-based sequential aggregate signatures. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 34–57. Springer, Heidelberg (Mar 2018). https://doi.org/10.1007/978-3-319-76581-5_2
24. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press (May 2008). <https://doi.org/10.1145/1374376.1374407>
25. Gentry, C., Ramzan, Z.: Identity-based aggregate signatures. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 257–273. Springer, Heidelberg (Apr 2006). https://doi.org/10.1007/11745853_17
26. Goubin, L., Cogliati, B., Faugère, J., Fouque, P., Larrieu, R., Macario-Rat, G., Minaud, B., Patarin, J.: PROV — P_Rovable unbalanced Oil and Vinegar. Tech.

- rep., National Institute of Standards and Technology (2023), available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>
27. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced Oil and Vinegar signature schemes. In: Stern, J. (ed.) EUROCRYPT'99. LNCS, vol. 1592, pp. 206–222. Springer, Heidelberg (May 1999). https://doi.org/10.1007/3-540-48910-X_15
 28. Kosuge, H., Xagawa, K.: Probabilistic hash-and-sign with retry in the quantum random oracle model. Cryptology ePrint Archive, Report 2022/1359 (2022), <https://eprint.iacr.org/2022/1359>
 29. Levitskaya, A.: Systems of random equations over finite algebraic structures. *Cybernetics and Systems Analysis* **41**, 67–93 (2005)
 30. Lu, S., Ostrovsky, R., Sahai, A., Shacham, H., Waters, B.: Sequential aggregate signatures and multisignatures without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 465–485. Springer, Heidelberg (May / Jun 2006). https://doi.org/10.1007/11761679_28
 31. Lysyanskaya, A., Micali, S., Reyzin, L., Shacham, H.: Sequential aggregate signatures from trapdoor permutations. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 74–90. Springer, Heidelberg (May 2004). https://doi.org/10.1007/978-3-540-24676-3_5
 32. Neven, G.: Efficient sequential aggregate signed data. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 52–69. Springer, Heidelberg (Apr 2008). https://doi.org/10.1007/978-3-540-78967-3_4
 33. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>
 34. Sakumoto, K., Shirai, T., Hiwatari, H.: On provable security of UOV and HFE signature schemes against chosen-message attack. In: Yang, B.Y. (ed.) Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011. pp. 68–82. Springer, Heidelberg (Nov / Dec 2011). https://doi.org/10.1007/978-3-642-25405-5_5
 35. Wang, Z., Wu, Q.: A practical lattice-based sequential aggregate signature. In: Steinfeld, R., Yuen, T.H. (eds.) ProvSec 2019. LNCS, vol. 11821, pp. 94–109. Springer, Heidelberg (Oct 2019). https://doi.org/10.1007/978-3-030-31919-9_6

A Missing Proofs

Lemma 2. *When a new node is added to the HTree as a result to a call to H, the additional value y' is chosen uniformly at random from \mathcal{Y} .*

Proof. When a new node is added to the HTree on Line 23 of H, there are two possibilities for the additional value y' . In both cases, y' is chosen uniformly at random from \mathcal{Y} and is independent of the view of \mathcal{A} . In fact, whenever the query to H is not the special random guess c^* chosen by the simulator, we have $y' \leftarrow \mathbf{G}(h') \oplus \alpha$. Here, $\mathbf{G}(h')$ is guaranteed to be a fresh uniformly random value since, otherwise, H would abort on Line 16 and the node would not be added to the HTree. If, on the other hand, the query c^* was made to H, then we set $y' \leftarrow y^*$ for one of the new nodes to be added. Since c^* was chosen randomly among all queries to H, the assignment of y^* is made independently of the view of \mathcal{A} and previous interactions with H. \square

Lemma 3. *For any $k > \psi$ functions $F_1, \dots, F_k: \mathcal{X} \rightarrow \mathcal{Y}$ and uniformly random $y_1, \dots, y_k \in \mathcal{Y}$, there exists $x \in \mathcal{X}$ such that $F_i(x) = y_i$, for every $i = 1, \dots, k$, with probability at most $|\mathcal{X}|/|\mathcal{Y}|^k$.*

Proof. Let $S_y^F = \{x \in \mathcal{X} : F(x) = y\}$ be the set of preimages of y under F . For a random choice of y_1 it holds that $|S_{y_1}^{F_1}| = \alpha$ for some $0 \leq \alpha \leq |\mathcal{X}|$. Then, there are at most α possible values for the tuple (y_2, \dots, y_k) , corresponding to $\{(F_2(x), \dots, F_k(x)) : x \in S_{y_1}^{F_1}\}$, such that $\bigcap S_{y_i}^{F_i} \neq \emptyset$. Since y_2, \dots, y_k are uniformly chosen in \mathcal{Y} , the probability of a non-empty intersection is at most $\alpha|\mathcal{Y}|^{1-k}$. Therefore, the desired probability is bounded by varying over the possible values of α :

$$\sum_{\alpha=0}^{|\mathcal{X}|} \frac{\alpha}{|\mathcal{Y}|^{k-1}} \Pr_{y_1 \leftarrow \mathcal{Y}}[|S_{y_1}^{F_1}| = \alpha] = \frac{1}{|\mathcal{Y}|^{k-1}} \sum_{\alpha=0}^{|\mathcal{X}|} \alpha \cdot \frac{|\{y_1 \in \mathcal{Y} : |S_{y_1}^{F_1}| = \alpha\}|}{|\mathcal{Y}|} = \frac{|\mathcal{X}|}{|\mathcal{Y}|^k}.$$

\square

Lemma 4. *If an input Q has not been entered in the HTree after being queried to H, the probability that it will ever become tethered to a node in HTree is at most $\psi \mathfrak{q}'/|\mathcal{Y}|$, where \mathfrak{q}' is the number of queries made to H after Q .*

Proof. Suppose that $Q = (F, m, r, x)$ was queried to H and was not added to the HTree, i.e. $\text{Lookup}(x) = \perp$. Now suppose that a query $Q' = (F', m', r', x')$ was subsequently sent to H and was added to HTree as part of a node N' with additional value y' . For Q to be tethered to N' , it must hold that $F'(x) = y'$. Following Lemma 2, when a new node is added to the HTree as a result to a call to H, the additional value y' is chosen uniformly at random from \mathcal{Y} . In particular, y' is random and independent of F' and x . Therefore, the probability of having $F'(x) = y'$ is $|\mathcal{Y}|^{-1}$. Since there are at most \mathfrak{q}' queries to H after Q and each query can add at most ψ nodes to the HTree, the desired probability follows by the union bound. \square

Game 4: Games for OAggSign($m, \rho = (h, x)$)		
Game ₀ :	Game ₁ -Game ₂ :	Game ₃ -Game ₅ :
1: $(\alpha, \beta) \leftarrow \text{enc}(x)$	1: $(\alpha, \beta) \leftarrow \text{enc}(x)$	1: $(\alpha, \beta) \leftarrow \text{enc}(x)$
2: repeat	2: repeat	2: $r \leftarrow_s \{0, 1\}^\lambda$
3: $r \leftarrow_s \{0, 1\}^\lambda$	3: $r \leftarrow_s \{0, 1\}^\lambda$	3: if $\text{HT}[\mathbf{F}^*, m, r, x] \neq \perp$
4: $\eta \leftarrow \mathbf{H}(\mathbf{F}^*, m, r, x)$	4: if $\text{HT}[\mathbf{F}^*, m, r, x] \neq \perp$ then	4: raise bad_{hcol}
5: $h' \leftarrow h \oplus \eta$	5: raise bad_{hcol}	5: $\eta \leftarrow_s \{0, 1\}^{2\lambda}$
6: $g' \leftarrow \mathbf{G}(h')$	6: $\eta \leftarrow_s \{0, 1\}^{2\lambda}$	6: $\text{HT}[\mathbf{F}^*, m, r, x] \leftarrow \eta$
7: $y' \leftarrow g' \oplus \alpha$	7: $\text{HT}[\mathbf{F}^*, m, r, x] \leftarrow \eta$	7: $h' \leftarrow h \oplus \eta$
8: $x' \leftarrow \mathbf{I}^*(y')$	8: $h' \leftarrow h \oplus \eta$	8: if $\text{GT}[h'] \neq \perp$ then
9: until $x' \neq \perp$	9: if $\text{GT}[h'] \neq \perp$ then	9: raise $\text{bad}_{\text{gcol1}}$
10: return $(r, \beta), (h', x')$	10: raise $\text{bad}_{\text{gcol1}}$	10: $x' \leftarrow \text{SampDom}(\mathbf{F}^*)$
	11: $y' \leftarrow_s \mathcal{Y}$	11: $y' \leftarrow \mathbf{F}^*(x')$
	12: $\text{GT}[h'] \leftarrow y' \oplus \alpha$	12: $\text{GT}[h'] \leftarrow y' \oplus \alpha$
	13: $x' \leftarrow \mathbf{I}^*(y')$	13: return $(r, \beta), (h', x')$
	14: until $x' \neq \perp$	
	15: return $(r, \beta), (h', x')$	

A.1 Proof for strong PS-HF-UF-CMA security (Theorem 1)

Proof. We prove the reduction by presenting a sequence of hybrid games, modifying the strong PS-HF-UF-CMA game (Game 2) until it can be simulated by the OW adversary \mathcal{B} . In the following, we use the notation $\Pr[\text{Game}_n(\mathcal{A}) = 1]$ to denote the probability that Game_n returns 1 when playing by \mathcal{A} . The game sequence Game₀-Game₃ for OAggSign is detailed in Game 4. The game sequence Game₃-Game₅ for H is detailed in Game 5.

Game₀ This is the original strong PS-HF-UF-CMA game against the HaS-HF-SAS scheme except that it uses programmable random oracles. At the start of the game, the challenger initializes two tables, HT for H and GT for G. When a query Q for H is received, if $\text{HT}[Q] = \perp$ it uniformly samples $\eta \leftarrow_s \{0, 1\}^{2\lambda}$ and stores $\text{HT}[Q] \leftarrow \eta$, finally it returns $\text{HT}[Q]$ (similarly for G). It follows that $\Pr[\text{Game}_0(\mathcal{A}) = 1] = \text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A})$.

Game₁ This game is identical to Game₀ except that OAggSign aborts by raising bad_{hcol} if on query $(m, \rho = (h, x))$ it samples a salt r such that the random oracle H was already queried at input $Q = (\mathbf{F}^*, m, r, x)$, i.e. $\text{HT}[Q] \neq \perp$. Otherwise it samples $\eta \leftarrow_s \{0, 1\}^{2\lambda}$ and programs $\text{HT}[Q] \leftarrow \eta$. It follows that $|\Pr[\text{Game}_0(\mathcal{A}) = 1] - \Pr[\text{Game}_1(\mathcal{A}) = 1]| \leq \Pr[\text{bad}_{\text{hcol}}]$.

Game₂ This game is identical to Game₁ except that OAggSign aborts by raising $\text{bad}_{\text{gcol1}}$ if on query $(m, \rho = (h, x))$, after sampling $\eta \leftarrow_s \{0, 1\}^{2\lambda}$ it computes $h' \leftarrow h \oplus \eta$ such that the random oracle G was already queried at input h' , i.e. $\text{GT}[h'] \neq \perp$. Otherwise it samples $y' \leftarrow_s \mathcal{Y}$ and programs $\text{GT}[h'] \leftarrow y' \oplus \alpha$. It follows that $|\Pr[\text{Game}_1(\mathcal{A}) = 1] - \Pr[\text{Game}_2(\mathcal{A}) = 1]| \leq \Pr[\text{bad}_{\text{gcol1}}]$.

Game₃ This game is identical to **Game₂** except that **OAggSign** directly samples $r \leftarrow_{\$} \{0,1\}^\lambda$, $x' \leftarrow \text{SampDom}(F^*)$ and computes $y' \leftarrow F^*(x')$ instead of computing $x' \leftarrow I^*(y')$ after sampling $y' \leftarrow_{\$} \mathcal{Y}$. The PS adversary \mathcal{D} can simulate both **Game₂** and **Game₃**, noticing that $y' = F^*(x')$ and programming **G** accordingly. More precisely, on receiving a query $Q = (m, \varrho = (h, x))$ for **OAggSign**, \mathcal{D} computes $(r, x') \leftarrow \text{Sample}_b$ and programs $\text{GT}[h'] \leftarrow F^*(x') \oplus \alpha$. Both **Game₂** and **Game₃** are equivalently modified by moving the programming step of **H** and **G** to the end of the **OAggSign**. It now follows that when \mathcal{D} is playing **PS₀** its simulation coincides with **Game₂**, while when it is playing **PS₁** it coincides with **Game₃**. Either way, \mathcal{D} simulates the games with at most the same running time of \mathcal{A} plus the time required for answering the queries to the sampling oracle. The latter takes $\mathcal{O}(\text{poly}(\text{len}(\mathcal{X}), \text{len}(\mathcal{Y})))$ and is repeated at most q_S times. Finally, we have that $|\Pr[\text{Game}_2(\mathcal{A}) = 1] - \Pr[\text{Game}_3(\mathcal{A}) = 1]| \leq \text{Adv}_T^{\text{PS}}(\mathcal{D})$.

Game₄ This game is identical to **Game₃** except that the random oracle **H** is simulated as follows. At the start of the game, the challenger initializes a labeled tree **HTree**, as described at the beginning of the proof. When **H** receives a query $Q = (F, m, r, x)$, if $\text{HT}[Q] \neq \perp$ it returns it. Otherwise, it samples a uniformly random $\eta \leftarrow_{\$} \{0,1\}^{2\lambda}$ and programs $\text{HT}[Q] \leftarrow \eta$. Then, it checks if Q can be added as a child node of existing nodes in **HTree**. To determine whether this is the case, it uses the **Lookup** function (see Algorithm 3) on input x that checks if it can be tethered to existing nodes, i.e. there exists a node $N_i \in \text{HTree}$ such that $F_i(x) = y_i$. If Q can be tethered to more than ψ nodes, the game aborts by raising bad_{tcol} . Otherwise, **H** add a new node N'_i with parent N_i for each node $N_i \in \text{HTree}$ returned by **Lookup**(x). N'_i contains the original query Q , the hash response η , the hash state $h'_i \leftarrow h_i \oplus \eta$ (where h_i is stored in N_i) and an additional value $y'_i \leftarrow G(h'_i) \oplus \alpha$ (where α is computed from $\text{enc}(x)$) that will be used to check if a future node can be tethered via **Lookup** queries. It holds that $|\Pr[\text{Game}_3(\mathcal{A}) = 1] - \Pr[\text{Game}_4(\mathcal{A}) = 1]| \leq \Pr[\text{bad}_{\text{tcol}}]$.

Game₅ This game is identical to **Game₄** except that the random oracle **H** is simulated as follows. At the beginning of the game, the challenger uniformly chooses an index $c^* \leftarrow_{\$} [q_H]$ among the queries to the random oracle **H**, initializes a counter $c \leftarrow 0$ and uniformly samples $y^* \leftarrow_{\$} \mathcal{Y}$. When **H** receives a query $Q = (F, m, r, x)$ it increments $c \leftarrow c + 1$. Then, if $F = F^*$ and $c = c^*$ it samples a random index i^* from the number of nodes in **NList**. If, for any of the new nodes to be added, it computes $h'_i \leftarrow h_i \oplus \eta$ such that the random oracle **G** was already queried at input h'_i , i.e. $\text{GT}[h'_i] \neq \perp$, it aborts by raising $\text{bad}_{\text{gcol2}}$. Otherwise, if $F = F^*$, $c = c^*$ and $i = i^*$, it sets $y'_i \leftarrow y^*$ and programs $\text{GT}[h'_i] \leftarrow y'_i \oplus \alpha$. It holds that $|\Pr[\text{Game}_4(\mathcal{A}) = 1] - \Pr[\text{Game}_5(\mathcal{A}) = 1]| \leq \Pr[\text{bad}_{\text{gcol2}}]$.

Game 5: Games for $H(F, m, r, x)$ **Game₀-Game₃:**

```

1:  $Q \leftarrow (F, m, r, x)$ 
2: if  $\text{HT}[Q] = \perp$  then
3:    $\eta \leftarrow_s \{0, 1\}^{2\lambda}$ 
4:    $\text{HT}[Q] \leftarrow \eta$ 
5: return  $\text{HT}[Q]$ 

```

Game₄:

```

1:  $Q \leftarrow (F, m, r, x)$ 
2: if  $\text{HT}[Q] = \perp$  then
3:    $\eta \leftarrow_s \{0, 1\}^{2\lambda}$ 
4:    $\text{HT}[Q] \leftarrow \eta$ 
5:    $\text{NList} \leftarrow \text{Lookup}(x)$ 
6:   for  $i \in [|\text{NList}|]$  do
7:      $N_i \leftarrow \text{NList}[i]$ 
8:      $N'_i \leftarrow$  new node with parent  $N_i$ 
9:     Retrieve  $h_i$  from  $N_i$ 
10:     $h'_i \leftarrow h_i \oplus \eta$ 
11:     $(\alpha, \beta) \leftarrow \text{enc}(x)$ 
12:     $g'_i \leftarrow G(h'_i)$ 
13:     $y'_i \leftarrow g'_i \oplus \alpha$ 
14:     $N'_i \leftarrow (Q, \eta, h'_i, y'_i)$ 
15: return  $\text{HT}[Q]$ 

```

Game₅:

```

1:  $Q \leftarrow (F, m, r, x)$ 
2:  $c \leftarrow c + 1$ 
3: if  $\text{HT}[Q] = \perp$  then
4:    $\eta \leftarrow_s \{0, 1\}^{2\lambda}$ 
5:    $\text{HT}[Q] \leftarrow \eta$ 
6:    $\text{NList} \leftarrow \text{Lookup}(x)$ 
7:   if  $F = F^* \wedge c = c^*$  then
8:      $i^* \leftarrow_s [|\text{NList}|]$ 
9:     for  $i \in [|\text{NList}|]$  do
10:       $N_i \leftarrow \text{NList}[i]$ 
11:       $N'_i \leftarrow$  new node with parent  $N_i$ 
12:      Retrieve  $h_i$  from  $N_i$ 
13:       $h'_i \leftarrow h_i \oplus \eta$ 
14:       $(\alpha, \beta) \leftarrow \text{enc}(x)$ 
15:      if  $\text{GT}[h'_i] \neq \perp$  then
16:        raise  $\text{bad}_{\text{gcol2}}$ 
17:      if  $F \neq F^* \vee c \neq c^* \vee i \neq i^*$ 
18:        then
19:           $g'_i \leftarrow G(h'_i)$ 
20:           $y'_i \leftarrow g'_i \oplus \alpha$ 
21:        else
22:           $y'_i \leftarrow y^*$ 
23:           $\text{GT}[h'_i] \leftarrow y'_i \oplus \alpha$ 
24:           $N'_i \leftarrow (Q, \eta, h'_i, y'_i)$ 
25:      return  $\text{HT}[Q]$ 

```

If none of the bad events happen, \mathcal{B} perfectly simulate Game_5 and we have that

$$\begin{aligned}
\text{Adv}_{\mathcal{T}}^{\text{OW}}(\mathcal{B}) &= \frac{1}{\psi_{\text{qH}}} \Pr[\text{Game}_5(\mathcal{A}) = 1] \\
&\geq \frac{1}{\psi_{\text{qH}}} (\text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A}) - \Pr[\text{bad}_{\text{hcol}}] - \Pr[\text{bad}_{\text{gcol1}}] \\
&\quad - \text{Adv}_{\mathcal{T}}^{\text{PS}}(\mathcal{D}) - \Pr[\text{bad}_{\text{tcol}}] - \Pr[\text{bad}_{\text{gcol2}}] - \Pr[\text{bad}_{\text{teth}}]).
\end{aligned}$$

\mathcal{B} can simulate Game_5 with at most the same running time of \mathcal{A} plus the time required for running AggVrfy and answering the queries to the random oracles H, G , and to the signing oracle OAggSign . These operations takes $\mathcal{O}(\text{poly}(\text{len}(\mathcal{X}), \text{len}(\mathcal{Y})))$ and are repeated at most $\text{qH} + \text{qS} + 1$ times.

In the following, we bound the probability of each bad event happening.

Probability of bad_{hcol} The event bad_{hcol} occurs on Line 5 of OAggSign on input $(m, \varrho = (h, x))$ when it samples $r \leftarrow_s \{0, 1\}^\lambda$ such that a value for $Q =$

(F^*, m, r, x) was already assigned in the HT. The table HT is populated by either OAggSign or H, so its entries are at most $q'_S + q_H$. The probability that a uniformly random r produces a collision with one of the entries is then at most $(q'_S + q_H)2^{-\lambda}$. Since at most q_S are made to OAggSign, then $\Pr[\text{bad}_{\text{hcol}}] \leq q_S(q'_S + q_H)2^{-\lambda}$.

Probability of $\text{bad}_{\text{gcol1}}$ The event $\text{bad}_{\text{gcol1}}$ occurs on Line 10 of OAggSign on input $(m, \varrho = (h, x))$ when, after sampling $\eta \leftarrow_{\$} \{0, 1\}^{2\lambda}$, it computes $h' \leftarrow h \oplus \eta$ such that a value for h' was already assigned in the GT. The table GT is populated by either OAggSign, H or G so its entries are at most $q'_S + q_H + q_G$. The probability that a uniformly random η produces a collision with one of the entries is then at most $(q'_S + q_H + q_G)2^{-2\lambda}$. Since at most q_S are made to OAggSign, then $\Pr[\text{bad}_{\text{gcol1}}] \leq q_S(q'_S + q_H + q_G)2^{-2\lambda}$.

Probability of bad_{tcol} The event bad_{tcol} occurs on Line 5 of Lookup on input x when the HTree contains $k > \psi$ nodes N_1, \dots, N_k such that $F_i(x) = y_i$ for $i = 1, \dots, k$, where F_i, y_i are stored in their respective nodes N_i . The HTree is populated by the simulation of the random oracle H. There are at most q_H queries to H and each query contributes a maximum of ψ nodes to the tree. Consequently, the total number of nodes in HTree does not exceed ψq_H . Therefore, we need to bound the probability that any $(\psi + 1)$ -tuple of nodes produce a collision on x .

To conclude, we prove that for any $(\psi + 1)$ -tuple (possibly adversarially chosen) of functions $F_i: \mathcal{X} \rightarrow \mathcal{Y}$ and uniformly random $y_i \in \mathcal{Y}$, there exists $x \in \mathcal{X}$ such that $F_i(x) = y_i$, for any $i = 1, \dots, \psi + 1$, with probability at most $|\mathcal{X}|/|\mathcal{Y}|^{\psi+1}$ (Lemma 3). Indeed, the adversary can issue $\psi + 1$ queries to H with inputs any functions F_i to be stored in $\psi + 1$ nodes N_i in the HTree. However, from Lemma 2, we know that when a new node is added to the HTree on Line 23 of H, the value y'_i is chosen uniformly at random from \mathcal{Y} and is independent of the view of \mathcal{A} . Therefore, the adversary would receive $\psi + 1$ random, independent values y_i .

Since the number of $(\psi + 1)$ -tuple of nodes in the HTree are at most $(\psi q_H)^{\psi+1}/(\psi + 1)!$, by the union bound, we obtain $\Pr[\text{bad}_{\text{tcol}}] \leq (\psi q_H)^{\psi+1} |\mathcal{X}| / ((\psi + 1)! \cdot |\mathcal{Y}|^{\psi+1})$.

Probability of $\text{bad}_{\text{gcol2}}$ The event $\text{bad}_{\text{gcol2}}$ occurs on Line 16 of H on input (F, m, r, x) when, after sampling $\eta \leftarrow_{\$} \{0, 1\}^{2\lambda}$ and retrieving h_{i-1} from the parent node N_{i-1} , it computes $h_i \leftarrow h_{i-1} \oplus \eta$ such that a value for h_i was already assigned in the GT. The same argument from the bound of $\Pr[\text{bad}_{\text{gcol}}]$ can be used to prove that $\Pr[\text{bad}_{\text{gcol2}}] \leq q_H(q'_S + q_H + q_G)2^{-2\lambda}$.

Probability of bad_{teth} The event bad_{teth} occurs on Line 8 of the simulation of \mathcal{B} when, after the adversary \mathcal{A} outputs a valid aggregate signature $\bar{\Sigma}_n$ for the history $L_n = (\text{pk}_1, m_1), \dots, (\text{pk}_n, m_n)$ the simulator recovers x_{i^*} , with $i^* \in [n]$ such that $\text{pk}_{i^*} = \text{pk}^*$ and $(m_{i^*}, \varsigma_{i^*}) \notin \mathcal{Q}$, but x_{i^*} cannot be tethered to any node in the HTree.

When bad_{teth} happens, the aggregate signature $\bar{\Sigma}_n$ must be valid on L_n . In particular, the inputs $Q_1 = (F_1, m_1, r_1, \varepsilon), Q_2 = (F_2, m_2, r_2, x_1), \dots, Q_{i^*} = (F_{i^*}, m_{i^*}, r_{i^*}, x_{i^*-1})$ have been queried to H in OAggVrfy. Let $\eta_1, \dots, \eta_{i^*}$ be the outputs of these queries, so that $\text{HT}[Q_j] = \eta_j$. Each of these entries has

been populated by H . In fact, the only exception could occur if (m_{i^*}, x_{i^*-1}) was queried to OAggSign . Suppose (r, β) is the complementary part of the signature produced by the oracle as a response. Since the forgery is valid, the complementary part $\varsigma_{i^*} = (r_{i^*}, \beta_{i^*-1})$ produced by \mathcal{A} must be different from (r, β) . However, both β_{i^*-i} and β must be the same partial encoding of x_{i^*-1} , so that $r_{i^*} \neq r$. Therefore, x_{i^*} must have been produced following a query to H with a fresh salt r_{i^*} .

Each step of OAggVrfy also recovers a value $h_j \leftarrow h_{j+1} \oplus \eta_j$ which is the input of the G query. Since the aggregate signature is correct, we obtain that $h_1 = \eta_1$. Observe that since Q_1 was queried to H , it must be tethered to the root of HTree and was therefore inserted as a node of HTree with additional values $\eta_1, h_1 = \eta_1, y_1 = G(h_1)$. Then, since $F(x_1) = y_1$, the query Q_2 is tethered to N_1 . Now we prove that either all Q_1, \dots, Q_{i^*} are part of a path of nodes in HTree , or there exists an input Q_j that was queried to H , is tethered to a node in HTree and is not itself in a node of HTree . We proceed by induction on $j \leq i^*$: we have already shown that Q_1 is in HTree ; suppose that Q_j is in the HTree , then, since $F_j(x_j) = y_j$, the query Q_{j+1} is tethered to Q_j and it may or may not be part of HTree . To conclude, we prove that if an input Q has not been entered in the HTree after being queried to H , the probability that it will ever become tethered to a node in HTree is at most $\psi \mathbf{q}' / |\mathcal{Y}|$, where \mathbf{q}' is the number of queries made to H after Q (Lemma 4). Since there are at most \mathbf{q}_H queries that add new nodes to HTree , we obtain, by the union bound, that $\Pr[\text{bad}_{\text{teth}}] \leq \psi \mathbf{q}_H^2 / (2|\mathcal{Y}|)$.

Combining the previous bound on bad events, we obtain the claimed estimate of $\text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A})$. \square

B Trapdoor Functions in Hash-and-Sign Signature Schemes

B.1 UOV Trapdoor Function

MQ-based trapdoor function consists of a multivariate quadratic map $\mathcal{P}: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ together with a secret information that allows to efficiently find a preimage. For a random map \mathcal{P} , the problem of finding a preimage is called *Multivariate Quadratic (MQ) problem*. The MQ problem is NP-hard over a finite field. Moreover, it is believed to be hard on average if $n \sim m$, both classically and quantumly.

Both UOV and MAYO are based on the same trapdoor function. For the description of the trapdoor function we mainly use the formalism introduced by Beullens in [7].

The trapdoor secret information is a linear subspace $O \subset \mathbb{F}_q^n$ of dimension $\dim(O) = m$. The trapdoor public function is a homogeneous multivariate quadratic map $\mathcal{P}: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on O . For key generation, an m -dimensional subspace $O \subset \mathbb{F}_q^n$ is randomly chosen, then a multivariate quadratic map $\mathcal{P}: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ is randomly chosen such that it vanishes on O . Given a

Algorithm 5: UOV Signature Scheme

<p>TrapGen_{uov}(1^λ):</p> <ol style="list-style-type: none"> 1: $O \leftarrow_s$ m-dimensional subspace of \mathbb{F}_q^m 2: $\mathcal{P} \leftarrow_s$ quadratic map $\mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on O 3: return $F_{uov} \leftarrow \mathcal{P}, I_{uov} \leftarrow (\mathcal{P}, O)$ <p>Sign(I_{uov}, m):</p> <ol style="list-style-type: none"> 1: $r \leftarrow_s \{0, 1\}^\lambda$ 2: $\sigma \leftarrow I_{uov}(H(m, r))$ 3: return (r, σ) 	<p>I_{uov}(\mathbf{t}):</p> <ol style="list-style-type: none"> 1: repeat 2: $\mathbf{v} \leftarrow_s \mathbb{F}_q^n$ 3: until $\{\mathbf{o} \in O \mid \mathcal{P}'(\mathbf{v}, \mathbf{o}) = \mathbf{t} - \mathcal{P}(\mathbf{v})\} \neq \emptyset$ 4: $\mathbf{o} \leftarrow_s \{\mathbf{o} \mid \mathcal{P}'(\mathbf{v}, \mathbf{o}) = \mathbf{t} - \mathcal{P}(\mathbf{v})\}$ 5: $\sigma \leftarrow \mathbf{v} + \mathbf{o}$ 6: return σ
--	--

target $\mathbf{t} \in \mathbb{F}_q^m$, the secret information O can be used to find a preimage $\mathbf{s} \in \mathbb{F}_q^n$, reducing the MQ problem to a linear system. For a map \mathcal{P} , we can define its *polar form* as $\mathcal{P}'(\mathbf{x}, \mathbf{y}) = \mathcal{P}(\mathbf{x} + \mathbf{y}) - \mathcal{P}(\mathbf{x}) - \mathcal{P}(\mathbf{y})$. It can be shown that the polar form of a multivariate quadratic map is a symmetric and bilinear map. Now, to find a preimage for \mathbf{t} , one randomly choose a vector $\mathbf{v} \in \mathbb{F}_q^n$ and solves $\mathcal{P}'(\mathbf{v} + \mathbf{o}) = \mathbf{t}$ for $\mathbf{o} \in O$. Since

$$\mathbf{t} = \mathcal{P}'(\mathbf{v} + \mathbf{o}) = \underbrace{\mathcal{P}'(\mathbf{v}, \mathbf{v})}_{\text{fixed}} + \underbrace{\mathcal{P}'(\mathbf{v}, \mathbf{o})}_{=0} + \underbrace{\mathcal{P}'(\mathbf{o}, \mathbf{o})}_{\text{linear in } \mathbf{o}},$$

the system reduce to the linear system $\mathcal{P}'(\mathbf{v}, \mathbf{o}) = \mathbf{t} - \mathcal{P}'(\mathbf{v}, \mathbf{v})$ of m equation and m variables \mathbf{o} . Notice that whenever the linear map $\mathcal{P}'(\mathbf{v}, \cdot)$ is non-singular⁷, the system has a unique solution $\mathbf{o} \in O$ and the preimage is $\mathbf{s} = \mathbf{v} + \mathbf{o}$.

B.2 Original Unbalanced Oil and Vinegar

Let $\mathsf{T}_{uov} = (\text{TrapGen}_{uov}, F_{uov}, I_{uov})$ be the TDF based on the description of the previous section. Unbalanced Oil and Vinegar (UOV) [27] is a HaS signature scheme based on T_{uov} . The key generation and the signing procedure with the trapdoor functions are shown in Algorithm 5.

In the original version of the UOV signature, the signer samples a random salt $r \leftarrow_s \{0, 1\}^\lambda$ and repeatedly samples $\mathbf{v} \leftarrow_s \mathbb{F}_q^n$ until there is a solution to the linear system $\mathcal{P}'(\mathbf{v}, \mathbf{o}) = H(m, r) - \mathcal{P}(\mathbf{v})$. Notice that the UOV signature lies in the HaS *without* retry paradigm, therefore $\mathbf{q}'_S = \mathbf{q}_S$ holds in Theorem 1. On the other hand, the PS advantage term $\text{Adv}_{\mathsf{T}_{uov}}^{\text{PS}}(\mathcal{D})$ cannot be omitted since signature simulation requires knowledge of the trapdoor function.

Corollary 1. *Let \mathcal{A} be a strong PS-HF-UF-CMA adversary against the HaS-HF-SAS scheme on T_{uov} in the random oracle model, which makes q_S signing queries, q_H queries to the random oracle H and q_G queries to the random oracle G . Then,*

⁷ This happens with probability approximately $1 - 1/q$

Algorithm 6: Provable UOV Signature Scheme

<p>TrapGen_{puov}(1^λ):</p> <ol style="list-style-type: none"> 1: $O \leftarrow_{\\$} o$-dimensional subspace of \mathbb{F}_q^n 2: $\mathcal{P} \leftarrow_{\\$}$ quadratic map $\mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on O 3: return $F_{\text{puov}} \leftarrow \mathcal{P}, I_{\text{puov}} \leftarrow (\mathcal{P}, O)$ <p>I_{puov}²(\mathbf{v}, \mathbf{t}):</p> <ol style="list-style-type: none"> 1: if $\{\mathbf{o} \in O \mid \mathcal{P}'(\mathbf{v}, \mathbf{o}) = \mathbf{t} - \mathcal{P}(\mathbf{v})\} \neq \emptyset$ then return \perp 2: $\mathbf{o} \leftarrow_{\\$} \{\mathbf{o} \mid \mathcal{P}'(\mathbf{v}, \mathbf{o}) = \mathbf{t} - \mathcal{P}(\mathbf{v})\}$ 3: $\boldsymbol{\sigma} \leftarrow \mathbf{v} + \mathbf{o}$ 4: return $\boldsymbol{\sigma}$ 	<p>I_{puov}¹(\cdot):</p> <ol style="list-style-type: none"> 1: $\mathbf{v} \leftarrow_{\\$} \mathbb{F}_q^n$ 2: return \mathbf{v} <p>Sign(I_{puov}, m):</p> <ol style="list-style-type: none"> 1: $\mathbf{v} \leftarrow I_{\text{puov}}^1(\cdot)$ 2: repeat 3: $r \leftarrow_{\\$} \{0, 1\}^\lambda$ 4: $\boldsymbol{\sigma} \leftarrow I_{\text{puov}}^2(\mathbf{v}, H(m, r))$ 5: until $\boldsymbol{\sigma} \neq \perp$ 6: return $(r, \boldsymbol{\sigma})$
---	---

there exist a OW adversary \mathcal{B} against T_{uov} , and a PS adversary \mathcal{D} against T_{uov} issuing q_S sampling queries, such that

$$\begin{aligned} \text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A}) &\leq (\psi q_H) \cdot \text{Adv}_{T_{\text{uov}}}^{\text{OW}}(\mathcal{B}) + \text{Adv}_{T_{\text{uov}}}^{\text{PS}}(\mathcal{D}) + \frac{(q_S + q_H)(q_S + q_H + q_G)}{2^{2\lambda}} \\ &\quad + \frac{q_S(q_S + q_H)}{2^\lambda} + \frac{\psi q_H^2}{2|\mathcal{Y}|} + \frac{(\psi q_H)^{\psi+1} |\mathcal{X}|}{(\psi + 1)! \cdot |\mathcal{Y}|^{\psi+1}}, \end{aligned}$$

where $\psi \geq \lceil \text{len}(\mathcal{X}) / \text{len}(\mathcal{Y}) \rceil$, and the running time of \mathcal{B} and \mathcal{D} are about that of \mathcal{A} .

The previous corollary can be applied to the UOV scheme [10] submitted to the NIST PQC Standardization of Additional Digital Signature. For typical parameters, n is chosen equal to $2.5m$. If we choose $\psi = 3$, the additive error terms in Corollary 1 are negligible for each parametrization in Table 1.

B.3 Provable Unbalanced Oil and Vinegar

By adopting the probabilistic HaS with retry paradigm, the UOV signature scheme can be proven EUF-CMA secure in the random oracle model [34]. To obtain uniform preimages over \mathbb{F}_q^n , the *provable* UOV (PUOV) signing procedure is slightly different from the generic one described in Algorithm 1. The signer starts by fixing a random $\mathbf{v} \leftarrow_{\$} \mathbb{F}_q^n$, then it repeatedly samples $r \leftarrow_{\$} \{0, 1\}^\lambda$ until there is a solution to the linear system $\mathcal{P}'(\mathbf{v}, \mathbf{o}) = H(m, r) - \mathcal{P}(\mathbf{v})$. Equivalently, the trapdoor I_{puov} can be split in two distinct functions I_{puov}^1 and I_{puov}^2 . The former is invoked only once and randomly chooses $\mathbf{v} \leftarrow_{\$} \mathbb{F}_q^n$. The latter is part of the repeat loop and tries to find a preimage \mathbf{s} of the corresponding linear system. The key generation and the signing procedure with the modified trapdoor functions are shown in Algorithm 6.

With this procedure, the authors of [34] proved that the preimages produced from $\text{Sign}(l_{\text{puov}}, \cdot)$ are indistinguishable from the output of $\text{SampDom}(F_{\text{puov}})$, so that in Theorem 1 we have $\text{Adv}_{\text{T}_{\text{puov}}}^{\text{PS}}(\mathcal{D}) = 0$.

Corollary 2. *Let \mathcal{A} be a strong PS-HF-UF-CMA adversary against the HaS-HF-SAS scheme on T_{puov} in the random oracle model, which makes q_S signing queries, q_H queries to the random oracle H and q_G queries to the random oracle G . Then, there exist a OW adversary \mathcal{B} against T_{puov} , such that*

$$\begin{aligned} \text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A}) &\leq (\psi q_H) \cdot \text{Adv}_{\text{T}_{\text{puov}}}^{\text{OW}}(\mathcal{B}) + \frac{(q_S + q_H)(q'_S + q_H + q_G)}{2^{2\lambda}} \\ &\quad + \frac{q_S(q'_S + q_H)}{2^\lambda} + \frac{\psi q_H^2}{2|\mathcal{Y}|} + \frac{(\psi q_H)^{\psi+1} |\mathcal{X}|}{(\psi+1)! \cdot |\mathcal{Y}|^{\psi+1}}, \end{aligned}$$

where $\psi \geq \lceil \text{len}(\mathcal{X}) / \text{len}(\mathcal{Y}) \rceil$, q'_S is a bound on the total number of queries to H in all the signing queries, and the running time of \mathcal{B} is about that of \mathcal{A} .

Unlike Corollary 1, we cannot explicitly take $q'_S = q_S$, since in l_{puov}^2 the probability of $\sigma \neq \perp$ depends on the fixed value of \mathbf{v} sampled in l_{puov}^1 . Depending on the concrete parameters of T_{puov} , we can give a meaningful bound on q'_S so that the probability of having a number of queries to H greater than q'_S is negligible. l_{puov}^2 returns \perp on input (\mathbf{v}, \mathbf{t}) if $\mathcal{P}'(\mathbf{v}, \cdot)$ does not have full rank and $\mathbf{t} - \mathcal{P}(\mathbf{v})$ does not belong to the image of $\mathcal{P}'(\mathbf{v}, \cdot)$. Let q_{ret} be a bound for the number of queries to H each signing query and let X_i be a random variable on the actual number of queries to H in the i -th query. Then

$$\Pr[X_i > q_{\text{ret}}] = \sum_{j=1}^m \Pr[\text{rank}(\mathcal{P}'(\mathbf{v}, \cdot)) = j] (1 - q^{j-m})^{q_{\text{ret}}}.$$

As done in [34], we can assume that for a random $\mathbf{v} \leftarrow^* \mathbb{F}_q^n$, $\mathcal{P}'(\mathbf{v}, \cdot)$ is distributed as a random $o \times m$ matrix. For $o \geq m$, the probability that a random $o \times m$ matrix over \mathbb{F}_q has rank $1 \leq j \leq m$ is given in [29]:

$$q^{(j-m)(o-j)} \frac{\prod_{k=o-j+1}^o (1 - q^{-k}) \prod_{k=m-j+1}^m (1 - q^{-k})}{\prod_{k=1}^j (1 - q^{-k})}. \quad (3)$$

Then, if we choose q_{ret} such that $q_S \Pr[X_i > q_{\text{ret}}]$ is negligible, we can use $q'_S = q_{\text{ret}} q_S$ in the bound of the corollary.

Corollary 2 can be applied to the PROV scheme [26] submitted to the NIST PQC Standardization of Additional Digital Signature. The parameters of PROV are selected so that the dimension of the trapdoor subspace is $o = m + \delta$. This choice significantly reduces the probability of Equation (3) whenever $j < m$. For instance, with the parameters of PROV-I we have $\Pr[X_i > 1] \leq 2^{-72}$ and $\Pr[X_i > 2^{14}] \leq 2^{-160}$. Similarly to Original UOV, if we choose $\psi = 3$, the additive error terms in Corollary 2 are negligible for each parametrization in [26].

Algorithm 7: MAYO Signature Scheme

<p>TrapGen_{mayo}(1^λ):</p> <ol style="list-style-type: none"> 1: $\mathbf{O} \leftarrow \mathbb{F}_q^{o \times (n-o)}$ 2: $O \leftarrow \text{RowSpace}(\mathbf{O}\mathbf{I}_o)$ 3: $\mathcal{P} \leftarrow$ quadratic map $\mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on O 4: return $\mathbf{F}_{\text{mayo}} \leftarrow \mathcal{P}, \mathbf{l}_{\text{mayo}} \leftarrow (\mathcal{P}, \mathbf{O})$ 	<p>l_{mayo}(\mathbf{t}):</p> <ol style="list-style-type: none"> 1: $\mathcal{P}^*(\mathbf{x}_1, \dots, \mathbf{x}_k) \leftarrow \sum_{i=1}^k \mathbf{E}_{i,i} \mathcal{P}(\mathbf{x}_i) + \sum_{1 \leq i < j \leq k} \mathbf{E}_{i,j} \mathcal{P}'(\mathbf{x}_i, \mathbf{x}_j)$ 2: $\mathbf{v}_1, \dots, \mathbf{v}_k \leftarrow \mathbb{F}_q^n \times \mathbf{0}_m)^k$ 3: if $\mathcal{P}^*(\mathbf{v}_1 + \mathbf{o}_1, \dots, \mathbf{v}_k + \mathbf{o}_k)$ does not have full rank then 4: return \perp 5: $\mathbf{o}_1, \dots, \mathbf{o}_k \leftarrow \{ \mathbf{o}_1, \dots, \mathbf{o}_k \in O \mid \mathcal{P}^*(\mathbf{v}_1 + \mathbf{o}_1, \dots, \mathbf{v}_k + \mathbf{o}_k) = \mathbf{t} \}$ 6: $\boldsymbol{\sigma} \leftarrow (\mathbf{v}_1 + \mathbf{o}_1, \dots, \mathbf{v}_k + \mathbf{o}_k)$ 7: return $\boldsymbol{\sigma}$
--	---

B.4 MAYO

MAYO [8] is a HaS signature scheme based on the UOV trapdoor function and employs a so-called *whipping* technique to use a smaller secret subspace O of dimension $\dim(O) = o < m$. Let $\mathbf{T}_{\text{mayo}} = (\text{TrapGen}_{\text{mayo}}, \mathbf{F}_{\text{mayo}}, \mathbf{l}_{\text{mayo}})$ be the TDF of MAYO. The key generation process is the same as for UOV and produces a multivariate quadratic map $\mathcal{P}: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on O . In the signing procedure, \mathcal{P} is deterministically transformed into a larger (whipped) map $\mathcal{P}^*: \mathbb{F}_q^{kn} \rightarrow \mathbb{F}_q^m$, for some $k > 1$, which vanishes on $O^k \subset \mathbb{F}_q^{kn}$ of dimension $ko \geq m$. In [8], the whipping transformation is obtained by choosing $k(k+1)/2$ random invertible matrices $\{ \mathbf{E}_{i,j} \in \text{GL}_m(\mathbb{F}_q) \}_{1 \leq i < j \leq k}$ and defining

$$\mathcal{P}^*(\mathbf{x}_1, \dots, \mathbf{x}_k) = \sum_{i=1}^k \mathbf{E}_{i,i} \mathcal{P}(\mathbf{x}_i) + \sum_{1 \leq i < j \leq k} \mathbf{E}_{i,j} \mathcal{P}'(\mathbf{x}_i, \mathbf{x}_j).$$

Similarly to UOV, to find a preimage for $\mathbf{t} \in \mathbb{F}_q^m$, we randomly choose $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{F}_q^{n-m} \times \mathbf{0}_m$. Then, $\mathcal{P}^*(\mathbf{v}_1 + \mathbf{o}_1, \dots, \mathbf{v}_k + \mathbf{o}_k) = \mathbf{t}$ is a system of m linear equations in $ko \geq m$ variables, so it will be solvable with high probability. The key generation and the preimage computation via \mathbf{l}_{mayo} are shown in Algorithm 7.

Instead of computing $\text{Adv}_{\mathbf{T}_{\text{mayo}}}^{\text{PS}}(\mathcal{D})$, we can use the result of [8, Lemma 2] that bounds the probability \mathbf{B} that $\mathcal{P}^*(\mathbf{v}_1 + \mathbf{o}_1, \dots, \mathbf{v}_k + \mathbf{o}_k)$ does not have full rank. It can be shown that if \mathbf{l}_{mayo} has never output \perp , then the preimages produced by $\text{Sign}(\mathbf{l}_{\text{mayo}}, \cdot)$ are indistinguishable from $\text{SampDom}(\mathbf{F}_{\text{mayo}})$. Therefore, we can modify the proof of Theorem 1 by introducing a new intermediate game Game_{2b} . This game is identical to Game_2 except that OAggSign aborts if \mathbf{l}_{mayo} outputs \perp . Since there are at most q_5 queries are made to OAggSign , the probability that Game_{2b} does not abort is at least $1 - q_5 \mathbf{B}$. It follows that $\Pr[\text{Game}_2(\mathcal{A}) = 1] \leq \frac{1}{1 - q_5 \mathbf{B}} \Pr[\text{Game}_{2b}(\mathcal{A}) = 1]$. Now, when Game_{2b} does not abort, the game is indistinguishable from Game_3 , so that $\Pr[\text{Game}_3(\mathcal{A}) = 1] = \Pr[\text{Game}_{2b}(\mathcal{A}) = 1]$. The remainder of the proof proceeds as the original. Finally, since MAYO now does not repeat any signature attempts, we can use $q'_5 = q_5$ in Theorem 1.

Algorithm 8: Wave Signature Scheme

TrapGen_{wave}(1^λ): 1: $\mathbf{H}_{\text{sk}} \in \mathbb{F}_q^{(n-k) \times n} \leftarrow_s \text{generalized}$ $(U, U + V)\text{-code}$ 2: $\mathbf{S} \leftarrow_s \text{GL}_{n-k}(\mathbb{F}_q)$ 3: $\mathbf{P} \leftarrow_s n \times n \text{ permutation matrix}$ 4: $\mathbf{H}_{\text{pk}} \leftarrow \mathbf{S}\mathbf{H}_{\text{sk}}\mathbf{P}$ 5: return $\mathbf{F}_{\text{mayo}} \leftarrow \mathbf{H}_{\text{pk}}, \mathbf{l}_{\text{mayo}} \leftarrow$ $(\mathbf{H}_{\text{sk}}, \mathbf{S}, \mathbf{P})$	l_{wave}(y): 1: $\mathbf{e} \leftarrow D_{\mathbf{H}_{\text{sk}}}(\mathbf{y}(\mathbf{S}^{-1})^\top)$ 2: $\mathbf{x} \leftarrow \mathbf{e}\mathbf{P}$ 3: return \mathbf{x}
---	--

Corollary 3. *Let \mathcal{A} be a strong PS-HF-UF-CMA adversary against the HaS-HF-SAS scheme on T_{mayo} in the random oracle model, which makes q_S signing queries, q_H queries to the random oracle H and q_G queries to the random oracle G . Then, there exist a OW adversary \mathcal{B} against T_{mayo} , such that*

$$\begin{aligned} \text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A}) &\leq \frac{\psi q_H}{1 - q_S \mathsf{B}} \cdot \text{Adv}_{\mathsf{T}_{\text{mayo}}}^{\text{OW}}(\mathcal{B}) + \frac{(q_S + q_H)(q_S + q_H + q_G)}{2^{2\lambda}} \\ &\quad + \frac{q_S(q_S + q_H)}{2^\lambda} + \frac{\psi q_H^2}{2|\mathcal{Y}|} + \frac{(\psi q_H)^{\psi+1} |\mathcal{X}|}{(\psi + 1)! \cdot |\mathcal{Y}|^{\psi+1}}, \end{aligned}$$

where $\psi \geq \lceil \text{len}(\mathcal{X}) / \text{len}(\mathcal{Y}) \rceil$, and the running time of \mathcal{B} is about that of \mathcal{A} .

The previous corollary can be applied to the MAYO scheme [9] submitted to the NIST PQC Standardization of Additional Digital Signature. In order to choose appropriate values for ψ , it is necessary to consider the whipped map $\mathcal{P}^*: \mathbb{F}_q^{kn} \rightarrow \mathbb{F}_q^m$, from which $\psi \geq \lceil kn/m \rceil$. If we consider the parameter sets in Table 2, we can choose ψ equal to 13, 7, 14 and 14 for MAYO₁, MAYO₂, MAYO₃ and MAYO₅ respectively, to obtain negligible additive terms in Corollary 3.

B.5 Wave

Wave [17] is a HaS signature scheme based on the family of the generalized $(U, U + V)$ -codes. Let $\mathsf{T}_{\text{wave}} = (\text{TrapGen}_{\text{wave}}, \mathbf{F}_{\text{wave}}, \mathbf{l}_{\text{wave}})$ be the TDF of Wave. The OW security of \mathbf{F}_{wave} is based on the indistinguishability of $(U, U + V)$ -codes from random codes and the Syndrome Decoding (SD) problem. The indistinguishability problem is NP-complete for large finite fields \mathbb{F}_q , while the SD problem is NP-hard for arbitrary finite fields. The trapdoor secret information is a random generalized $(U, U + V)$ -code over \mathbb{F}_q of length n and dimension $k = k_U + k_V$, described by its parity check matrix $\mathbf{H}_{\text{sk}} \in \mathbb{F}_q^{(n-k) \times n}$, an invertible matrix $\mathbf{S} \in \mathbb{F}_q^{(n-k) \times (n-k)}$ and a permutation matrix $\mathbf{P} \in \mathbb{F}_q^{n \times n}$. Using the underlying structure of the $(U, U + V)$ -code, an efficient decoding algorithm $D_{\mathbf{H}_{\text{sk}}}$ is produced. The public function \mathbf{F}_{wave} is obtained from the parity check matrix $\mathbf{H}_{\text{pk}} = \mathbf{S}\mathbf{H}_{\text{sk}}\mathbf{P}$. Let $S_{w,n}$ be the subset of vectors in \mathbb{F}_q^n

with Hamming weight w . The weight w is chosen such that the public function $F_{\text{wave}}: e \in S_{w,n} \mapsto e\mathbf{H}_{\text{pk}}^T \in \mathbb{F}_q^{n-k}$ is a surjection. To find a preimage for $\mathbf{y} \in \mathbb{F}_q^{n-k}$, the signer uses the decoding algorithm $D_{\mathbf{H}_{\text{sk}}}$ on $\mathbf{y}(\mathbf{S}^{-1})^T$ to find $e \in S_{w,n}$, and finally returns $e\mathbf{P}$. The key generation and the preimage computation via l_{wave} are shown in Algorithm 8.

Wave can be described in the HaS *without* retry paradigm, therefore $\mathbf{q}'_{\mathcal{S}} = \mathbf{q}_{\mathcal{S}}$ holds in Theorem 1. In [14], T_{wave} is described in the context of ATPSF, a weaker notion of PSF where the uniformity property on preimages is required to hold only on average. In particular, for any $(F, l) \leftarrow \text{TrapGen}_{\text{wave}}(1^\lambda)$, consider the statistical distance $\varepsilon_{F,l} = \Delta(\text{SampDom}(F), l(U(\mathcal{Y})))$. Then, it holds that $\mathbb{E}_{(F,l)}[\varepsilon_{F,l}] \leq \varepsilon$, where ε is negligible in the security parameter λ . In Theorem 1 we can use this condition and [14, Prop. 1] to bound the distinguishing advantage on PS with ε , obtaining $\text{Adv}_{\mathsf{T}_{\text{wave}}}^{\text{PS}}(\mathcal{D}) \leq \mathbf{q}_{\mathcal{S}}\varepsilon$.

Corollary 4. *Let \mathcal{A} be a strong PS-HF-UF-CMA adversary against the HaS-HF-SAS scheme on T_{wave} in the random oracle model, which makes $\mathbf{q}_{\mathcal{S}}$ signing queries, $\mathbf{q}_{\mathcal{H}}$ queries to the random oracle \mathcal{H} and $\mathbf{q}_{\mathcal{G}}$ queries to the random oracle \mathcal{G} . Then, there exist a OW adversary \mathcal{B} against T_{wave} , such that*

$$\begin{aligned} \text{Adv}_{\text{HaS-HF-SAS}}^{\text{PS-HF-UF-CMA}}(\mathcal{A}) &\leq (\psi\mathbf{q}_{\mathcal{H}}) \cdot \text{Adv}_{\mathsf{T}_{\text{wave}}}^{\text{OW}}(\mathcal{B}) + \mathbf{q}_{\mathcal{S}}\varepsilon + \frac{(\mathbf{q}_{\mathcal{S}} + \mathbf{q}_{\mathcal{H}})(\mathbf{q}_{\mathcal{S}} + \mathbf{q}_{\mathcal{H}} + \mathbf{q}_{\mathcal{G}})}{2^{2\lambda}} \\ &\quad + \frac{\mathbf{q}_{\mathcal{S}}(\mathbf{q}_{\mathcal{S}} + \mathbf{q}_{\mathcal{H}})}{2^\lambda} + \frac{\psi\mathbf{q}_{\mathcal{H}}^2}{2|\mathcal{Y}|} + \frac{(\psi\mathbf{q}_{\mathcal{H}})^{\psi+1}|\mathcal{X}|}{(\psi+1)! \cdot |\mathcal{Y}|^{\psi+1}}, \end{aligned}$$

where $\psi \geq \lceil \text{len}(\mathcal{X})/\text{len}(\mathcal{Y}) \rceil$, and the running time of \mathcal{B} is about that of \mathcal{A} .

In Corollary 4, we can choose $\psi = 3$ to have negligible additive error terms with respect to the parametrization of Table 3.

C PSF-based signatures

In Section 4, we briefly discussed the applicability of the HaS-HF-SAS scheme to Falcon [33] and PSF-based signatures. In general, lattice-based signatures within the GPV framework [24] require the use of PSF. In this section, we discuss how the construction of Section 3 can be modified in the presence of PSF.

Definition 11. *A TDF $\mathsf{T} = (\text{TrapGen}, F, l, \text{SampDom})$ is a Preimage Sampleable Function (PSF) if it satisfies the following properties:*

1. $y \leftarrow F(\text{SampDom}(F))$ is uniformly distributed over \mathcal{Y} .
2. $x \leftarrow l(y)$, with $y \leftarrow_s \mathcal{Y}$, is distributed as $x \leftarrow \text{SampDom}(F)$ conditioned on $F(x) = y$.
3. For any $y \in \mathcal{Y}$, $l(y)$ always returns $x \in \mathcal{X}$ such that $F(x) = y$.

A collision-resistant PSF satisfies the following additional properties:

4. For any $y \in \mathcal{Y}$, the conditional min-entropy of $x \leftarrow \text{SampDom}(F)$ conditioned on $F(x) = y$ is at least $\omega(\log \lambda)$.
5. For an adversary \mathcal{A} , the probability of $\mathcal{A}(F)$ returning two distinct $x, x' \in \mathcal{X}$ such that $F(x) = F(x')$ is negligible in λ .

If we consider a PSF without collision resistance, we can apply Theorem 1 and observe that the distinguishing advantage of the PS adversary is 0. In fact, the PS notion of Definition 5 is a weaker condition for indistinguishability on preimages than property 2 of PSFs. Furthermore, following property 3, the signature associated with the PSF can be described in the HaS without retry paradigm. As a result, we can modify the proof of Theorem 1 by merging Game₂ and Game₃, without the need to introduce the PS adversary. Unfortunately, this would not change the tightness of the reduction.

Conversely, if we consider a collision-resistant PSF, we can further modify Theorem 1 to obtain a tighter reduction from collision resistance (CR) to PS-HF-UF-CMA. In the simulation of H in Algorithm 3, modify Line 21 by taking a random input $x' \leftarrow X$, assigning $y' \leftarrow F^*(x')$ and programming the random oracle G on input h' with $y' \oplus \alpha$. Here the simulation of G is correct since, from property 1 of PSF, y' is uniformly distributed in \mathcal{Y} . After the adversary returns a forged aggregate signature, if none of the bad events happen, the value x_{i^*} , from the forgery, and the value x' , produced by H and stored in the HTree, constitute a collision for F^* . When the reduction is performed from the OW game as in the original proof of Theorem 1, the OW adversary provides its challenge to the PS-HF-UF-CMA adversary in one of the q_H queries to the random oracle H. This results in a multiplicative loss of advantage by a factor q_H . However, when the reduction is performed as above, the CR adversary can prepare responses that will lead to a collision in each query to H involving the target public key. As a result, we get a tight reduction with only negligible losses from additive terms.