

HHGS: Forward-secure Dynamic Group Signatures from Symmetric Primitives^{*}

Xuelian Cao¹, Zheng Yang², Daniel Reijsbergen³, Jianting Ning^{4**}, Junming Ke⁵, Zhiqiang Ma⁶, and Jianying Zhou³

¹ Tsinghua University, Beijing, China

`xl-cao@mail.tsinghua.edu.cn`

² Southwest University, Chongqing, China

³ Singapore University of Technology and Design, Singapore

`{daniel_reijsbergen,jianying_zhou}@sutd.edu.sg`

⁴ Wuhan University, Wuhan, China, and City University of Macau

`jtning88@gmail.com`

⁵ Hangzhou Research Institute of AI and Holographic Technology, Hangzhou, China

`junmingke@zjuqx.com`

⁶ Fujian Normal University, Fuzhou, China

`zhiqiang-ma@hotmail.com`

Abstract. Group signatures allow a group member to sign messages on behalf of the group while preserving the signer’s anonymity, making them invaluable for privacy-sensitive applications. As quantum computing advances, post-quantum security in group signatures becomes essential. Symmetric primitives (SP) offer a promising pathway due to their simplicity, efficiency, and well-understood security foundations. In this paper, we introduce the first *forward-secure dynamic group signature* (FSDGS) framework relying solely on SP. We begin with *hierarchical hypertree group signatures* (HHGS), a basic scheme that securely organizes keys of one-time signatures (OTS) in a hypertree using puncturable pseudorandom functions to enable on-demand key generation and forward security, dynamic enrollment, and which provides resilience against attacks that exploit registration patterns by obfuscating the assignment and usage of keys. We then extend this foundation to HHGS⁺, which orchestrates multiple HHGS instances in a generic way, significantly extending the total signing capacity to $O(2^{60})$, which outperforms HHGS’s closest competitors while keeping signatures below 8 kilobytes. We prove the security of both schemes in the standard model. Our results outline a practical SP-driven pathway toward post-quantum-secure group signatures suitable for resource-constrained client devices.

Keywords: Group signatures · Forward security · Post-quantum security · Symmetric primitives · Standard model.

1 Introduction

Group signatures (GS) serve as an essential component in privacy-preserving applications such as anonymous credentials [10], trusted computing via direct anonymous attestation [15], enhanced privacy identification [25], digital rights management [35], and biometric authentication [16]. GS allow a group member to sign messages on behalf of the group while preserving anonymity, ensuring that a verifier can confirm the signature originates from a legitimate group member without uncovering the signer’s identity. This powerful feature makes GS a fundamental cryptographic primitive for modern security architectures.

As quantum computing threatens the foundations of classical public-key cryptographic systems, the development of post-quantum secure group signatures has become imperative. Various post-quantum constructions have been proposed, leveraging lattice [43,9], code [38], isogeny [9], multivariate [39,42], and symmetric primitives (SP) [2,11,18,20,44]. Each approach has its strengths and limitations. For instance, lattice-based schemes offer resistance against quantum attacks and asymptotic efficiency, but the resulting group signatures still have a signature size of several hundred kilobytes [37]. In contrast, symmetric primitives offer

^{*} This is the full version of the paper that will appear at CT-RSA 2026.

^{**} Jianting Ning is the corresponding author.

simplicity and efficiency and, because the best known quantum attacks (e.g. [27,14]) against them are generic and predictably bounded.

Related Work. There are two primary ways for constructing GS schemes from SP. The first approach lies in the generic GS framework by Bellare et al. [4] based on non-interactive zero-knowledge (NIZK) proofs. The core design of the state-of-the-art NIZK-and-SP based GS [11,20] leverages the “MPC-in-the-Head” technique of Ishai et al. [32] being SP-compatible for achieving post-quantum security. However, the introduction of NIZK in [11,20] results in large signature sizes, with group signatures in [11,20] reaching sizes in the order of megabytes. Moreover, the NIZK-based GS schemes [11,20] are only proven secure in the random oracle model because of the use of Fiat-Shamir-like transformations [22]. However, most commonly-deployed hash functions deviate significantly from the ideal properties assumed in the random oracle model. This gap between the theoretical model and practical implementations can lead to invalid security guarantees [19,21]. Accordingly, although these NIZK-based GS schemes can provide full dynamic security properties [13] (such as full anonymity), they cannot achieve forward security because of the static signing keys of participants. Forward security is a fundamental requirement for group signatures, first formalized by Song [41], ensuring that compromise of a (current) signing key does not permit the forgery of valid group signatures for any earlier time period. NIZK-based GS schemes often require substantial computations on the client device (where the group credential is stored) to produce the zero-knowledge proofs. This high cost in both bandwidth and local processing can be particularly prohibitive for resource-constrained devices (e.g., smart cards or IoT devices whose computational and storage resources are limited).

The alternative approach explored in our paper leverages one-time signatures (OTS) [2,18,44] in combination with symmetric primitive-based authentication mechanisms, such as Merkle trees, to generate group authentication credentials for members’ OTS keys. However, compared to NIZK-based schemes, OTS-based group signature schemes may offer weaker security (or fewer functionalities) and have smaller signing capacities. The G-Merkle scheme [2] supports only static group management, as it relies on a single Merkle tree and a pseudorandom permutation (PRP) scheme with fixed inputs. The enhanced GM^{MT} [44] employs a hypertree structure of multiple sub-trees to enable dynamic group management (**DGMA**), yet it guarantees anonymity only within a single sub-tree [44, §4.2]. That is, GM^{MT} cannot achieve anonymity across the entire group management lifecycle because an adversary can mount the OTS keys of two challenged parties in different sibling sub-trees rather than the same sub-tree (referred to as a *sibling attack*, see also in Section 4). Consequently, any signatures from sub-trees created before a member is joined is necessarily attributed to the older member. Additionally, G-Merkle and GM^{MT} offer limited signing capacities of $O(2^{20})$ and $O(2^{26})$, respectively. Another work following G-Merkle, called dynamic G-Merkle (DGM) [18], allows new users to join the group at any time and distributes members’ OTS keys across different sub-trees. After each sub-tree is generated, its root r_{SMT} is transformed into an independent fallback key F_k via the symmetric encryption SE, i.e., $F_k := \text{SE.Dec}(F_n, r_{\text{SMT}})$ with an encryption key F_n . However, DGM requires the group manager to store all fallback keys and remain online to handle verifier queries concerning the validity of the corresponding fallback keys F_k in signatures. Consequently, it fails to provide members with *self-authenticable* group credentials (**SAGC**), which would allow them to independently validate and generate signatures without relying on the group manager for real-time support. In the latter, DGMT [24] extends DGM to enable SAGC, i.e., allowing non-interactive verification of signatures, by pre-publishing fallback keys. However, this design introduces a centralization threat like G-Merkle: the manager, possessing all signing keys of members, could potentially forge signatures or suffer single-point failures, thus limiting applicability in settings requiring decentralized trust and user-controlled keys (**UCK**). Moreover, to the best of our knowledge, none of the existing OTS-based GS schemes can achieve forward security.

Motivations. To address the shortcomings of existing schemes, we focus on the fundamental research problem of how to design a forward-secure dynamic group signature (FSDGS) scheme using symmetric primitives that achieves the following properties: i) forward security, ii) DGMA, iii) SAGC, iv) UCK, v) small signature size, and vi) provable security in the standard model. We will follow the line of OTS-based approach, since it is friendly to resource-constrained devices requiring few OTS signing operations, rather than expensive NIZK operations. Although group credentials (such as Merkle proofs) in OTS-based schemes [2,18,44] might be large, they can be stored on any honest auxiliary device, rather than the device for storing the secret keys in NIZK-based GS schemes. That is, our design will target such a *resource constrained-auxiliary collaboration* (RCAC) scenario, including smart cards with mobile devices, drones with

ground stations, and IoT devices with smart gateways. Meanwhile, our aim is to improve the signing capacity with both SAGC and UCK, and design techniques to avoid sibling attacks (unlike during dynamic group management).

Our Contributions and Techniques. This work presents the first SP-based FSDGS framework, addressing the motivated research problem. First, we propose HHGS, an innovative FSDGS scheme that leverages a hierarchical OTS-key management structure to enable efficient DGMA. HHGS is the first secure OTS-based GS to achieve DGMA, SAGC and UCK, alongside the resilience against sibling attacks (a security gap in existing schemes). By leveraging puncturable pseudorandom functions (PPRF), HHGS ensures forward security and significantly expands signing capacity to $O(2^{40})$, surpassing previous OTS-based GS schemes with SAGC and UCK by over 2^{14} times under similar join-time efficiency.⁷ Second, we extend this design with HHGS⁺, a scalable enhancement that integrates modular HHGS instances to significantly improve signing capabilities and accommodate larger group sizes. Consequently, the total signing capacity of HHGS⁺ reaches $O(2^{60})$, accommodating most practical applications, while both schemes can maintain signature sizes under 8 kilobytes. Lastly, we rigorously prove the security of these schemes in the standard model, ensuring that they meet anonymity and traceability.

The core of our design lies in HHGS, which organizes OTS keys of a group manager (GM) and group members into an OTS hypertree (as in [7,30]) to support dynamic group management. The group credential of each member’s OTS key is a signature from the hypertree indexed by a unique leaf. The key innovation of HHGS is a tailored obfuscated leaf layer of sub-trees mounted to the hypertree, which consists of two sub-layers (mixed sub-layer and member sub-layer). This design bridges the gap between the determinism of each sub-tree and the dynamic OTS mounting requirements within the sub-tree. All OTS keys in the member sub-layer are from different members joined at any time, each of which is mounted (signed) by its parent, an OTS key, in the mixed sub-layer. The mixed sub-layer encompasses the OTS keys from both members and GM to obfuscate the joining time of these OTS keys. We make use of the fact that the OTS keys from members and GM are generated from the same distribution, so they can be indistinguishable when the owners are uncorrupted. Thus, deterministic sub-trees enable GM to handle dynamic and unpredictable joining requests from members. Our construction addresses sibling attacks by customizing a randomized leaf assignment mechanism. This mechanism employs a pseudorandom function (PRF) to ensure that signing keys are independently distributed across members, preventing adversaries from exploiting shared structural information in signatures to identify members. Since the outputs of existing random functions including PRF are not truly random, potential collisions may occur during the key assignment. To address this, we utilize puncturable pseudorandom functions (PPRFs) in our customized random selection methods to track assigned signing keys. By puncturing previously assigned keys, PPRFs effectively prevent key reuse and resolve collisions. Additionally, PPRFs enable GM to update its secret key after each joining operation, ensuring forward security for its generated signatures. For resource-constrained members, pseudorandom generators serve as a lightweight alternative for generating independent OTS keys and updating secret keys to achieve forward security.

Building on the foundation of HHGS, we introduce HHGS⁺, an extension designed to efficiently accommodate large group sizes and dynamic workloads. HHGS⁺ adopts a modular framework by initializing multiple HHGS instances and organizing them within a unified structure. Specifically, a short-range PRF is used to randomly assign an OTS key to a suitable HHGS instance, utilizing collisions among PRF values. A key challenge in HHGS⁺ lies in randomizing the assignment of OTS keys across HHGS instances while avoiding biases caused by the non-uniform distribution of PRF values, which could lead to earlier full-instance scenarios. To address this, HHGS⁺ integrates redundancy within each HHGS instance, countering potential biases and mitigating adversarial exploitation of uneven assignments. Through formal analysis, we establish a redundancy model that ensures secure and balanced key distribution while maintaining minimal overhead, enabling HHGS⁺ to handle dynamic group operations with high scalability and robustness.

Comparison. Table 1 summarizes a comparison of our proposed schemes with some existing ones according to underlying construction mechanisms. The first compared category involves group signatures (i.e., G-Merkle [2] and GM^{MT} [44]) which are composed of OTS signatures and self-authenticable group membership

⁷ Theoretically, all state-of-the-art OTS-based GS schemes supporting SAGC and UCK can achieve signing capacities approaching $O(2^{60})$. However, due to differences in underlying design principles, their practical performance can vary significantly. We aim to compare them under similar (reported) performance conditions.

credentials (which exclude DGM [18]). G-Merkle, GM^{MT} and DGMT differ from the group credential and the features of supporting dynamic group management (DGMT) and user-centric key (UCK) management. G-Merkle uses a single Merkle tree to create group credentials, so it does not provide DGMA. Additionally, both G-Merkle and DGMT does not support UCK, making them less secure and considerably easier to construct compared to other schemes. In contrast, the group credentials in GM^{MT} are generated using XMSS^{T} [29] equipped with a hypertree. While it supports DGMA, it fails to provide resilience against sibling attacks (SA) [44, §4.2]. SA applies only to DGMA-enabled GS schemes.

Table 1: Comparison among SP-based Group Signature Schemes.

Schemes	Sig. Technique	Credential Structure	UCK Support	Group Management	Standard Model	Forward Security	SA Resilience	Group Size	Each Mem. Sig. Capacity	Sig. Size
G-Merkle [2]	OTS	Merkle tree	×	Static	✓	×	-	2^6	2^{14}	1.9
GM^{MT} [44]	OTS	Hypertree	✓	Dynamic	✓	×	×	2^{16}	2^{10}	3.3
DGMT [24]	OTS	Hypertree	×	Dynamic	✓	×	✓	2^7	$2^{26.76}$	5.344
BEB [11]	NIZK	Merkle tree	✓	Static	×	×	-	2^{40}	unlimit	6650
SPHINX [20]	NIZK	Hypertree	✓	Dynamic	×	×	✓	2^{60}	unlimit	2000
HHGS	OTS	Hypertree	✓	Dynamic	✓	✓	✓	2^{20}	2^{20}	6.41
HHGS ⁺	OTS	Hypertree	✓	Dynamic	✓	✓	✓	2^{30}	2^{30}	7.04

In the second category, a group signature comprises of a SP-based NIZK proof of the knowledge of a group membership credential, such as BEB [11] and SPHINX [20]. The credential in BEB [11], is a signature of GM based on a single Merkle tree, so it only provides static group management. SPHINX [20] is a fully dynamic group signature with a hash-based signature, called F-SPHINCS+, and optimized for efficient NIZK proofs. Unlike previous works, it supports large group sizes (up to 2^{60}).

It is important to note that the comparison between OTS-based group signatures and NIZK-based schemes involves different design goals and security guarantees. In particular, NIZK-based schemes such as BEB and SPHINX can achieve strong anonymity properties, for example full anonymity that holds even against adversaries that obtain all users' secret keys, and they typically provide more compact user-side secret keys. This comes at the cost of larger signature sizes and higher computational overhead for signing and verification. In contrast, our OTS-based constructions are tailored to RCAC scenarios with forward security for resource-constrained devices. They avoid executing computationally intensive NIZK protocols on the client side, at the expense of larger credentials and relatively weaker anonymity guarantees under full key exposure.

Moreover, all compared schemes cannot provide forward security. Unlike NIZK-based GS schemes, our schemes are more suitable for RCAC scenarios involving resource-constrained devices, as they do not require running NIZK operations on devices that store the secret key. Instead, our schemes only need to execute one PRG operation, one OTS key generation, and two OTS signing operations on the client device (e.g., smart card). Although the group credentials are not small, they can be stored on an auxiliary storage device (e.g., mobile phone), and the leakage of group credentials does not compromise traceability. To further reduce storage costs, the client can register OTS keys on-demand, rather than registering all of them at once. Our schemes make a trade-off between group size and signature capacity. The total OTS keys' amount of our scheme is $O(2^{60})$, and it can *adaptively adjust the ratio of members to their OTS keys*. For simplicity, we consider an equal split (as an example), which can meet the requirements of most practical applications.

We also provide a rough comparison of the signature sizes (based on results from the relevant literature) across different schemes, considering the corresponding group sizes. While the signatures in our schemes are slightly larger than those of G-Merkle, GM^{MT} , and DGMT, this increase stems from our simultaneous support for dynamic group management, forward security, user-controlled key management, and sibling attack resilience. However, they are significantly smaller than those of the NIZK-based schemes.

2 Preliminaries

Notation. We let κ be the security parameter and \emptyset be an empty string. The set of integers between 1 and n is represented by $[n] = \{1, \dots, n\} \subset \mathbb{N}$. We use \parallel to denote the string concatenation operation, and $\#$ to represent an operation to calculate the size of an element. Moreover, we denote by $y \leftarrow \mathcal{A}(x)$ the execution

of algorithm \mathcal{A} on input x , resulting in output y . The notation $x \xleftarrow{\$} X$ represents the operation of sampling x uniformly at random from a set X . Let $\text{negl}(\kappa) : \mathbb{N} \rightarrow \mathbb{R}^+$ denote a negligible function, defined such that for every polynomial $P(\kappa)$ there exists a $e_0 \in \mathbb{N}$ s.t. for all $e > e_0$, $\text{negl}(e) \leq 1/P(e)$. We write $\mathcal{A}^\mathcal{O}$ to denote the an algorithm \mathcal{A} with access to oracle \mathcal{O} where \mathcal{O} may represent multiple oracles, i.e., $\mathcal{O} = \{\mathcal{O}_1, \dots, \mathcal{O}_n\}$. For security definitions, we use $\text{Exp}_{\mathcal{A}, \Sigma}^\Psi(1^\kappa, \rho) \Rightarrow 1$ to denote that the experiment Exp , when instantiated with security parameter κ , system parameters ρ , under an adversary \mathcal{A} attacking the security property Ψ of the primitive Σ , returns 1 (indicating success). The parameters ρ , \mathcal{A} , and Ψ may be omitted when clear from context. We define the advantage of \mathcal{A} in this experiment as

$$\text{Adv}_{\mathcal{A}, \Sigma, \lambda}^\Psi := \left| \Pr[\text{Exp}_{\mathcal{A}, \Sigma}^\Psi(1^\kappa, \rho) \Rightarrow 1] - \lambda \right|,$$

where λ may be omitted if $\lambda = 0$.

In the following, we manly review the main cryptographic building blocks of our upcoming constructions. We assume the *parameters* generated by the setup algorithm of a primitive may be implicitly used by its other algorithms.

Digital Signature. A digital signature scheme $\text{SIG} = (\text{Setup}, \text{KGen}, \text{Sign}, \text{Verify})$ involves four algorithms. $\text{Setup}(1^\kappa)$ initializes parameters pm_{SIG} , defining the randomness space $\mathcal{RS}_{\text{SIG}}$, secret key $\mathcal{SK}_{\text{SIG}}$, public key space $\mathcal{PK}_{\text{SIG}}$, and signature space \mathcal{S}_{SIG} . $\text{KGen}(rs)$ uses randomness $rs \in \mathcal{RS}_{\text{SIG}}$ to produce a secret key $sk \in \mathcal{SK}_{\text{SIG}}$ and its corresponding public key $pk \in \mathcal{PK}_{\text{SIG}}$. $\text{Sign}(sk, m)$ generates a signature $\sigma \in \mathcal{S}_{\text{SIG}}$ for a message $m \in \mathcal{M}_{\text{SIG}}$ using the secret key sk . $\text{Verify}(pk, m, \sigma)$ checks if σ is a valid signature on m under pk , outputting 1 if valid and 0 otherwise. One-time signature (OTS) schemes, designed for securely signing a single message with a unique key pair, ensure authenticity and resistance to forgery, with security analyzed under adaptive chosen message attacks.

We say that a SIG scheme \mathbf{S} is correct if for all $m \in \mathcal{M}_{\text{SIG}}$ it holds that

$$\Pr \left[\mathbf{S}.\text{Verify}(pk, m, \sigma) = 1 \mid \begin{array}{l} pm_{\text{SIG}} \leftarrow \mathbf{S}.\text{Setup}(1^\kappa); \quad rs \xleftarrow{\$} \mathcal{RS}_{\text{SIG}}; \quad (sk, pk) \leftarrow \mathbf{S}.\text{KGen}(rs) \\ \sigma \leftarrow \mathbf{S}.\text{Sign}(sk, m) \end{array} \right] = 1.$$

To capture existential unforgeability under adaptive chosen message attacks (EUF-CMA), we define a security experiment $\text{Exp}_{\mathcal{A}, \mathbf{S}}^{\text{EUF-CMA}}(\kappa, q)$ as below:

$\text{Exp}_{\mathcal{A}, \mathbf{S}}^{\text{EUF-CMA}}(\kappa, q) :$ $pm_{\text{SIG}} \leftarrow \mathbf{S}.\text{Setup}(1^\kappa); \quad rs \xleftarrow{\$} \mathcal{RS}_{\text{SIG}}; \quad (sk, pk) \leftarrow \mathbf{S}.\text{KGen}(rs)$ $ct := 0; \quad L_S := \emptyset$ $(m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{HMSign}(\cdot)}(\kappa, pm_{\text{SIG}}, pk);$ $\text{OUTPUT } ((m^*, \sigma^*) \notin L_S) \wedge (\mathbf{S}.\text{Verify}(pk, m^*, \sigma^*) = 1)$	$\text{HMSign}(m) :$ $\text{IF } ct \geq q, \text{ OUTPUT } \perp$ $ct := ct + 1; \quad \sigma \leftarrow \mathbf{S}.\text{Sign}(sk, m)$ $\text{APPEND } (m, \sigma) \rightarrow L_S$ $\text{OUTPUT } \sigma$
--	---

Definition 1 (EUF-CMA for SIG). We say that \mathbf{S} is an EUF-CMA secure SIG scheme if for any PPT \mathcal{A} , the advantage $\text{Adv}_{\mathcal{A}, \mathbf{S}}^{\text{EUF-CMA}}(\kappa, q)$ of \mathcal{A} in $\text{Exp}_{\mathcal{A}, \mathbf{S}}^{\text{EUF-CMA}}(\kappa, q)$ is negligible. An EUF-CMA secure SIG scheme \mathbf{S} with $q = 1$ is referred to as a one-time signature (OTS) scheme.

Merkle Tree. A Merkle tree scheme consists of four algorithms, $\text{MT} = (\text{Setup}, \text{Build}, \text{GetPrf}, \text{Verify})$. $\text{Setup}(1^\kappa)$ takes as input 1^κ and outputs the parameter pm_{MT} defining leaf space \mathcal{L}_{M} and proof space \mathcal{PF}_{M} . $\text{Build}(\{\text{lf}_i\}_{i \in [N]})$ constructs a Merkle tree instance Tr based on these leaves in \mathcal{L}_{M} , and outputs the initial state st_{BDS} for the tree traversal algorithm described in [17] (commonly referred to as the BDS algorithm). Moreover, we let $\text{st}_{\text{BDS}}^{\text{lf}_i} \subseteq \text{st}_{\text{BDS}}$ denote the BDS state for a specific leaf lf_i . $\text{GetPrf}(\text{st}_{\text{BDS}}, \text{lf}_i)$ takes as input the state st_{BDS} and a leaf lf_i , and out outputs a Merkle proof $\text{pf}_{\text{lf}_i} \in \mathcal{PF}_{\text{M}}$ that attests to the inclusion of lf_i in the tree, along with the updated BDS state st_{BDS} . $\text{Verify}(\text{Tr}.\text{Rt}, \text{lf}_i, \text{pf}_{\text{lf}_i})$ takes as input the root $\text{Tr}.\text{Rt}$, a leaf node lf_i , and the corresponding proof pf_{lf_i} , and returns 1 if the proof validates lf_i 's inclusion in $\text{Tr}.\text{Rt}$, and 0 otherwise. The Merkle proofs of a secure MT scheme should satisfy unforgeability.

We say that a MT scheme \mathbf{M} is correct if for a polynomial number N defined by κ , any $\{\text{lf}_i\}_{i \in [N]}$ (s.t. $\text{lf}_i \in \mathcal{L}_{\text{M}}$) and $\text{lf}_j \in \{\text{lf}_i\}_{i \in [N]}$, it holds that

$$\Pr \left[\mathbf{M}.\text{Verify}(\text{Tr}.\text{Rt}, \text{lf}_j, \text{pf}_{\text{lf}_j}) = 1 \mid \begin{array}{l} pm_{\text{MT}} \leftarrow \mathbf{M}.\text{Setup}(1^\kappa); \quad (\text{Tr}, \text{st}_{\text{BDS}}) \leftarrow \mathbf{M}.\text{Build}(\{\text{lf}_i\}_{i \in [N]}) \\ \text{pf}_{\text{lf}_j} \leftarrow \mathbf{M}.\text{GetPrf}(\text{st}_{\text{BDS}}, \text{lf}_j) \end{array} \right] = 1.$$

We define a security experiment $\text{Exp}_{\mathcal{A}, \mathbf{M}}^{\text{UF}}(\kappa)$ regarding unforgeability (UF) for a MT scheme \mathbf{M} as below:

$$\begin{array}{l}
\text{Exp}_{\mathcal{A},\mathcal{M}}^{\text{UF}}(\kappa) : \\
pms_{\text{MT}} \leftarrow \text{M.Setup}(1^\kappa) ; (st, \{\text{lf}_i\}_{i \in [N]}) \leftarrow \mathcal{A}_1(\kappa, pms_{\text{MT}}) \\
(\text{Tr}, \text{st}_{\text{BDS}}) \leftarrow \text{M.Build}(\{\text{lf}_i\}_{i \in [N]}) ; (\text{lf}^*, \text{pf}^*) \leftarrow \mathcal{A}_2(st, \text{Tr}) \\
\text{OUTPUT } (\text{M.Verify}(\text{Tr.Rt}, \text{lf}^*, \text{pf}^*) = 1) \wedge (\text{lf}^* \notin \{\text{lf}_i\}_{i \in [N]})
\end{array}$$

Definition 2 (Unforgeability of MT). We say that M is a secure MT scheme if the advantage $\text{Adv}_{\mathcal{A},\mathcal{M}}^{\text{UF}}(\kappa)$ of any PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in $\text{Exp}_{\mathcal{A},\mathcal{M}}^{\text{UF}}(\kappa)$ is negligible.

Authenticated Encryption. An authenticated encryption (AE) scheme consists of four algorithms $\text{AE} = (\text{Setup}, \text{KGen}, \text{Enc}, \text{Dec})$. $\text{Setup}(1^\kappa)$ initializes the scheme with security parameter κ , outputting public parameters pms_{AE} that define the key space \mathcal{K}_{AE} , plaintext space \mathcal{M}_{AE} , and ciphertext space \mathcal{C}_{AE} . $\text{KGen}(rk)$ generates a symmetric key $k \in \mathcal{K}_{\text{AE}}$ using randomness $rk \in \mathcal{RK}_{\text{AE}}$. $\text{Enc}(k, m)$ encrypts a message $m \in \mathcal{M}_{\text{AE}}$ with key k , producing ciphertext $c \in \mathcal{C}_{\text{AE}}$. $\text{Dec}(k, c)$ decrypts $(k, c) \in \mathcal{K}_{\text{AE}} \times \mathcal{C}_{\text{AE}}$, outputting plaintext m or failure symbol \perp . For security of AE, we consider the security properties regarding indistinguishability under adaptive chosen-ciphertext attacks, and the integrity of ciphertext.

In the following, we define two experiments for security properties regarding indistinguishability under adaptive chosen-ciphertext attacks (IND-CCA) and integrity of ciphertext (CT-INT), respectively.

An AE scheme E is said to be correct if for any $m \in \mathcal{M}_{\text{AE}}$ it holds that

$$\Pr \left[\text{E.Dec}(k, c) = m \mid \begin{array}{l} pms_{\text{AE}} \xleftarrow{\$} \text{E.Setup}(1^\kappa); rk \xleftarrow{\$} \mathcal{RK}_{\text{AE}}; k \leftarrow \text{E.KGen}(rk) \\ re \xleftarrow{\$} \mathcal{R}_{\text{AE}}; c \leftarrow \text{E.Enc}(k, m; re) \end{array} \right] = 1.$$

In the following, we first define the experiment for the security property regarding indistinguishability under adaptive chosen-ciphertext attacks (IND-CCA).

$$\begin{array}{l}
\text{Exp}_{\mathcal{A},\text{E}}^{\text{IND-CCA}}(\kappa) : \\
pms_{\text{AE}} \xleftarrow{\$} \text{E.Setup}(1^\kappa); b \xleftarrow{\$} \{0, 1\}; rk \xleftarrow{\$} \mathcal{RK}_{\text{AE}}; k \leftarrow \text{E.KGen}(rk) \\
(m_0, m_1, st) \leftarrow \mathcal{A}_1^{\text{O}_{\text{Dec}}(\cdot)}(\kappa, pms_{\text{AE}}); re^* \xleftarrow{\$} \mathcal{R}_{\text{AE}}; c^* \leftarrow \text{E.Enc}(k, m_b; re^*) \\
b' \leftarrow \mathcal{A}_2^{\text{O}_{\text{Dec}}(\cdot)}(c^*, m_0, m_1, st); \\
\text{OUTPUT } b = b'
\end{array}
\quad \left| \quad \begin{array}{l}
\mathcal{O}_{\text{Dec}}(c) : \\
\text{If } c = c^*, \text{ OUTPUT } \perp \\
m \leftarrow \text{E.Dec}(k, c) \\
\text{OUTPUT } m
\end{array}
\right.$$

Definition 3 (IND-CCA for AE). We say that E is an IND-CCA-secure AE scheme if the advantage $\text{Adv}_{\mathcal{A},\text{E},\frac{1}{2}}^{\text{IND-CCA}}(\kappa)$ of all PPT adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in $\text{Exp}_{\mathcal{A},\text{E}}^{\text{IND-CCA}}(\kappa)$ is negligible.

Also, we need the AE to provide integrity CT-INT of ciphertext against adaptive-chosen message attacks, which is formalized by the following experiment.

$$\begin{array}{l}
\text{Exp}_{\mathcal{A},\text{E}}^{\text{CT-INT}}(\kappa) : \\
pms_{\text{AE}} \xleftarrow{\$} \text{E.Setup}(1^\kappa); L_{\text{AE}} := \emptyset; rk \xleftarrow{\$} \mathcal{RK}_{\text{AE}}; k \xleftarrow{\$} \text{E.KGen}(rk) \\
c^* \leftarrow \mathcal{A}^{\text{O}_{\text{Enc}}(\cdot)}(\kappa, pms_{\text{AE}}); m^* \leftarrow \text{Dec}(k, c^*); \\
\text{OUTPUT } c^* \notin L_{\text{AE}} \wedge m^* \neq \perp
\end{array}
\quad \left| \quad \begin{array}{l}
\mathcal{O}_{\text{Enc}}(m) : \\
re \xleftarrow{\$} \mathcal{R}_{\text{AE}} \\
c \leftarrow \text{E.Enc}(k, m; re) \\
\text{APPEND } c \rightarrow L_{\text{AE}} \\
\text{Return } c
\end{array}
\right.$$

Definition 4 (CT-INT for AE). We say that E is a CT-INT-secure AE if the advantage $\text{Adv}_{\mathcal{A},\text{E}}^{\text{CT-INT}}(\kappa)$ of all PPT adversaries \mathcal{A} in $\text{Exp}_{\mathcal{A},\text{E}}^{\text{CT-INT}}(\kappa)$ is negligible.

Pseudorandom Generators. A pseudorandom generator consists of three algorithms $\text{PRG} = (\text{Setup}, \text{SGen}, \text{Eval})$. $\text{Setup}(1^\kappa)$ initializes the scheme with security parameter κ , producing parameters pms_{PRG} that define the input space \mathcal{S}_{PRG} and range space \mathcal{R}_{PRG} . $\text{SGen}(pms_{\text{PRG}})$ uses these parameters to output a random seed $s \xleftarrow{\$} \mathcal{S}_{\text{PRG}}$, while the randomness generation algorithm $\text{Eval}(s)$ takes a seed $s \in \mathcal{S}_{\text{PRG}}$ and outputs a pseudorandom value $r \in \mathcal{R}_{\text{PRG}}$. The output of PRG must be indistinguishable from a truly random value. For a specific PRG scheme G , we may use the shorthand notation $\text{G}(s)$ to represent $\text{G.Eval}(s)$. We define the following security experiment $\text{Exp}_{\mathcal{A},\text{G}}^{\text{IND}}(\kappa)$ to formalize the security of a PRG scheme G based on indistinguishability (IND):

$\text{Exp}_{\mathcal{A},\mathcal{G}}^{\text{IND}}(\kappa) :$
 $b \xleftarrow{\$} \{0,1\}; pms_{\text{PRG}} \leftarrow \mathcal{G}.\text{Setup}(1^\kappa); s \xleftarrow{\$} \mathcal{G}.\text{SGen}(pms_{\text{PRG}}); r_0 \xleftarrow{\$} \mathcal{R}_{\text{PRG}}; r_1 \leftarrow \mathcal{G}.\text{Eval}(s)$
 $b^* \leftarrow \mathcal{A}(\kappa, pms_{\text{PRG}}, r_b);$
 OUTPUT $b = b^*$

Definition 5 (IND for PRG). We say that \mathcal{G} is a secure PRG scheme if for any probabilistic polynomial time (PPT) adversary \mathcal{A} , it holds that the advantage $\text{Adv}_{\mathcal{A},\mathcal{G}}^{\text{IND}}(\kappa)$ of \mathcal{A} in $\text{Exp}_{\mathcal{A},\mathcal{G}}^{\text{IND}}(\kappa)$ is negligible.

(Puncturable) Pseudo-random Functions. We define a pseudorandom function (PRF) family with three algorithms $\text{PRF} = (\text{Setup}, \text{KGen}, \text{Eval})$. $\text{Setup}(1^\kappa)$ initializes the scheme with security parameter κ , producing parameters pms_{PRF} that define the key space \mathcal{K}_{PRF} , message space \mathcal{M}_{PRF} , and range space \mathcal{R}_{PRF} . $\text{KGen}(pms_{\text{PRF}})$ generates a random secret key $k \xleftarrow{\$} \mathcal{K}_{\text{PRF}}$. $\text{Eval}(k, x)$ computes the evaluation result $r \in \mathcal{R}_{\text{PRF}}$ for a message $x \in \mathcal{M}_{\text{PRF}}$ using the secret key k .

A puncturable pseudorandom function (PPRF) family is defined by four algorithms, $\text{PPRF} = (\text{Setup}, \text{KGen}, \text{Eval}, \text{Punc})$, with the same spaces as PRF for simplicity. The syntax of the first two algorithms mirrors that of PRF. The puncturing algorithm $\text{Punc}(k, x)$ takes as input a key $k \in \mathcal{K}_{\text{PRF}}$ and a message $x \in \mathcal{M}_{\text{PRF}}$, and outputs a punctured key in \mathcal{K}_{PRF} . For a specific (P)PRF function \mathcal{F} , we may use the shorthand notation $\mathcal{F}(k, x)$ to represent $\mathcal{F}.\text{Eval}(k, x)$.

We say that a PPRF scheme \mathcal{F} is correct if for every subset $\mathcal{S} = \{x_1, \dots, x_q\} \subseteq \mathcal{M}_{\text{PRF}}$ and any $x \in \mathcal{M}_{\text{PRF}} \setminus \mathcal{S}$ it holds that

$$\Pr \left[\mathcal{F}.\text{Eval}(k_0, x) = \mathcal{F}.\text{Eval}(k_t, x) \mid \begin{array}{l} pms_{\text{PRF}} \leftarrow \mathcal{F}.\text{Setup}(1^\kappa); k \xleftarrow{\$} \mathcal{F}.\text{KGen}(pms_{\text{PRF}}) \\ \text{for } i \in [q] : k_i \leftarrow \mathcal{F}.\text{Punc}(k_{i-1}, x_i) \end{array} \right] = 1.$$

We define a security experiment $\text{Exp}_{\mathcal{A},\mathcal{F}}^{\text{IND-CMA}}(\kappa, q_f)$ for (P)PRF below. The boxed lines are exclusively used in the security experiment for PPRF.

$\text{Exp}_{\mathcal{A},\mathcal{F}}^{\text{IND-CMA}}(\kappa, q_f) :$ $b \xleftarrow{\$} \{0,1\}; ct := 0; pms_{\text{PRF}} \leftarrow \mathcal{F}.\text{Setup}(1^\kappa); k \xleftarrow{\$} \mathcal{F}.\text{KGen}(pms_{\text{PRF}})$ $(x^*, st) \leftarrow \mathcal{A}_1^{\text{OPRF}(\cdot)}(\kappa, pms_{\text{PRF}}); r_0 \xleftarrow{\$} \mathcal{R}_{\text{PRF}}; r_1 \leftarrow \mathcal{F}.\text{Eval}(k, x^*)$ IF $x^* \in \mathcal{L}_{\text{PRF}}$, OUTPUT \perp <div style="border: 1px solid black; padding: 2px; display: inline-block;"> $k \leftarrow \mathcal{F}.\text{Punc}(k, x^*)$ </div> ; $b' \leftarrow \mathcal{A}_2^{\text{OPRF}(\cdot)}(st, r_b, \boxed{k})$ OUTPUT $(b' = b)$	$\mathcal{O}_{\text{PRF}}(x) :$ IF $ct \geq q_f$, OUTPUT \perp APPEND $x \rightarrow \mathcal{L}_{\text{PRF}}$ $r \leftarrow \mathcal{F}.\text{Eval}(k, x)$ <div style="border: 1px solid black; padding: 2px; display: inline-block;"> $k \leftarrow \mathcal{F}.\text{Punc}(k, x)$ </div> OUTPUT r
--	---

Definition 6 (IND-CMA for (P)PRF). We say that \mathcal{F} is a secure (P)PRF scheme if, in the experiment $\text{Exp}_{\mathcal{A},\mathcal{F}}^{\text{IND-CMA}}(\kappa, q_f)$, the advantage $\text{Adv}_{\mathcal{A},\mathcal{F}}^{\text{IND-CMA}}(\kappa, q_f)$ of any PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is negligible.

In addition, we define the statistical distance (SD) of \mathcal{F} (with a random key $k \xleftarrow{\$} \mathcal{K}_{\text{PRF}}$) to a uniform random function $\mathcal{U} : \mathcal{M}_{\text{PRF}} \rightarrow \mathcal{R}_{\text{PRF}}$, as $\Delta(\mathcal{F}, \mathcal{U}) = \frac{1}{2} \sum_{x \in \mathcal{M}_{\text{PRF}}} |\Pr[\mathcal{F}(k, x) = y] - \Pr[\mathcal{U}(x) = y]|$.

Collision-resistant Hash Functions. We define a collision-resistant (CR) hash function $\text{CRH} = (\text{Setup}, \text{Eval})$ as a keyed hash function with two algorithms. $\text{Setup}(1^\kappa)$ initializes the scheme with security parameter κ , producing parameters pms_{CRH} and a random key $hk \xleftarrow{\$} \mathcal{K}_{\text{CRH}}$, where pms_{CRH} defines the key space \mathcal{K}_{CRH} , message space \mathcal{M}_{CRH} , and hash value space \mathcal{Y}_{CRH} . $\text{Eval}(hk, m)$ computes a hash value $y \in \mathcal{Y}_{\text{CRH}}$ for a message $m \in \mathcal{M}_{\text{CRH}}$ using key hk . For a given CRH scheme \mathcal{H} , we use $\mathcal{H}(m)$ as shorthand for $\mathcal{H}.\text{Eval}(hk, m)$ when the hash key hk is clear from the context.

We define the security experiment $\text{Exp}_{\mathcal{A},\mathcal{H}}^{\text{CR}}(\kappa)$ for a CRH scheme \mathcal{H} as follows.

$\text{Exp}_{\mathcal{A},\mathcal{H}}^{\text{CR}}(\kappa) :$
 $(hk, pms_{\text{CRH}}) \leftarrow \mathcal{H}.\text{Setup}(1^\kappa); (m_0, m_1) \leftarrow \mathcal{A}(\kappa, pms_{\text{CRH}}, hk)$
 OUTPUT $(m_0 \neq m_1) \wedge (\mathcal{H}.\text{Eval}(hk, m_0) = \mathcal{H}.\text{Eval}(hk, m_1))$

Definition 7. We say \mathcal{H} is secure if no PPT adversary has non-negligible advantage $\text{Adv}_{\mathcal{A},\mathcal{H}}^{\text{CR}}(\kappa)$ in $\text{Exp}_{\mathcal{A},\mathcal{H}}^{\text{CR}}(\kappa)$.

3 Security Model of FSDGS

Syntax. A FSDGS scheme involves a group manager (GM), verifiers, and group members. The GM is a trusted third party responsible for system initialization and group member management. Each group member has a unique identity ID and can anonymously create a group signature. Verifier, who might be an outsider or a member of the group, can verify the validity of a group signature. We define the syntax of FSDGS via the following algorithms $\Sigma = (\text{GMInit}, \text{MInit}, \text{MRegGen}, \text{Join}, \text{Sign}, \text{Verify}, \text{Open}, \text{Revoke})$:

- $(\text{Pms}, \text{gpk}, \text{sk}_{\text{GM}}, \text{st}_{\text{GM}}, \text{RL}) \leftarrow \text{GMInit}(1^\kappa, \text{SP})$: The GM initialization algorithm takes as input the security parameter 1^κ and a setup parameter SP required by a specific scheme. This algorithm outputs a system parameter Pms, a group public and private key pair $(\text{gpk}, \text{sk}_{\text{GM}})$, the group management state st_{GM} of the group manager GM, and a revocation list RL.
- $(\text{sk}_{\text{ID}}, \text{st}_{\text{ID}}) \leftarrow \text{MInit}(\text{ID})$: A group member ID runs the member initialization algorithm to generate its secret key sk_{ID} and the initial state st_{ID} .
- $\text{RF}_{\text{ID}} \leftarrow \text{MRegGen}(\text{sk}_{\text{ID}})$: A group member ID uses its secret key sk_{ID} to generate the registration file $\text{RF}_{\text{ID}} \in \mathcal{RF}$ from the corresponding space \mathcal{RF} .
- $(\text{Vt}, \text{RR}_{\text{ID}}) \leftarrow \text{Join}(\text{sk}_{\text{GM}}, \text{st}_{\text{GM}}, \text{ID}, \text{RF}_{\text{ID}})$: The join protocol is an interactive procedure between the group manager and a member ID seeking to join the group with its registration file RF_{ID} . The group manager GM utilizes the group secret key sk_{GM} and its local state st_{GM} . Upon completion of the protocol, GM updates the group secret key sk_{GM} and its management state st_{GM} , respectively. ID will update its state st_{ID} using the registration result RR_{ID} obtained from GM. Furthermore, the transcript of the protocol's execution is recorded in the variable Vt which may be \perp indicating any join failure. A group member ID can register different registration files many times.
- $(\text{sk}'_{\text{ID}}, \text{st}'_{\text{ID}}, \sigma) \leftarrow \text{Sign}(\text{sk}_{\text{ID}}, \text{st}_{\text{ID}}, m)$: The signing algorithm takes as input the secret key sk_{ID} of ID, along with its state st_{ID} , and a message $m \in \mathcal{M}$. The algorithm generates a group signature σ on the message m and also outputs an updated secret key sk'_{ID} and state st'_{ID} , for ID.
- $\{0, 1\} \leftarrow \text{Verify}(\text{gpk}, m, \sigma, \text{RL})$: The signature verification algorithm takes as input the group public key gpk, a message m , a group signature σ on m , and a public revocation list RL. It outputs 1 if the signature is valid, and 0 otherwise.
- $\text{ID} \leftarrow \text{Open}(\text{gpk}, \text{sk}_{\text{GM}}, m, \sigma)$: The opening algorithm takes as input the group public key gpk, the group secret key sk_{GM} , a message m , and a group signature σ for m to return an identity ID or the symbol \perp to indicate failure.
- $\text{RL}' \leftarrow \text{Revoke}(\text{gpk}, \text{sk}_{\text{GM}}, \text{RL}, \text{ID}, \text{RSS})$: The revocation algorithm takes as input the group public key gpk, the group secret key sk_{GM} , the revocation list RL, the revoking identity ID, and a set of group signatures RSS to be revoked. This algorithm updates the revocation list RL to RL' .

Correctness. A correct FSDGS scheme should ensure that all honest members' unrevoked signatures on any messages should pass the verification algorithm, i.e., $\Pr[\text{Exp}_{\mathcal{A}, \Sigma}^{\text{Correct}}(\kappa, q_s, \text{SP}) \Rightarrow 1] = 1$, where is $\text{Exp}_{\mathcal{A}, \Sigma}^{\text{Correct}}(\kappa, q_s, \text{SP})$ the correctness (Correct) experiment defined in Figure 1 for $\text{exp} = \text{Correct}$.

Security Definition. The security properties that we consider in this work, are adapted from [4,5,13,20,36], including two critical properties: forward-secure traceability (Trace) and anonymity (Anony) with forward privacy, in a dynamic group management setting. For the security definition, we formalize these properties through the security experiments defined in Figure 1. These experiments are indexed by the variable $\text{exp} \in \{\text{Trace}, \text{Anony}\}$, respectively. Moreover, we let $\text{GetCT}()$ be a generic timestamp function that outputs the current system time, e.g., by counting the steps executed by the Turing Machine.

Informally speaking, forward-Secure traceability requires that even an adversary, capable of corrupting the tracing manager and potentially some or all group members (via **Corrupt** queries), cannot produce a valid signature under two conditions (**Trace-Cons**): i) The signature cannot be opened to a non-corrupted group member, ensuring that honest users remain protected; ii) The signature cannot be traced to a corrupted member if it was created in a time period before the adversary accessed that the secret key of a member or the group manager. These properties strengthen the standard traceability requirement, as established in [4,5], to include forward-security, ensuring that prior signatures remain secure even after either a member or group manager is compromised. Anonymity with forward privacy ensures that a PPT adversary cannot distinguish (**Anony-Cons**) which of two honest group members signed a challenged message, provided that GM remains

uncorrupted, the signatures stay unopened, and the signers were uncorrupted when they generated their signatures. This holds even when the adversary selects the members (adding honest and malicious ones via AddHM and AddMM queries, respectively), the signing message of honest members (using HMSign queries), and the time period. Meanwhile, an adversary may exploit the dynamics of the group management (e.g., registering malicious members, revoking members, opening identities of signatures) trying to gain extra advantage for breaking these two properties.

Definition 8 (FSDGS Security). *We say that a correct FSDGS scheme Σ is secure if the advantages, $\text{Adv}_{\mathcal{A}, \Sigma}^{\text{Trace}}(\kappa, q_s, \text{SP})$ and $\text{Adv}_{\mathcal{A}, \Sigma, \frac{1}{2}}^{\text{Anony}}(\kappa, q_s, \text{SP})$, of any PPT adversaries are negligible in the corresponding security experiments.*

$\text{Exp}_{\mathcal{A}, \Sigma}^{\text{exp}}(\kappa, q_s, \text{SP}) :$ $b \xleftarrow{\$} \{0, 1\}; T_p := \infty \quad \triangleright \text{Anonymity test bit and time}$ $t_s := 0 \quad \triangleright \text{Number of asked signing queries}$ $T_c^{\text{GM}} := \infty \quad \triangleright \text{Corrupt time of GM}$ $L_C := L_{\text{HM}} := \emptyset \quad \triangleright \text{Lists of corrupted and honest IDs}$ $L_{\text{Op}} := L_S := \emptyset \quad \triangleright \text{Lists of opened and queried signatures}$ $L_R := \emptyset \quad \triangleright \text{Lists of revoked IDs and signatures}$ $(\text{Pms}, \text{gpk}, \text{sk}_{\text{GM}}, \text{st}_{\text{GM}}, \text{RL}) \leftarrow \text{GMInit}(1^\kappa, \text{SP})$ $(b^*, m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{FSDGS}}(\cdot), \mathcal{O}_{\text{PriTest}}(\cdot, \cdot, \cdot)}(\text{Pms}, \text{gpk})$ $T^* := \text{GetCT}()$ $vr^* \leftarrow \text{Verify}(\text{gpk}, m^*, \sigma^*, \text{RL})$ $\text{ID}^* \leftarrow \Sigma.\text{Open}(\text{gpk}, \text{sk}_{\text{GM}}, m^*, \sigma^*)$ IF exp = Correct IF $(\text{ID}^* \notin L_{\text{HM}}) \vee (\text{ID}^* \in L_R) \vee (\sigma^* \in L_R)$, OUTPUT 1 OUTPUT $((vr^* = 1) \wedge ((\text{ID}^*, m^*, \sigma^*) \in L_S))$ IF exp = Trace \triangleright Trace-Cons IF $(\text{ID}^* \notin L_{\text{HM}}) \vee ((\text{ID}^* \in L_{\text{HM}}) \wedge (T_c^{\text{ID}^*} > T^*))$, $idc := 1$ IF $T_c^{\text{GM}} > T^*$, $gmc := 1$ OUTPUT $(vr^* = 1) \wedge idc \wedge gmc$ IF exp = Anony \triangleright Anony-Cons IF $((\text{ID}_0^*, \text{ID}_1^*) \notin L_{\text{HM}}) \vee (T_p \geq T_c^{\text{ID}_0^*}) \vee (T_p \geq T_c^{\text{ID}_1^*})$ IF $(\text{GM} \in L_C) \vee (\hat{\sigma}_0 \in L_{\text{Op}}) \vee (\hat{\sigma}_1 \in L_{\text{Op}})$, OUTPUT \perp OUTPUT $b^* = b$ <hr/> $\mathcal{O}_{\text{PriTest}}(\text{ID}_0^*, \text{ID}_1^*, m) :$ IF $(\text{exp} \neq \text{Anony}) \wedge ((\text{ID}_0^*, \text{ID}_1^*) \notin L_{\text{HM}})$, OUTPUT \perp $T_p := \text{GetCT}()$ $(\text{sk}_{\text{ID}_0^*}', \text{st}_{\text{ID}_0^*}', \hat{\sigma}_0) \leftarrow \text{Sign}(\text{sk}_{\text{ID}_0^*}, \text{st}_{\text{ID}_0^*}, m)$ $(\text{sk}_{\text{ID}_1^*}', \text{st}_{\text{ID}_1^*}', \hat{\sigma}_1) \leftarrow \text{Sign}(\text{sk}_{\text{ID}_1^*}, \text{st}_{\text{ID}_1^*}, m)$ OUTPUT $\hat{\sigma}_b$	$\mathcal{O}_{\text{FSDGS}}(\text{Query}) :$ $T_\delta := \text{GetCT}()$ IF Query = AddHM(ID): IF $\text{ID} \in L_C \cup \perp$, OUTPUT \perp IF $\text{ID} \notin L_{\text{HM}}$, $(\text{sk}_{\text{ID}}, \text{st}_{\text{ID}}) \leftarrow \text{MInit}(\text{ID})$ $T_c^{\text{ID}} := \infty$; $\text{RF}_{\text{ID}} \leftarrow \text{MRegGen}(\text{sk}_{\text{ID}})$ $(\text{Vt}, \text{RR}_{\text{ID}}) \leftarrow \text{Join}(\text{sk}_{\text{GM}}, \text{st}_{\text{GM}}, \text{ID}, \text{RF}_{\text{ID}})$ OUTPUT $\text{Vt} \neq \perp$ IF Query = AddMM(ID, RF_{ID}) IF $\text{ID} \in L_{\text{HM}} \cup \perp$, OUTPUT \perp $T_c^{\text{ID}} := T_\delta$; $L_C := L_C \cup (\text{ID}, T_c^{\text{ID}}) \rightarrow L_C$ OUTPUT $\text{Join}(\text{sk}_{\text{GM}}, \text{st}_{\text{GM}}, \text{ID}, \text{RF}_{\text{ID}})$ IF Query = HMSign(ID, m) IF $(\text{ID} \notin L_{\text{HM}}) \vee (t_s \geq q_s)$, OUTPUT \perp $(\text{sk}_{\text{ID}}', \text{st}_{\text{ID}}', \sigma) \leftarrow \Sigma.\text{Sign}(\text{sk}_{\text{ID}}, \text{st}_{\text{ID}}, m)$ $L_S := L_S \cup (\text{ID}, m, \sigma)$; $t_s := t_s + 1$ OUTPUT σ IF Query = Corrupt(P) IF $(P \in L_C) \vee (P \notin L_{\text{HM}} \cup \text{GM})$, OUTPUT \perp $T_c^P := T_\delta$; $L_C := L_C \cup (P, T_c^P)$ OUTPUT sk_P IF Query = OpenID(m, σ) $L_{\text{Op}} := L_{\text{Op}} \cup \sigma$ OUTPUT $\Sigma.\text{Open}(\text{gpk}, \text{sk}_{\text{GM}}, m, \sigma)$ IF Query = Revoke(ID, RSS) $L_R := L_R \cup (\text{ID}, \text{RSS})$ OUTPUT $\text{Revoke}(\text{gpk}, \text{sk}_{\text{GM}}, \text{RL}, \text{ID}, \text{RSS})$
--	---

Fig. 1: Security Experiments for Forward Secure Dynamic Group Signature.

4 Our FSDGS Constructions

In this section, we introduce two novel FSDGS schemes based on OTS. The first scheme is built upon an OTS hypertree, while the second securely leverages the first as a building block to further extend the signing capability.

4.1 An Efficient Solution HHGS with Hypertree

Construction Overview. The goals of HHGS include: post-quantum security (G1), efficient dynamic group management (G2), forward security (G3), and provable security in standard model (G4).

Technical Challenges. One approach for achieving the above goals is to follow the research line of OTS-based group signatures. A member requests to join a group with its OTS public key and the group manager GM adds the member to the group by signing that member’s key using GM’s OTS secret key. The member can anonymously sign a message using the OTS secret key corresponding to the public key certified with signatures from GM. A dynamic group with a large size demands a large number of OTS key pairs of GM. Updating the group public key whenever the group changes might pose an inconvenience for existing members. A natural way for achieving G2 is to use a Merkle tree to manage GM’s OTS key pairs, where each leaf corresponds to a key pair of GM. Note that the authentication paths of the signatures generated by GM would share nodes in the tree. An adversary can launch a **sibling attack** by exploiting the shared information to infer the identity of the member who generates the group signature and thus break the anonymity. This naturally leads to the question of achieving resilience against sibling attacks in a dynamic group with a large group size. To answer this question, we need to ensure that the OTS key pairs of GM are randomly assigned to members to be joined. A straightforward randomization approach using pseudorandom permutations is impractical since the member can dynamically join the group with unpredictable keys and handling permutations over a large set becomes computationally infeasible. Hence, realizing random selection of a large number of GM’s OTS key pairs remains a challenge.

To meet the requirements of G2 and G3, we cannot simply adopt the stateless hypertree used in previous signature schemes, such as those in the SPHINCS+ family [7,30], where secret keys remain static throughout the scheme’s lifetime and can derive all keys within the hypertree. However, incorporating dynamic group management (G2) or key-evolving techniques for achieving PFS (G3) inevitably introduces dynamic states for managing the group structure or evolving keys, thereby violating their stateless nature. For instance, the FORS+C layer of SPHINCS+C [30] involves more than 2^{77} leaves, which may require extensive state management. Thus, we need to seek an optimized stateful hypertree structure suitable for G2 and G3 from scratch. Meanwhile, we should avoid the repeated selection of the same key pair in the random selection process. In addition, all construction details should satisfy G4, which makes it harder to build, especially when considering adversaries who can adaptively register both honest and malicious members and compromise honest members.

Construction Ideas. Our construction starts from organizing OTS key pairs to align with all the aforementioned design goals, following the approach of stateful signature schemes like XMSS-T [31], which offer a more efficient signature structure compared to stateless alternatives such as SPHINCS+C [30]. Our main idea is to adopt a hypertree structure which consists of a GM’s OTS hypertree and multiple mixed OTS mount sub-trees (MMT) with obfuscated OTS layers. The leaves of the trees in the hypertree are used to sign the root of the trees or the sub-trees below. This structure allows GM to generate one tree (sub-tree) on each layer on demand, significantly reducing signature sizes and computational overhead compared to a single large tree. To further optimize GM’s computational cost and members’ storage requirements, the leaves computed during constructing MMT not only derive from OTS key pairs of GM but also from members’ OTS key pairs. This design results in a mixed sub-layer in MMT. However, a security issue arises if the members’ OTS secret keys corresponding to their OTS public keys in the mixed sub-layer are used to sign messages directly. In this case, the group signatures generated from these two types of leaves differ in form and can reflect a certain chronological order. An adversary can infer the identity of the member who generates the challenge group signature according the joining order of members. Thus, all keys in the above mixed sub-layer are used to sign the OTS public keys (in the bottom member sub-layer) from members, which will be used to sign messages. To realize the indistinguishability of the group signatures, we bind the OTS public keys of GM and members with fake Identity-OPK ciphertexts (IOCs) and members’ IOCs, respectively. The hashes generated from these OTS public keys and ciphertexts serve as the real leaves of MMT.

To realize G3, we use PPRFs to generate OTS key pairs for GM. Within each layer of the hypertree structure, all OTS key pairs of GM are derived using a PPRF specific to that layer. GM maintains the punctured key and state for each PPRF to track which indices and corresponding key pairs have already been used. A used index will be punctured to prevent the reuse of assigned key pairs, thereby ensuring forward security. After each puncturing operation, the BDS state of each tree is updated accordingly, enabling on-demand Merkle proof generation for the remaining available OTS key pairs. Simultaneously, each member uses a pseudorandom generator to generate its OTS key pairs, ensuring the forward security of the signatures.

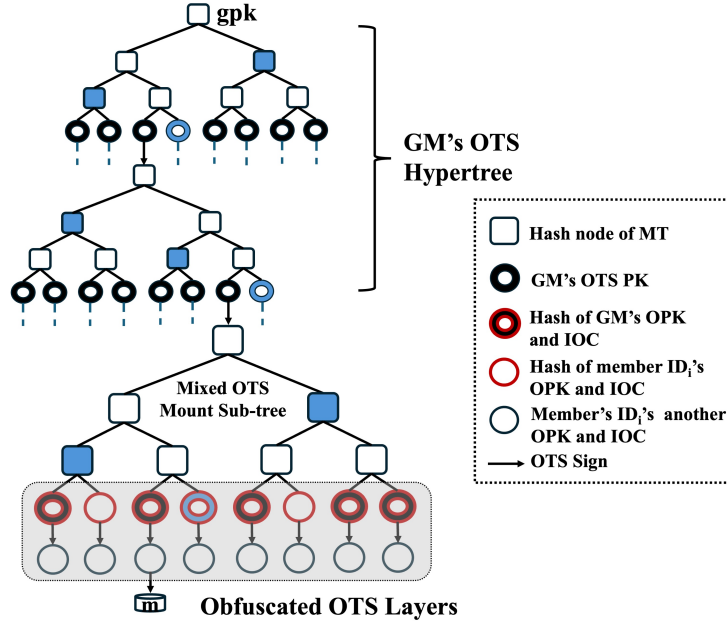


Fig. 2: A visualization of HHGS with a compact structure. The elements highlighted in blue are included in the group credential used for signing m .

A critical component of our design is to assign OTS key pairs in the mixed sub-layer to members dynamically. To achieve resilience against sibling attacks, the assignment must be random and independent. Random functions can be used to realize the random selection of OTS key pairs, i.e., determine the index of a leaf in the mixed sub-layer randomly. However, the outputs of existing random functions (like keyed hash functions) are not uniformly random, potentially leading to collisions in indices chosen by GM. There are two main challenges for resolving collisions: how to know which index corresponding to a OTS key pair has been used yet (Q1)? and how to randomly choose an unused key pair from the remaining available key pairs (Q2)? For Q1, PPRFs can be used to trace allocated indices, thereby supporting the subsequent random selection approaches. To answer Q2, we explore two approaches: recomputation and random range. Recomputation iteratively modifies the input of the random function (e.g., appending a counter) and repeats the calculation to generate a new index until an unused index is found. This approach becomes computationally expensive as the number of available indices diminishes. Alternatively, random range allows direct sampling from unused indices tracked by the PPRFs, significantly reducing the overhead associated with collision resolution.

Construction Preliminaries. To facilitate the description of our construction, we first review the hypertree structure [7,30] and define functions for obtaining/verifying the authentication credential of a message based on the hypertree.

Hypertree. A hypertree (i.e., a tree of trees) of total height $h \in \mathbb{N}$ consists of d layers of trees, each having height $h' = \frac{h}{d}$. On layer $i \in [d]$, the hypertree has $TN_i = 2^{(d-i)h'}$ trees. Let HI_i denote the index of a OTS key pair at the i -th layer. The OTS key pair (osk^{HI_1}, opk^{HI_1}) belonging to a tree at the bottom layer (i.e., layer 1), which is to be used to sign a message m . We denote HI_i as the first $(d+1-i) \times h'$ bits of HI_1 for $i \in [d]$. On layer $1 < i \leq d$, the OTS key pair (osk^{HI_i}, opk^{HI_i}) with the index HI_i is used to sign the root of the tree which has the OTS public key $osk^{HI_{i-1}}$ as a leaf node. For a message m , its authentication credential $\text{Auth}[m] = \{(opk^{HI_i}, p_{lf}^{HI_i}, \sigma^{HI_i})\}_{i \in [d]}$ can be generated (verified) by running the following **GetAuth** (**AuthVrfy**) algorithm. The i -th tuple in $\text{Auth}[m]$ contains the OTS public key opk^{HI_i} for verifying σ^{HI_i} , the Merkle proof $p_{lf}^{HI_i}$ for authenticating opk^{HI_i} , and the signature σ^{HI_i} on the message m ($i = 1$) or on the root of a below tree including $opk^{HI_{i-1}}$ ($i > 1$).

–**GetAuth**(sk, m, HI_1). Given a secret key sk , a message m , and an index HI_1 , this algorithm outputs the authentication credential $\text{Auth}[m]$ for m . For layers 1 to d , this algorithm repeatedly uses sk to generate the tree at the i -th layer which include the OTS key pair (osk^{HI_i}, opk^{HI_i}) and computes the Merkle proof $p_{lf}^{HI_i}$ for the leaf lf associated with opk^{HI_i} . Meanwhile, the root of the tree generated at the end of the previous

iteration is signed by the OTS secret key computed in the next iteration. The signature $\sigma^{\text{H}1}$ on m is signed by $\text{osk}^{\text{H}1}$. Finally, this algorithm outputs $\text{Auth}[m] := \{(\text{opk}^{\text{H}1}, \text{pf}_{\text{lf}}^{\text{H}1}, \sigma^{\text{H}1})\}_{i \in [d]}$.

– $\text{AuthVrfy}(pk, m, \text{Auth}[m])$. Given a public key pk , a message m , and an authentication credential $\text{Auth}[m]$, this algorithm checks the validity of $\text{Auth}[m]$. For layers 1 to d , this algorithm gets repeated to check whether the leaf lf derived from $\text{opk}^{\text{H}1}$ belongs to the tree on that layer and whether $\sigma^{\text{H}1}$ is a valid signature on the message m ($i = 1$) or the root Rt_i ($i > 1$). During the verification, the root of each tree is computed with the corresponding leaf and Merkle proof in $\text{Auth}[m]$ by following the root-computation algorithm in [6]. In the final repetition, the generated Rt_d for the root of the top-most tree on layer d is compared to pk . Finally, this algorithm outputs 1 if $\text{Rt}_d = pk$ and 0 otherwise.

Detailed Algorithms of HHGS. We define the algorithms of HHGS as follows:⁸.

• $\text{GMInit}(1^\kappa, \text{SP})$: Given the security parameter 1^κ and a setup parameter $\text{SP} = (h, d)$, the group manager GM initializes a set of PPRF schemes as $\{pms_{F_i} := F_p^i.\text{Setup}(1^\kappa)\}_{0 \leq i \leq d}$, a OTS scheme S as $pms_S := S.\text{Setup}(1^\kappa)$, a MT scheme M as $pms_M := M.\text{Setup}(1^\kappa)$, a AE scheme E as $pms_E := E.\text{Setup}(1^\kappa)$, a PRG scheme G as $pms_G := G.\text{Setup}(1^\kappa)$, a CRH scheme H as $(hk, pms_H) := H.\text{Setup}(1^\kappa)$, the PRF schemes F_1 and F_2 as $pms_{F_i} := F_i.\text{Setup}(1^\kappa)$ for $i \in [2]$. We assume that the spaces of each building block are compatible with those of other building blocks when their algorithms interact or invoke each other. The system parameter is set as $\text{Pms} := (\{pms_{F_i}\}_{0 \leq i \leq d}, pms_S, pms_M, pms_E, pms_G, pms_H, pms_{F_1}, pms_{F_2}, hk)$.

For $0 \leq i \leq d$, GM computes a random secret key $k_p^i := F_p^i.\text{KGen}(pms_{F_i})$ to generate OTS key pairs. The i -th layer for $1 \leq i \leq d$ is a hypertree structure, and the 0-th layer is used to mount the members' OTS keys (which will be defined later). GM also generates a random secret key $k_{GM} := F_1.\text{KGen}(pms_{F_1})$ for F_1 . The group secret key sk_{GM} consists of the random secret keys of PPRFs and F_1 , i.e., $\text{sk}_{GM} := (\{k_p^i\}_{0 \leq i \leq d}, k_{GM})$. Then, GM computes the group public key gpk through the following steps:

- Compute the randomnesses $\{rs_{GM}^{d,1,v}\}_{v \in [2^{h'}]}$ to be used in the generation of leaves $\{\text{lf}_{d,1,v}\}_{v \in [2^{h'}]}$, where each randomness value is defined as $rs_{GM}^{d,1,v} := F_p^d(k_p^d, v)$, with v representing a h' -bit index of the leaf $\text{lf}_{d,1,v}$ in the top-most tree of the hypertree structure.
- Compute the OTS key pairs $\{(\text{osk}_{GM}^{d,1,v}, \text{opk}_{GM}^{d,1,v}) := S.\text{KGen}(rs_{GM}^{d,1,v})\}_{v \in [2^{h'}]}$ and the leaves $\{\text{lf}_{d,1,v} := \text{opk}_{GM}^{d,1,v}\}_{v \in [2^{h'}]}$.
- Generate the group public key $\text{gpk} := \text{Tr}_{d,1}.\text{Rt}$, where top-most tree $\text{Tr}_{d,1}$ is built via the tree-building algorithm $(\text{Tr}_{d,1}, \text{st}_{\text{BDS}}^{d,1}) := M.\text{Build}(\{\text{lf}_{d,1,v}\}_{v \in [2^{h'}]}).$

Meanwhile, GM initializes the group management state st_{GM} to store a registered identity list $\text{RegIDL} := \emptyset$ recording the identities of registered group members and a revocation list $\text{RL} := \emptyset$.

• $\text{MInit}(\text{ID})$: The member ID computes a secret seed $sd_{\text{ID}} := G.\text{SGen}(pms_G)$ and defines its secret key as $\text{sk}_{GM} := sd_{\text{ID}}$. Meanwhile, ID initializes $\text{st}_{\text{ID}} := \emptyset$.

• $\text{MRegGen}(\text{sk}_{\text{ID}})$: ID runs this algorithm to generate $\ell_r \in \mathbb{N}$ OTS public keys to be registered. The member first gets the current secret seed sd_{ID} from its secret key sk_{GM} . For $v \in [\ell_r]$, ID sets the PRG seed as $sd_{\text{ID}}^0 := sd_{\text{ID}}$ and generates the v -th OTS key pair $(\text{osk}_{\text{ID}}^v, \text{opk}_{\text{ID}}^v) := S.\text{KGen}(rs_{\text{ID}}^{v,1})$, where $(rs_{\text{ID}}^{v,1}, rs_{\text{ID}}^{v,2}) := G(rs_{\text{ID}}^v)$ and $(sd_{\text{ID}}^v, rs_{\text{ID}}^v) := G(sd_{\text{ID}}^{v-1})$. Finally, the registration file is set as $\text{RF}_{\text{ID}} := \{\text{opk}_{\text{ID}}^v\}_{v \in [\ell_r]}$, and the secret key is updated to $\text{sk}_{\text{ID}} := sd_{\text{ID}}^{\ell_r}$.

• $\text{Join}(\text{sk}_{GM}, \text{st}_{GM}, \text{ID}, \text{RF}_{\text{ID}})$: For each $\text{opk}_{\text{ID}} \in \text{RF}_{\text{ID}}$, GM randomly selects an unused *mount index* $\text{MI} := \text{RandSelect}(\text{sk}_{GM}, \text{ID}, \text{opk}_{\text{ID}})$ that is to mount opk_{ID} (at the layer 0 below the hypertree), where the concrete steps of the random selection function RandSelect is defined later. We further decompose MI into two sub-indices as $\text{MI} = \text{sTI}||\text{sMI}$, where $\text{sTI} := \text{MI}(h)$ comprises the first h bits of MI , indexing a MMT. The remaining $|\text{MI}| - h$ bits specify the leaf index, sMI , within the corresponding MMT. Then, GM adds the tuple $(\text{ID}, \text{opk}_{\text{ID}}^{\text{MI}}, \text{MI})$ to a registration cache file $\text{RFCa}_{\text{MMT}}^{\text{sTI}}$ for the sub-tree $\text{MMT}[\text{sTI}]$ with the index sTI . Note that sTI is also the index of the leaf in the hypertree that is used to sign the root of $\text{MMT}[\text{sTI}]$.

GM checks whether $\text{MMT}[\text{sTI}]$ has been created by computing $F_p^1(k_p^1, \text{sTI})$. It then generates authentication credentials for the member's OTS keys in $\text{RFCa}_{\text{MMT}}^{\text{sTI}}$ based on one of the following two cases for constructing the obfuscated OTS layer in MMT.

⁸ We present the pseudocodes of the main algorithms in Appendix D.

- **Case 1** ($F_p^1(k_p^1, s\text{TI}) \neq \perp$) for MMT initialization: GM needs to build $\text{MMT}[s\text{TI}]$ using member's keys in $\text{RFCa}_{\text{MMT}}^{s\text{TI}}$ and the supplementary OTS keys generated by itself (for future usage). Let L_{mI} be a list storing the indices of members in $\text{RFCa}_{\text{MMT}}^{s\text{TI}}$ and $\bar{\text{L}}_{\text{mI}}$ be the complement of L_{mI} such that $\text{L}_{\text{mI}} \cup \bar{\text{L}}_{\text{mI}}$ covers all indices of leaves in $\text{MMT}[s\text{TI}]$. GM prepares the leaves of $\text{MMT}[s\text{TI}]$ based on L_{mI} and $\bar{\text{L}}_{\text{mI}}$ respectively as follows.
 - For each index $s\text{MI} \in \text{L}_{\text{mI}}$, GM generates an Identity-OPK ciphertext (IOC) $c_{\text{MI}}^1 := \text{E.Enc}(ke_{\text{MI}}^1, \text{ID}; re_{\text{MI}}^1)$, where $ke_{\text{MI}}^1 := F_1(k_{\text{GM}}, \text{"K-Mix"} || \text{MI})$ and $re_{\text{MI}}^1 := F_1(k_{\text{GM}}, \text{"R-Mix"} || \text{MI})$. It computes a leaf $\text{lf}_{s\text{MI}} := H(\text{opk}_{\text{ID}}^{\text{MI}} || c_{\text{MI}}^1)$ in $\text{MMT}[s\text{TI}]$. It sets $c_{\text{MI}}^2 := \emptyset$, $\text{opk}_{\text{GM}}^{\text{MI}} := \emptyset$ and $\sigma_{\text{GM}}^{\text{MI}} := \emptyset$. These values should be completed by the member during signing a message. Additionally, GM also punctures MI as $k_p^{0'} := F_p^0.\text{Punc}(k_p^0, \text{MI})$ at level 0 (in a MMT).
 - For each index $s\text{MI}' \in \bar{\text{L}}_{\text{mI}}$, GM sets $\text{MI}' := s\text{TI} || s\text{MI}'$ and computes the key pair $(\text{osk}_{\text{GM}}^{\text{MI}'}, \text{opk}_{\text{GM}}^{\text{MI}'}) := \text{S.KGen}(rs_{\text{GM}}^{\text{MI}'})$, where $rs_{\text{GM}}^{\text{MI}'} := F_p^0(k_p^0, \text{MI}')$. It computes the encryption key $ke_{\text{MI}'}^1 := F_1(k_{\text{GM}}, \text{"K-Mix"} || \text{MI}')$ and randomness $re_{\text{MI}'}^1 := F_1(k_{\text{GM}}, \text{"R-Mix"} || \text{MI}')$ to generate a fake IOC (encrypting GM) for $\text{opk}_{\text{GM}}^{\text{MI}'}$, as $c_{\text{MI}'}^1 := \text{E.Enc}(ke_{\text{MI}'}^1, \text{GM}; re_{\text{MI}'}^1)$.
- Then, GM runs $(\text{MMT}[s\text{TI}], \text{st}_{\text{BDS}}^{\text{MMT}[s\text{TI}]}) := \text{M.Build}(\{\text{lf}_v\}_{v \in [2^{|\text{sMI}|}]})$ to build the sub-tree $\text{MMT}[s\text{TI}]$, and obtains $\text{Auth}[\text{MMT}[s\text{TI}].\text{Rt}] := \text{GetAuth}(\{k_p^i\}_{i \in [d]}, \text{MMT}[s\text{TI}].\text{Rt}, s\text{TI})$. Next, it derives the Merkle proof of each $\text{lf}_{s\text{MI}}$ associated with ID's OTS key as $\text{pf}_{\text{lf}_{s\text{MI}}} := \text{M.GetPrf}(\text{st}_{\text{BDS}}^{\text{MMT}[s\text{TI}]}, \text{lf}_{s\text{MI}})$. To achieve forward security, GM punctures the used leaves on the hypertree as $k_p^{i'} := F_p^i.\text{Punc}(k_p^i, \text{MI}(w))$ for the PPRF F_p^i associated with the tree on the i -th layer, where $1 \leq i \leq d$ and $\text{MI}(w)$ denotes the first $w = \frac{(d+1-i)h}{d}$ bits of MI.
- In order to compute the Merkle proofs of the GM's OTS keys indexed by any $s\text{MI}' \in \bar{\text{L}}_{\text{mI}}$ in the future, GM stores the corresponding state $\text{st}_{\text{BDS}}^{\text{lf}_{s\text{MI}'}}$ and $\text{Auth}[\text{MMT}[s\text{TI}].\text{Rt}]$ into st_{GM} .
- **Case 2** ($F_p^1(k_p^1, s\text{TI}) = \perp$) for MMT usage: In this scenario, each $\text{opk}_{\text{ID}}^{\text{MI}}$ in $\text{RFCa}_{\text{MMT}}^{s\text{TI}}$ must reside within the member OTS sub-layer and should therefore be authenticated using GM's OTS key in the mixed sub-layer. GM generates the keys $(\text{osk}_{\text{GM}}^{\text{MI}}, \text{opk}_{\text{GM}}^{\text{MI}}) := \text{S.KGen}(rs_{\text{GM}}^{\text{MI}})$, where $rs_{\text{GM}}^{\text{MI}} := F_p^0(k_p^0, \text{MI})$ and the IOC c_{MI}^1 for $\text{opk}_{\text{GM}}^{\text{MI}}$ as in Case 1. Then, GM computes the leaf $\text{lf}_{s\text{MI}} := H(\text{opk}_{\text{GM}}^{\text{MI}} || c_{\text{MI}}^1)$ and gets its Merkle proof $\text{pf}_{\text{lf}_{s\text{MI}}} := \text{M.GetPrf}(\text{st}_{\text{BDS}}^{\text{lf}_{s\text{MI}}}, \text{lf}_{s\text{MI}})$, where $\text{st}_{\text{BDS}}^{\text{lf}_{s\text{MI}}} \in \text{st}_{\text{GM}}$. It also retrieves $\text{Auth}[\text{MMT}[s\text{TI}].\text{Rt}]$ from st_{GM} , and generates an IOC $c_{\text{MI}}^2 := \text{E.Enc}(ke_{\text{MI}}^2, \text{ID}; re_{\text{MI}}^2)$ with $ke_{\text{MI}}^2 := F_1(k_{\text{GM}}, \text{"K-Mem"} || \text{MI})$ and $re_{\text{MI}}^2 := F_1(k_{\text{GM}}, \text{"R-Mem"} || \text{MI})$. Lastly, GM creates the signature $\sigma_{\text{GM}}^{\text{MI}} := \text{S.Sign}(\text{osk}_{\text{GM}}^{\text{MI}}, \text{opk}_{\text{ID}}^{\text{MI}} || c_{\text{MI}}^2)$ over $\text{opk}_{\text{ID}}^{\text{MI}}$. GM removes $\text{st}_{\text{BDS}}^{\text{lf}_{s\text{MI}}}$ and punctures MI as $k_p^{0'} := F_p^0.\text{Punc}(k_p^0, \text{MI})$.

The authentication credential for each $\text{opk}_{\text{ID}}^{\text{MI}} \in \text{RFCa}_{\text{MMT}}^{s\text{TI}}$ is set as $\text{Auth}[\text{opk}_{\text{ID}}^{\text{MI}}] := (\text{opk}_{\text{GM}}^{\text{MI}}, c_{\text{MI}}^1, c_{\text{MI}}^2, \text{pf}_{\text{lf}_{s\text{MI}}}, \sigma_{\text{GM}}^{\text{MI}}, \text{MMT}[s\text{TI}].\text{Rt}, \text{Auth}[\text{MMT}[s\text{TI}].\text{Rt}])$. GM iterates through all OTS keys in RF_{ID} , processing each as described. It then appends every $\text{Auth}[\text{opk}_{\text{ID}}^{\text{MI}}]$ to the registration result RR_{ID} and returns it to ID, which updates its state st_{ID} accordingly.⁹ During the execution of this protocol, the sent registration file and the received registration result of ID are recorded in the transcript Vt . Meanwhile, GM adds ID to the registered identity list RegIDL if ID has not previously submitted a join request.

We assume that the join sessions for processing different registration requests from members share the state st_{GM} . These sessions can be collectively handled by GM through a unified join-hyper routine, reducing the GM's computational and storage costs.

- **Sign**($sk_{\text{ID}}, \text{st}_{\text{ID}}, m$): A member ID computes the next secret seed and the current key generation randomness as $(sd_{\text{ID}}^v, rs_{\text{ID}}^v) := G(sd_{\text{ID}}^{v-1})$ and $(rs_{\text{ID}}^{v,1}, rs_{\text{ID}}^{v,2}) := G(rs_{\text{ID}}^v)$, respectively. Using the randomness $rs_{\text{ID}}^{v,1}$, ID generates a OTS key pair $(\text{osk}_{\text{ID}}, \text{opk}_{\text{ID}}) := \text{S.KGen}(rs_{\text{ID}}^{v,1})$. Next, ID must determine the sub-layer (either the mixed sub-layer or the member sub-layer) where the OTS public key opk_{ID} resides within the obfuscated layer. To accomplish this, ID retrieves the first (i.e., most recent) authentication credential $\text{Auth}[\text{opk}_{\text{ID}}]$ from st_{ID} . If $\sigma_{\text{GM}}^{\text{MI}} \in \text{Auth}[\text{opk}_{\text{ID}}]$ is non-empty, this indicates that opk_{ID} belongs to the member sub-layer, and osk_{ID} can be directly used to sign the message m , producing the signature $\sigma_{\text{ID}} := \text{S.Sign}(\text{osk}_{\text{ID}}, m)$. Eventually, ID constructs the group signature as $\sigma := (\text{opk}_{\text{ID}}, \text{Auth}[\text{opk}_{\text{ID}}], \sigma_{\text{ID}})$.

For $\sigma_{\text{GM}}^{\text{MI}} = \emptyset$, opk_{ID} must reside in the mixed sub-layer. In this case, ID generates an additional key pair $(\text{osk}_{\text{ID}}', \text{opk}_{\text{ID}}') := \text{S.KGen}(rs_{\text{ID}}^{v,2})$ for signing the message, i.e., $\sigma_{\text{ID}} := \text{S.Sign}(\text{osk}_{\text{ID}}', m)$. Additionally, ID

⁹ GM may optionally return ke_{MI}^1 and ke_{MI}^2 to enable the member to check the corresponding ciphertext.

randomly generates an IOC c'_{ID} for opk'_{ID} as in the Join protocol but uses the random key $ke' \xleftarrow{\$} \mathcal{K}_{AE}$ and the randomness $re' \xleftarrow{\$} \mathcal{ER}_{AE}$. Next, ID produces the signature $\sigma'_{ID} := \text{S.Sign}(osk_{ID}, opk'_{ID} || c'_{ID})$ and completes the values regarding $(opk'_{GM}, \sigma'_{GM}, c'_{MI})$ in the authentication credential $\text{Auth}[opk_{ID}]$ by incorporating opk_{ID} , σ'_{ID} , and c'_{ID} . The group signature is then defined as $\sigma := (opk'_{ID}, \text{Auth}[opk_{ID}], \sigma_{ID})$.

Finally, ID updates its secret key with the new secret seed, i.e., $sk'_{ID} := sd_{ID}$, and gets an updated state $st'_{ID} := st_{ID} \setminus (opk_{ID}, \text{Auth}[opk_{ID}])$.

- **Verify(gpk, m, σ , RL)**: The verifier parses $\sigma = (opk'_{ID}, \text{Auth}[opk_{ID}], \sigma_{ID})$ and $\text{Auth}[opk_{ID}] = (opk'_{GM}, c'_{MI}, c'_{MI}, pf_{If_{sMI}}, \sigma'_{GM}, \text{MMT}[sTI].Rt, \text{Auth}[\text{MMT}[sTI].Rt])$. This algorithm outputs 0 if one of the following conditions is met: i) $c'_{MI} \in \text{RL}$; ii) $\text{AuthVrfy}(gpk, \text{MMT}[sTI].Rt, \text{Auth}[\text{MMT}[sTI].Rt]) = 0$; iii) $\text{M.Verify}(\text{MMT}[sTI].Rt, If_{sMI}, pf_{If_{sMI}}) = 0$, where $If_{sMI} = H(opk'_{GM} || c'_{MI})$; iv) $\text{S.Verify}(opk'_{GM}, opk'_{ID} || c'_{MI}, \sigma'_{GM}) = 0$; v) $\text{S.Verify}(opk'_{ID}, m, \sigma_{ID}) = 0$. Otherwise, the verifier outputs 1.

- **Open(gpk, sk_{GM} , m, σ)**: The algorithm outputs 0 if the group signature is not valid, i.e., $\text{Verify}(gpk, m, \sigma, T, \emptyset) = 0$. GM gets the index MI from σ according to its authentication credential. Next, it computes keys randomness (ke'_{MI}, re'_{MI}) and (ke'_{MI}, re'_{MI}) as in Join protocol based on k_{GM} . GM decrypts $\hat{ID}_1 := \text{E.Dec}(ke'_{MI}, c'_{MI}; re'_{MI})$ and $\hat{ID}_2 := \text{E.Dec}(ke'_{MI}, c'_{MI}; re'_{MI})$. If $\hat{ID}_1 = \text{GM}$ then return \hat{ID}_2 ; otherwise \hat{ID}_1 is returned.

- **Revoke(gpk, sk_{GM} , RL, ID, RSS)**: For each tuple $(\sigma, m) \in \text{RSS}$, GM returns 0 if $\text{Verify}(gpk, m, \sigma, \text{RL}) = 0$ or $\text{Open}(gpk, sk_{GM}, m, \sigma) \neq \text{ID}$. Otherwise, GM gets c'_{MI} from each $\sigma \in \text{RSS}$ and puts it into RL.

The correctness of HHGS is presented in Appendix B, which is basically guaranteed by the correctness of building blocks.

Instantiations of RandSelect(sk_{GM} , ID, opk_{ID}). One method for random index selection applies a pseudo-random permutation to shuffle all leaf indices in a HHGS instance [2]. However, this becomes impractical for trees with over 2^{20} leaves due to the scaling costs of permutation [2, §4.2].

Alternatively, we propose using a short-range PRF scheme, F_2 , with range $\{0, 1\}^{\ell_I}$ to generate indices for OTS mounting. The first approach mimics random sampling without replacement, while the second directly selects from remaining available indices. In the first approach, collisions may occur due to the birthday paradox. To address this, the F_2 process can be repeated with different inputs to avoid collisions (referred to as **AP1**). However, as the number of available indices, denoted by ϕ , decreases, the likelihood of collisions increases. When this happens, we can leverage the PPRF keys, which encode the available keys, to directly select an index from the remaining indices (referred to as **AP2**), provided the number of remaining indices is not excessively large (e.g., $\leq 2^{20}$). We define ξ as the threshold for the number of available indices, marking the point at which we switch from the initial random selection method to the latter approach. To support this second selection method, we introduce an additional algorithm for PPRF, denoted as $\text{GetMsg}(k, i)$. This algorithm enables the retrieval of the i -th unpunctured message. Specifically, $\text{GetMsg}(k, i)$ takes as input the current key $k \in \mathcal{K}_{PRF}$ and returns the i -th accessible message x_i from the remaining computable message set. The realization of this algorithm is detailed in Appendix A.

Specifically, GM generates a random selection key, $k_s := F_1(k_{GM}, \text{"RandSel"})$, and computes the index KI as follows:

- **AP1**. If $\phi < \xi$, compute $\text{KI} := F_2(k_s, \text{ID} || opk_{ID} || ctr)$ using an incremental counter ctr (initialized to 0) and repeat this process until $F_p^0.\text{Eval}(k_p^0, \text{KI}) \neq \perp$. Thereafter, $\phi := \phi - 1$.
- **AP2**. Otherwise, compute $tmp := F_2(k_s, \text{ID} || opk_{ID})$, then set $tmp := tmp \bmod \phi$, and derive $\text{KI} := F_p^0.\text{GetMsg}(k_p^0, tmp)$.

The complexity of the above instantiation of RandSelect is determined by the signing $Y = 2^{\tilde{h}}$ of HHGS, where $\tilde{h} = h + \frac{h}{d}$ is the total height of the hypertree structure. The computational cost of **AP1** is determined by the number n_r of PRF recomputations for selecting an unused index. Let ϕ denote the remaining indices that can be chosen. The probability of selecting an occupied index is about $1 - \frac{\phi}{Y}$, while considering PRF as a random function. The number of PRF recomputations follows a geometric distribution with success probability $p = \frac{\phi}{Y}$. The expected value of n_r is $\mathbb{E}[n_r] = \frac{Y}{\phi}$, which can be used as the threshold for switching between AP1 and AP2. As ϕ decreases, $\mathbb{E}[n_r]$ grows inversely proportional to ϕ . The total computational cost of **AP1** scales with the reduction in the number of remaining indices, ϕ . In contrast, the cost of **AP2** is directly proportional to ϕ , as determined by the cost of GetMsg (analyzed in Appendix A). Consequently, as ϕ decreases, the overhead of **AP2** also reduces.

In a nutshell, the signing capacity, Y , of HHGS is inherently limited. For instance, when $\mathbb{E}[n_r] = \phi = O(2^{20})$, Y is in $O(2^{40})$.

Security Analysis. The correctness of HHGS (analyzed in Appendix B) follows from the correctness of its building blocks. The security of HHGS is demonstrated through the following theorems with detailed proofs in Appendix C.

Theorem 1. *Assume that the OTS scheme S , the PPRF functions $\{F_p^i\}_{0 \leq i \leq d}$, the PRF functions F_1 and F_2 , the MT scheme M , the AE scheme E , the PRG scheme G , and the hash function H are secure as defined in Section 2, then HHGS satisfies traceability.*

The proof employs a sequence of game transformations, starting with the **Trace** experiment, to incrementally restrict the adversary’s ability to forge signatures under HHGS. Starting from the original **Trace** experiment, the challenger incrementally modifies the game by introducing abort conditions to enable the reduction of security to the hardness of underlying cryptographic primitives.

Theorem 2. *Under the same assumptions as Theorem 1, HHGS achieves anonymity.*

The core idea involves progressively replacing PRF values, random selection keys, encryption keys, and randomness with random values, leveraging the security of the underlying primitives such as PRF and PRG. Consequently, the OTS keys and ciphertexts associated with the $\mathcal{O}_{\text{PriTest}}(\text{ID}_0^*, \text{ID}_1^*, m)$ query for the challenged members and **GM** are shown to follow identical distributions, preventing the adversary from distinguishing between them. Finally, the mount indices of the target OTS keys are selected randomly, limiting the adversary’s success to a random guess due to the independence of the selection process.

The security of HHGS also ensures its resilience to sibling attacks. The key insight lies in the randomness of the selection process for the leaves at level 0 of the hypertree structure, which prevents the attacker from distinguishing between the positions of the two candidate challenge signatures on the tree. Although the attacker knows the results of the other $2^h - 2$ selections, the two remaining unknown selections concerning the challenged identities’ OTS keys are drawn randomly and without replacement. Since the selection process is independent and unbiased in the final game, the distribution of these two unknown selections appears identical from the attacker’s perspective.

4.2 Signing Capability Extension

This section focuses on improving the signing capability of HHGS. A group signature with a large signing capacity is vital for applications with numerous users or frequent signing demands, such as privacy-preserving IoT networks in smart cities or industrial IoT systems. A naive solution to support larger groups and signing capabilities is to increase the parameters of HHGS. However, this is not efficient due to the complexity of the **RandSelect** algorithm analyzed in the previous section. Alternatively, using multiple independent PPRF keys could address this inefficiency, but it introduces the challenge of privately determining which member’s OTS key should be associated with which PPRF key. This requires a method to securely partition members’ OTS keys into subgroups.

Construction Overview. We address this challenge with HHGS^+ , an improved FSDGS scheme that generically builds on HHGS. Specifically, HHGS^+ initializes $W \in \mathbb{N}$ copies of HHGS and employs the MT scheme M to construct a tree, where the leaves are the group public keys $\{\text{gpk}_i\}_{i \in [W]}$ generated by the HHGS instances. To ensure efficiency, we use a single Merkle tree over a hypertree structure. This design minimizes additional tree layers, avoiding the need to store more PPRF keys. Moreover, using multiple independent HHGS instances could enable parallelization of group management.

To securely group members’ OTS keys, we randomly associate each OTS key with a specific HHGS instance using a short-range PRF function, which takes as input a OTS key and its owner’s identity. Due to the short range of the PRF, different OTS keys may yield the same PRF value (referred to as OTS collision), enabling the grouping feature. However, as OTS collisions are not uniformly distributed among PRF values, the volume of each HHGS instance poses a challenge. That is, the non-uniform distribution of OTS collisions creates challenges in *balancing the load across HHGS instances and ensuring the security*.¹⁰ In an insecure

¹⁰ Note that, the **RandSelect** function selects nodes randomly from all available nodes (independent of MMTs), while the approach here introduces a structured process, selecting HHGS instances first, creating dependencies in the selection sequence.

scenario, referred to as the **Full-instance Attack (FIA)**, if a HHGS instance becomes full early in the process (e.g., while registering ID_0^*), it will not be selected in subsequent steps (e.g., during the registration of ID_1^*), potentially creating a bias in the remaining selections. This temporal dependency in the selection process allows an attacker to exploit the evolving probabilities as trees fill, compromising the randomness and enabling prediction of future instance choices. To counter this, each HHGS instance should contain sufficient redundancy (see Lemma 1) to accommodate the non-uniform distribution of PRF values, ensuring they can handle varying numbers of OTS keys securely.

Detailed Algorithms. We define the algorithms of HHGS^+ as follows:

- **GMInit($1^\kappa, \text{SP}$):** For $i \in [W]$, GM creates $W \in \mathbb{N}$ HHGS instances as $(\text{Pms}_i, \text{gpk}_i, \text{sk}_{\text{GM}}^i, \text{st}_{\text{GM}}^i, \text{RL}_i) := \text{HHGS.GMInit}(1^\kappa, \text{SP})$. Subsequently, it builds the Merkle tree instance as $(\text{Tr}', \text{st}'_{\text{BDS}}) := \text{M.Build}(\{\text{gpk}_i\}_{i \in [W]})$. GM sets $\text{gpk} := \text{Tr}'.\text{Rt}$, $\text{sk}_{\text{GM}} := \{\text{sk}_{\text{GM}}^i\}_{i \in [W]}$, $\text{Pms} := \{\text{Pms}_i\}_{i \in [W]}$, $\text{st}_{\text{GM}} := \{\text{st}_{\text{GM}}^i\}_{i \in [W]}$, and $\text{RL} := \{\text{RL}_i\}_{i \in [W]}$. We will use the building blocks initiated in HHGS instances such as the PRG scheme G , the PRF schemes F_1 and F_2 , and the MT scheme M . Additionally, GM samples a random PRF key for $k_{\text{GM}} := F_1.\text{KGen}(\text{pms}_{\text{PRF}})$. To save space, GM replaces with $k_{\text{GM}}^i := F_1(k_{\text{GM}}, "i||\text{SubInstance}")$, so that it can be recovered from k_{GM} .

- The algorithms MInit and MRegGen in HHGS^+ are identical to those in HHGS.

- **Join($\text{sk}_{\text{GM}}, \text{st}_{\text{GM}}, \text{ID}, \text{RF}_{\text{ID}}$):** Although GM can utilize a hyper-join routine to process join requests of members in batches, as implemented in HHGS, we focus here solely on ID for simplicity, to convey the main idea. Let LR_{ID} denote the list recording the indices of HHGS instances involved, and RF_{ID}^i represent the registration file used in the i -th HHGS instance.

For each opk_{ID}^j in RF_{ID} (where $j \in [\#\text{RF}_{\text{ID}}]$), GM randomly assigns it to a HHGS instance and fills RF_{ID}^i and LR_{ID} accordingly. Specifically, it computes the random selection key $k_j := F_1(k_{\text{GM}}, \text{ID}||\text{opk}_{\text{ID}}^j)$ and determines the HHGS instance index as $\text{il} := F_2(k_j, \text{ID}||\text{opk}_{\text{ID}}^j)$. Next, GM inserts opk_{ID}^j into $\text{RFCa}_{\text{ID}}^{\text{il}}$ and adds il to LR_{ID} if $\text{il} \notin \text{LR}_{\text{ID}}$.

For $i \in \text{LR}_{\text{ID}}$, GM executes $(\text{Vt}_i, \text{RR}_{\text{ID}}^i) \leftarrow \text{Join}(\text{sk}_{\text{GM}}^i, \text{st}_{\text{GM}}^i, \text{ID}, \text{RF}_{\text{ID}}^i, T_\delta)$ with ID, where $\text{RR}_{\text{ID}}^i = (k_{\text{ID}}, \{\text{Auth}[\text{opk}_{\text{ID}}^{i,j}]\}_{j \in [\#\text{RF}_{\text{ID}}^i]})$ is the authentication credentials generated for $\text{opk}_{\text{ID}}^{i,j}$, and $\text{opk}_{\text{ID}}^{i,j}$ is the j -th OTS key in the registration file RF_{ID}^i for the ID in the i -th HHGS instance. The overall transcript of the protocol is the union of the sub-transcripts from each join session, i.e., $\text{Vt} = \bigcup_{i \in \text{LR}_{\text{ID}}} \text{Vt}_i$. Meanwhile, GM appends each authentication credential in RR_{ID}^i with the corresponding index i to yield a modified registration result set $\{\hat{\text{RR}}_{\text{ID}}^i\}_{i \in \text{LR}_{\text{ID}}} = \{(i, \text{Auth}[\text{opk}_{\text{ID}}^{i,j}])\}_{i \in \text{LR}_{\text{ID}}, j \in [\#\text{RF}_{\text{ID}}^i]}$.

GM builds the states $\hat{\text{st}}_{\text{ID}} = \{\hat{\text{RR}}_{\text{ID}}^i\}_{i \in \text{LR}_{\text{ID}}}$ and $\text{st}_{\text{ID}}^{\text{pf}} := \{\text{gpk}_i, \text{pf}_{\text{gpk}_i}\}_{i \in \text{LR}_{\text{ID}}}$ for ID. Finally, it defines $\text{RR}_{\text{ID}} = (k_{\text{ID}}, \text{st}_{\text{ID}}^{\text{pf}}, \hat{\text{st}}_{\text{ID}})$, which is returned to ID. Upon receiving RR_{ID} , ID updates its states $\text{st} := \text{st} \cup (\text{st}_{\text{ID}}^{\text{pf}}, \hat{\text{st}}_{\text{ID}})$ and secret key $\text{sk}_{\text{ID}} := \text{sk}_{\text{ID}} \cup k_{\text{ID}}$. If necessary, the member may delete duplicate states and secret keys to save storage.

- **Sign($\text{sk}_{\text{ID}}, \text{st}_{\text{ID}}, m$):** ID first pops up the first tuple $(i, \text{Auth}[\text{opk}_{\text{ID}}^j])$ from $\hat{\text{st}}_{\text{ID}}$, where opk_{ID}^j is the OTS public key opk_{ID}^j (to be generated within HHGS.Sign later) for verifying the signature about to generate. It then retrieves $(\text{gpk}_i, \text{pf}_{\text{gpk}_i})$ from $\text{st}_{\text{ID}}^{\text{pf}}$ based on i . To proceed with the HHGS instance, ID constructs a sub-state $\text{st}_{\text{ID}}^i := \text{Auth}[\text{opk}_{\text{ID}}^j]$. For signing the message m , ID executes $(\text{sk}_{\text{ID}}^i, \sigma_{\text{ID}}^i) := \text{HHGS.Sign}(\text{sk}_{\text{ID}}, \text{st}_{\text{ID}}, T_\delta, m)$ and composes the final signature as $\sigma_{\text{ID}} := (\sigma_{\text{ID}}^i, \text{gpk}_i, \text{pf}_{\text{gpk}_i})$. Lastly, the tuple $(i, \text{Auth}[\text{opk}_{\text{ID}}^j])$ is removed from st_{ID} .

- **Verify($\text{gpk}, m, \sigma_{\text{ID}}, \text{RL}$):** The verifier parses $\sigma_{\text{ID}} = (\sigma_{\text{ID}}^i, \text{gpk}_i, \text{pf}_{\text{gpk}_i})$, and it performs two verification steps: i) $\text{mtV} := \text{M.Verify}(\text{gpk}, \text{lf}_{\text{gpk}_i}, \text{pf}_{\text{gpk}_i})$, where $(\text{gpk}_i, \text{pf}_{\text{gpk}_i}) \in \sigma_{\text{ID}}$; ii) $\text{dsV} := \text{HHGS.Verify}(\text{gpk}_i, m, \sigma_{\text{ID}}^i, \text{RL}_i)$. The index i can be inferred analogously from σ_{ID} . If $\text{mtV} = \text{dsV} = 1$, the algorithm outputs 1; otherwise it outputs 0.

- **Open($\text{gpk}, \text{sk}_{\text{GM}}, m, \sigma_{\text{ID}}$):** If $\text{M.Verify}(\text{gpk}, \text{gpk}_i, \text{pf}_{\text{gpk}_i}) = 1$, it outputs the execution result $\text{HHGS.Open}(\text{gpk}_i, \text{sk}_{\text{GM}}^i, m, \sigma_{\text{ID}}^i)$, where $(\text{gpk}_i, \text{pf}_{\text{gpk}_i}, \sigma_{\text{ID}}^i) \in \sigma_{\text{ID}}$.

- **Revoke**(gpk, sk_{GM}, RL, ID, RSS): For each signature $\sigma \in \text{RSS}$, GM gets the index i of the HHGS instance according to the Merkle proofs in σ , and runs $\text{RL}'_i := \text{HHGS.Revoke}(\text{gpk}_i, \text{sk}_{\text{GM}}^i, \text{RL}_i, \text{ID}, \sigma)$. After this, the revocation list is updated to $\text{RL}' := \{\text{RL}'_i\}$.

Security Analysis. The correctness of HHGS^+ is inherently ensured by the correctness of HHGS, MT, and PRF. Therefore, the analysis is omitted for simplicity. The security of HHGS^+ is essentially implied by that of HHGS, along with the security of PRF and M. Here, we mainly present the key security results through the following theorems and outline their proof strategies. We analyze the redundancy in HHGS to resist the FIA, as formally established by Lemma 1, thereby laying the security foundation for HHGS^+ .

Lemma 1 (Resilience of FIA). *Suppose each HHGS instance has a maximum capacity of Y leaves, and is expected to register at least Z OTS keys. Given $B = W \cdot Z$ total OTS keys from members, and F_2 is a secure PRF whose output distribution has a statistical distance of at most ϵ_{F_2} from the uniform distribution. Then, the probability that no HHGS instance exceeds its capacity is bounded by a negligible probability $\epsilon_{\text{FIA}} = \epsilon_{\text{FIA}}(\kappa)$, provided that Y satisfies the condition $Y \geq Z + \sqrt{2Z \ln((1 + O(\epsilon_{F_2}))/\epsilon_{\text{FIA}})}$.*

Proof. The scenario resembles the Coupon Collector's Problem [8], where we aim to collect Z values in each of W instances, and we wish to ensure that no instance exceeds its capacity. The number of OTS keys involved in each instance can be approximated using a binomial distribution. Let X_j denote the number of values assigned to the j -th instance. We aim to find the value of Y such that the probability of any instance exceeding its capacity is negligible. Formally, we want the event $A_j = \{X_j \geq Y\}$ (i.e., the j -th HHGS instance is full) to satisfy $\Pr(A_j) \leq \epsilon_{\text{FIA}}$. The expected number of OTS keys assigned to any HHGS instance is about $\mathbb{E}[X_j] \approx \frac{B}{W} = Z$. Meanwhile, the statistical distance ϵ_{F_2} between the output of F_2 and a uniform distribution introduces a slight deviation from perfect uniformity. This deviation affects the variance, thereby introducing a correction term in the overflow probability bound. Using the Chernoff bound [8], for large Z , the probability of the event A_j is given by $\Pr(X_j \geq Y) \leq \exp\left(-\frac{(Y-Z)^2}{2Z}\right) (1 + O(\epsilon_{F_2}))$, where \exp refers to the exponential function. To ensure that this probability is negligible, we require $\exp\left(-\frac{(Y-Z)^2}{2Z}\right) (1 + O(\epsilon_{F_2})) \leq \epsilon_{\text{FIA}}$. Taking the natural logarithm of both sides $-\frac{(Y-Z)^2}{2Z} \geq \ln((1 + O(\epsilon_{F_2}))/\epsilon_{\text{FIA}})$. Solving for Y , we have $Y \geq Z + \sqrt{2Z \ln((1 + O(\epsilon_{F_2}))/\epsilon_{\text{FIA}})}$. This complete the proof.

For our target $Z \in O(2^{40})$, Y can be configured around $2^{40} + 2^{20.5}$, ensuring a negligible overflow probability of $\epsilon_{\text{FIA}} \leq 2^{-128}$.

Theorem 3. *Assume that HHGS provides traceability and the MT scheme M is secure, then HHGS^+ satisfies traceability.*

The proof of Theorem 3 closely follows that of Theorem 1, with the primary distinction being the need for an additional game to reduce the security of HHGS^+ to that of MT.

Table 2: Performance. Runtime in milliseconds (ms) and size in kilobytes (KB).

κ	HHGS						HHGS ⁺					
	Signing Capacity (d)	Setup	Join	Sign	Verify	Sig. Size	Signing Capacity (d)	Setup	Join	Sign	Verify	Sig. Size
128	$2^{39}(12)$	1.53	301.68	2.3	1.08	7.95	$2^{59}(12)$	16×10^5	303.73	2.3	1.09	8.58
	$2^{40}(9)$	3.14	318.89	2.3	0.82	6.41	$2^{60}(9)$	33×10^5	323.07	2.3	0.83	7.04
	$2^{40}(19)$	0.74	294.59	2.3	1.64	11.58	$2^{60}(19)$	7.7×10^5	295.29	2.3	1.65	12.21
256	$2^{39}(12)$	7.29	654.62	12.87	5.86	20.25	$2^{59}(12)$	76×10^5	676.22	12.87	5.87	20.88
	$2^{40}(9)$	14.68	713.46	12.87	4.40	16.24	$2^{60}(9)$	153×10^5	730.85	12.87	4.41	16.87
	$2^{40}(19)$	3.62	632.91	12.87	8.77	29.72	$2^{60}(19)$	38×10^5	637.34	12.87	8.78	30.35

Theorem 4. *Assume that HHGS provides anonymity and the PRF families (F_1, F_2) are secure, then HHGS^+ satisfies anonymity.*

To prove this theorem, it suffices to demonstrate that random selection among HHGS instances does not leak identity-related information, since OTS keys within each HHGS instance are indistinguishable due to

its security. It is straightforward to observe that if no HHGS instance becomes fully occupied during the lifetime of HHGS⁺ (i.e., the resilience against FIA attacks), then OTS keys are randomly assigned to W independent HHGS instances with nearly identical probability (with negligible differences due to the security of F). Using Lemma 1, we show that by properly configuring the size of each HHGS instance, the probability of a successful FIA attack is negligible.

5 Evaluation

In this section, we evaluate the performance of our proposed FSDGS schemes. In the performance analysis, we instantiate the OTS scheme using the state-of-the-art OTS scheme WOTS+C [30,40], and use AES-GCM-SIV [28] for the authenticated encryption E.¹¹ For the PRG, we use counter-mode AES [3]. The hash functions in our schemes are implemented using the SHA-3 family [23]. The PPRF is implemented with the GGM construction [12]. Meanwhile, we consider 128-bit security against classical attacks and (roughly) double the security parameter to ensure post-quantum security.

For the parameters of WOTS+C, we aim to strike a trade-off between signature size and computational cost by selecting a relatively small number of 2^5 of signature chains, with a Winternitz parameter of 2^4 for the length of each chain, and a 9-bit checksum size. The signature size of HHGS is primarily determined by the signing capacity ($Y = 2^{\tilde{h}}$), selected as either 2^{39} or 2^{40} to accommodate the corresponding group size. The cost of authenticating one OTS in the Join protocol, involve at most $d+1$ OTS signing operations and the cost of RandSelect (about $O(2^{20})$ F_2 operations in the worst case). Let B (e.g., $B \in O(2^{15})$) denote the maximum number of OTS keys a member retains. Consequently, the credentials stored in the member's cache have a size of $O(B \cdot h)$. Members can adjust B based on the storage capacity of their auxiliary storage devices. We consider a practical scenario where, for the first join request, a member must submit at least B OTS keys, and the GM processes join requests from U members in a batch. For simplicity, we assume $B \cdot U = \sqrt{Y}$. Thus, the storage cost of the GM's state, dominated by the PPRF keys, is of magnitude $O(\sqrt{Y})$. The signing process on the member's device is efficient, with the worst-case computational cost dominated by two WOTS+C key generation and signing operations.

We benchmark the performance of the algorithms in our proposed schemes using a Raspberry Pi 3 to simulate the client's signing operation, and a PC with an Intel(R) Core(TM) i5-13500H, 2.60 GHz, and 16 GB RAM, which serves as the verifier and GM. Table 2 provides a summary of the performance of our schemes across key metrics, as in prior works [2,20,44].¹² For simplicity, we do not benchmark the full schemes. The performance of Open and Revoke can be inferred from that of Verify. Specifically, the cost of Open is primarily determined by a single Verify operation along with one encryption and one decryption. To revoke a single signature, the overhead of Revoke consists of the combined costs of one Verify and one Open operation. Additionally, the runtime of the join protocol is measured based on handling a single OTS key in the worst-case scenario, dominated by $O(2^{20})$ F_2 operations. The height of the upper layer Merkle tree of HHGS⁺ is $W = 2^{20}$ in our benchmark. As a result, the setup time (about a few minutes) for HHGS⁺ is W times that of HHGS.

6 Conclusions

We introduced the first FSDGS framework based on symmetric primitives, addressing key challenges in post-quantum secure group signatures. Our scheme, HHGS, employs a hierarchical OTS-key management structure with PPRF and obfuscation layers to enable efficient dynamic group management, ensure forward security, and defend against sibling attacks. Building on this, we developed HHGS⁺, a scalable extension integrating multiple HHGS instances, expanding the signing capacity to $O(2^{60})$ while capable of maintaining compact signature sizes below 8 kilobytes. Through security proofs in the standard model, both HHGS and HHGS⁺ demonstrate strong guarantees of anonymity, unforgeability, and forward security against adaptive adversaries.

¹¹ Although the newer signature scheme [33] provides shorter signatures and more efficient verification, this work adopts WOTS+C for its superior signing efficiency, making it better suited to computation-constrained signers nor powerful verifier.

¹² The source code of our implementation is available at <https://github.com/WithoutNi/GS>.

References

1. Aviram, N., Gellert, K., Jager, T.: Session resumption protocols and efficient forward security for TLS 1.3 0-rtt. *J. Cryptol.* **34**(3), 20 (2021)
2. Bansarkhani, R.E., Misoczki, R.: G-merkle: A hash-based group signature scheme from standard assumptions. In: *PQCrypto*. pp. 441–463. Springer (2018)
3. Barker, E.B., Kelsey, J.M.: Recommendation for random number generation using deterministic random bit generators (revised). US Department of Commerce, Technology Administration, NIST (2007)
4. Bellare, M., Micciancio, D., Warinschi, B.: Foundations of group signatures: formal definitions, simplified requirements, and a construction based on general assumptions. In: *EUROCRYPT*. pp. 614–629. Springer (2003)
5. Bellare, M., Shi, H., Zhang, C.: Foundations of group signatures: The case of dynamic groups. In: *CT-RSA*. pp. 136–153. Springer (2005)
6. Bernstein, D.J., Hopwood, D., Hülsing, A., Lange, T., Niederhagen, R., Papachristodoulou, L., Schneider, M., Schwabe, P., Wilcox-O’Hearn, Z.: SPHINCS: practical stateless hash-based signatures. In: *EUROCRYPT*. pp. 368–397. Springer (2015)
7. Bernstein, D.J., Hülsing, A., Kölbl, S., Niederhagen, R., Rijneveld, J., Schwabe, P.: The sphincs⁺ signature framework. In: *CCS*. pp. 2129–2146. ACM (2019)
8. Bertsekas, D., Tsitsiklis, J.N.: Introduction to probability, vol. 1. Athena Scientific (2008)
9. Beullens, W., Dobson, S., Katsumata, S., Lai, Y., Pintore, F.: Group signatures and more from isogenies and lattices: Generic, simple, and efficient. In: *EUROCRYPT*. pp. 95–126. Springer (2022)
10. Bobolz, J., Diaz, J., Kohlweiss, M.: Foundations of anonymous signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In: *Financial Crypto*. pp. 121–139. Springer (2024)
11. Boneh, D., Eskandarian, S., Fisch, B.: Post-quantum EPID signatures from symmetric primitives. In: *CT-RSA*. pp. 251–271. Springer (2019)
12. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: *ASIACRYPT*. pp. 280–300. Springer (2013)
13. Bootle, J., Cerulli, A., Chaidos, P., Ghadafi, E., Groth, J.: Foundations of fully dynamic group signatures. *J. Cryptol.* **33**(4), 1822–1870 (2020)
14. Brassard, G., Høyer, P., Tapp, A.: Quantum cryptanalysis of hash and claw-free functions. In: *LATIN*. pp. 163–169. LNCS, Springer (1998)
15. Brickell, E.F., Camenisch, J., Chen, L.: Direct anonymous attestation. In: *CCS*. pp. 132–145. ACM (2004)
16. Bringer, J., Chabanne, H., Pointcheval, D., Zimmer, S.: An application of the boneh and shacham group signature scheme to biometric authentication. In: *IWSEC*. pp. 219–230. Springer (2008)
17. Buchmann, J., Dahmen, E., Schneider, M.: Merkle tree traversal revisited. In: *PQCrypto*. pp. 63–78. Springer (2008)
18. Buser, M., Liu, J.K., Steinfeld, R., Sakzad, A., Sun, S.: DGM: A dynamic and revocable group merkle signature. In: *ESORICS*. pp. 194–214. Springer (2019)
19. Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. *J. ACM* **51**(4), 557–594 (2004)
20. Chen, L., Dong, C., Newton, C.J.P., Wang, Y.: Sphinx-in-the-head: Group signatures from symmetric primitives. *ACM Trans. Priv. Secur.* **27**(1), 11:1–11:35 (2024)
21. Cheval, V., Cremers, C., Dax, A., Hirschi, L., Jacomme, C., Kremer, S.: Hash gone bad: Automated discovery of protocol attacks that exploit hash function weaknesses. In: *USENIX Security Symposium*. pp. 5899–5916. USENIX Association (2023)
22. Dagdelen, Ö., Fischlin, M., Gagliardoni, T.: The fiat-shamir transformation in a quantum world. In: *ASIACRYPT*. pp. 62–81. Springer (2013)
23. Dworkin, M.J.: Sha-3 standard: Permutation-based hash and extendable-output functions (2015)
24. Fadavi, M., Karati, S., Erfanian, A., Safavi-Naini, R.: DGMT: A fully dynamic group signature from symmetric-key primitives. *Cryptogr.* **9**(1), 12 (2025)
25. Faonio, A., Fiore, D., Nizzardo, L., Soriente, C.: Subversion-resilient enhanced privacy ID. In: *CT-RSA*. pp. 562–588. Springer (2022)
26. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. *J. ACM* **33**(4), 792–807 (1986)
27. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: *STOC*. pp. 212–219. ACM (1996)
28. Gueron, S.: Aes-gcm-siv implementations (128 and 256 bit) (2018), <https://github.com/Shay-Gueron/AES-GCM-SIV>
29. Hülsing, A., Butin, D., Gazdag, S., Rijneveld, J., Mohaisen, A.: XMSS: extended merkle signature scheme. *RFC* **8391**, 1–74 (2018)
30. Hülsing, A., Kudinov, M.A., Ronen, E., Yogev, E.: SPHINCS+C: compressing SPHINCS+ with (almost) no cost. In: *SP*. pp. 1435–1453. IEEE (2023)
31. Hülsing, A., Rijneveld, J., Song, F.: Mitigating multi-target attacks in hash-based signatures. In: *PKC*. pp. 387–416. LNCS, Springer (2016)

32. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge proofs from secure multiparty computation. *SIAM J. Comput.* **39**(3), 1121–1152 (2009)
33. Khovratovich, D., Kudinov, M., Wagner, B.: At the top of the hypercube - better size-time tradeoffs for hash-based signatures. In: Tauman Kalai, Y., Kamara, S.F. (eds.) *CRYPTO 2025*. pp. 93–123. Springer Nature Switzerland (2025)
34. Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: *CCS*. pp. 669–684. ACM (2013)
35. Kiayias, A., Yung, M.: Extracting group signatures from traitor tracing schemes. In: *EUROCRYPT*. pp. 630–648. Springer (2003)
36. Ling, S., Nguyen, K., Wang, H., Xu, Y.: Forward-secure group signatures from lattices. In: *PQCrypto*. pp. 44–64. Springer (2019)
37. Lyubashevsky, V., Nguyen, N.K., Plançon, M., Seiler, G.: Shorter lattice-based group signatures via “almost free” encryption and other optimizations. In: *ASIACRYPT*. pp. 218–248. Springer (2021)
38. Nguyen, K., Tang, H., Wang, H., Zeng, N.: New code-based privacy-preserving cryptographic constructions. In: *ASIACRYPT*. pp. 25–55. Springer (2019)
39. Omar, S., Padhye, S.: Multivariate linkable group signature scheme. In: *Proceedings of the International Conference on Computing and Communication Systems*. pp. 623–632. Springer, Singapore (2021)
40. Ronen, E.: Software for the sphincs+c scheme (2022), <https://github.com/eyalr0/sphincsplusc>
41. Song, D.X.: Practical forward secure group signature schemes. In: *CCS*. pp. 225–234. ACM (2001)
42. Yang, G., Tang, S., Yang, L.: A novel group signature scheme based on MPKC. In: *ISPEC*. pp. 181–195. Springer (2011)
43. Yang, R., Au, M.H., Zhang, Z., Xu, Q., Yu, Z., Whyte, W.: Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In: *CRYPTO*. pp. 147–175. Springer (2019)
44. Yehia, M., AlTawy, R., Gulliver, T.A.: Gm^{mt} : A revocable group merkle multi-tree signature scheme. In: *CANS*. pp. 136–157. Springer (2021)

A Tree-based PPRF

In this section, we briefly revisit the key concept of tree-based PPRFs [12,34,1], built from the well-known one-way function based PRF construction, proposed by Goldreich, Goldwasser, and Micali (GGM) [26].

The GGM construction leverages on a pseudorandom generator (PRG) $G : \{0, 1\}^\ell \rightarrow \{0, 1\}^{2\ell}$. The output string is then split into two equal parts, denoted as $G_0(k)$ and $G_1(k)$, representing the first and second halves, respectively. To define the PRF $F : \{0, 1\}^\ell \times \{0, 1\}^n \rightarrow \{0, 1\}^\ell$, the GGM construction organizes its domain as a binary tree. The PRF key k , generated by $F.Setup(1^\kappa)$, serves as the root seed. Each leaf corresponds to a unique PRF output, while edges and internal nodes facilitate computation. Edges are labeled 0 or 1, indicating connections to left or right children, respectively, and each node is identified by the binary string of edge labels along the path from the root. For a message $x = x_1x_2 \dots x_n \in \{0, 1\}^n$, the PRF value is computed by traversing the tree from the root to the leaf specified by x . Starting with the root seed k , the PRG G is iteratively applied along the path dictated by the bits of x . At each step, G_{x_i} (either G_0 or G_1) is applied to the current intermediate value. This process yields the final output:

$$F.Eval(k, x) = G_{x_n} \circ G_{x_{n-1}} \circ \dots \circ G_{x_1}(k) \in \{0, 1\}^\ell.$$

The hierarchical structure of this tree-based construction ensures pseudo-randomness and allows for efficient operations. Moreover, its flexibility enables the introduction of puncturing mechanisms. To puncture at $x = x_1x_2 \dots x_n$, realizing $PPRF.Punc(k, x)$, the initial root key k is replaced with intermediate node evaluations for prefixes $\bar{x}_1, \bar{x}_1\bar{x}_2, \bar{x}_1\bar{x}_2\bar{x}_3, \dots, \bar{x}_1\bar{x}_2 \dots \bar{x}_n$. That is, these values are taken as the new key k' output by $F.Punc$ algorithm, allowing recomputation of the PRF for all inputs except x .

For implement the **RandSelect** function in our **FSDGS** constructions, we need an additional algorithm of PPRF, i.e., $GetMsg(k, i)$ to enable the retrieval of the i -th unpunctured message. To implement the $F.GetMsg(k', i)$ algorithm based on the GGM construction, we devise the following steps:

1. Initialize $currentPrefix = \emptyset$.
2. For each level $j = 1$ to n :
 - Compute the number ν of accessible leaves in the left subtree:
 - Count the number μ of punctured message according to the prefixes with $currentPrefix||0$;

- Compute $\nu := 2^{n-j} - \mu$.
 - If $i \leq \nu$, append $0 \rightarrow \text{currentPrefix}$. Otherwise, append $1 \rightarrow \text{currentPrefix}$ and sets $i := i - \nu$.
3. Return currentPrefix as the binary representation of the i -th accessible input message.

The complexity of F.GetMsg primarily depends on the depth of the binary tree, which is determined by the bit-length n of the key, and the number of stored intermediate evaluations, ϕ , corresponding to the punctured points. The algorithm retrieves the i -th accessible input message by iterating through the nodes in the path from the root to the leaf representing the i -th unpunctured message. For each level in the tree, the algorithm verifies whether the current node or path is part of the stored intermediate evaluations. This lookup is $O(m)$ for each level since we may need to search among ϕ intermediate evaluations. Since there are n levels in the tree, the overall complexity of the algorithm is $O(n \cdot \phi)$. This complexity reflects the worst-case scenario where all ϕ punctured points must be checked at each level.

B Correctness of HHGS

According to the correctness experiment $\text{Exp}_{\mathcal{A}, \text{HHGS}}^{\text{Correct}}$, the correctness of HHGS encompasses two key aspects: (i) the adversary can register honest parties under chosen identities, and all signatures generated by honest parties must pass verification; and (ii) signatures produced by honest parties must be accurately opened (traced) to their legitimate owners.

The registration process in HHGS primarily involves locating a mount point on the mixed OTS mount subtree (MoMSt). Since each OTS is uniquely associated with a leaf node, successful registration is guaranteed as long as there are available vacancies in the tree structure. The validity of an honest party's signature during verification is inherently ensured by the correctness of OTS and MT. Moreover, the oracle queries in the correctness game are simulated using secret keys (including sub-secret keys of the building blocks, such as AE, PPRF, and PRF) generated by the challenger. The correctness of these building blocks ensures that all queries and protocol executions involving them are processed correctly in accordance with the protocol's specifications.

C Security Proofs of HHGS

C.1 Proof of Theorem 1

We denote by $\text{Win}_i^{\text{Trace}}$ the event where the i -th (modified) Trace experiment (game) outputs 1.

Game 0. This game corresponds to the original Trace game. Specifically,

$$\Pr[\text{Win}_0^{\text{Trace}}] = \text{Adv}_{\mathcal{A}, \text{HHGS}}^{\text{Trace}}(\kappa, \text{SP}).$$

Game 1. This game proceeds as before, but the challenger aborts if the adversary \mathcal{A} outputs a forgery σ^* containing a value $\text{opk}_{\text{ID}}^* || c_{\text{ID}}^*$ such that: i) $\text{opk}_{\text{ID}}^* || c_{\text{ID}}^*$ is not generated by the challenger during the simulations of join protocol instances (in handling AddHM and AddMM queries); and ii) $\text{H}(\text{opk}_{\text{ID}}^* || c_{\text{ID}}^*) = \text{H}(\text{opk}_{\text{ID}}^* || c_{\text{ID}}^*)$, where $\text{opk}_{\text{ID}}^* || c_{\text{ID}}^*$ is generated by the challenger within the game. That is, \mathcal{A} succeeds in finding a collision of H . It is not hard to see that such a forgery can be used to break the collision-resistance of H . Thus, we have that

$$\Pr[\text{Win}_0^{\text{Trace}}] \leq \Pr[\text{Win}_1^{\text{Trace}}] + \text{Adv}_{\mathcal{A}, \text{H}}^{\text{CRH}}(\kappa).$$

Game 2. This game is modified to evaluate the security of the Merkle tree scheme. Specifically, the challenger \mathcal{C} aborts if the adversary \mathcal{A} produces a forgery σ^* containing a Merkle proof for a subtree (originating either from the OTS hypertree or the mixed OTS mount subtrees) that was not generated by \mathcal{C} during the game, yet the Merkle proof successfully yields a Merkle root generated by \mathcal{C} . It is worth noting that there exists at least one honest Merkle root (i.e., gpk). If such an event occurs with non-negligible probability, it indicates a breach of the Merkle tree's unforgeability. Notably, all Merkle trees associated with the OTS hypertree and the mixed OTS mount subtrees are known to the challenger at the time of their creation. To construct an algorithm \mathcal{B} capable of breaking the security of M by leveraging \mathcal{A} , \mathcal{B} must correctly guess which subtree \mathcal{A} intends to forge. The probability of making a correct guess is approximately $\frac{1}{TN}$, where $TN = \frac{1-2^{\frac{h}{d}}}{1-2^{\frac{h}{d}}}$.

represents the total number of subtrees within each HHGS instance and $\tilde{h} = h + \frac{h}{d}$ is the total height of the hypertree structure. Importantly, this guess is part of \mathcal{B} 's internal setup and does not influence the subsequent progression of the games. If \mathcal{B} successfully identifies the target subtree, it can act as the challenger for \mathcal{A} , simulating the game by preparing the tree's leaves (either from \mathcal{A} 's AddMM queries or generated by \mathcal{B} itself) while interacting with the MT-challenger. Should \mathcal{A} successfully forge a Merkle proof, \mathcal{B} can replicate this success to break the MT with a success probability $\frac{1}{TN}$. Thus, we have:

$$\Pr[\text{Win}_1^{\text{Trace}}] \leq \Pr[\text{Win}_2^{\text{Trace}}] + TN \cdot \text{Adv}_{\mathcal{A}, \mathcal{M}}^{\text{UF}}(\kappa).$$

If \mathcal{C} does not abort in this game, the HHGS instance must contain at least one honest subtree with all leaves (i.e., OTS keys) generated by the challenger.

Game 3. In this game, the challenger \mathcal{C} aborts if it fails to correctly guess the index of the OTS key contained in the forgery σ^* , which resides at the lowest level of the honest subtree. Given that there are at most $2^{\tilde{h}}$ leaves, the probability of making a correct guess is bounded by

$$\Pr[\text{Win}_2^{\text{Trace}}] = 2^{\tilde{h}} \cdot \Pr[\text{Win}_3^{\text{Trace}}].$$

The adversary might attempt to forge the signature of the guessed OTS key or compromise the integrity of the Identity-OPK ciphertext associated. In the subsequent games, several preparations are necessary to reduce the analysis to the security of these primitives, such as replacing their keys with random values.

Game 4. In this game, the challenger \mathcal{C} aborts if it fails to correctly guess any forgery cases (when such cases exist): i) **ACase 1**—The adversary \mathcal{A} forges the signature of a OTS key and determines which OTS key was forged. ii) **ACase 2**—The adversary \mathcal{A} forges the Identity-OPK ciphertext and determines which ciphertext was forged. Since the leaf of the forged target (either a OTS or a ciphertext) is already assumed to be known by the challenger (based on the correct guess in the previous game), and each leaf corresponds to at most two OTS keys and two Identity-OPK key ciphertexts (within the obfuscated OTS layer), the probability of making a correct guess is bounded by:

$$\Pr[\text{Win}_3^{\text{Trace}}] = 4 \cdot \Pr[\text{Win}_4^{\text{Trace}}].$$

Game 5. The challenger \mathcal{C} changes this game from the previous game according to the guessed attacking cases in the previous games.

- For **ACase 1** and the guessed OTS key is from GM, \mathcal{C} replace the randomness re (output by a PPRF function) for generating the OTS key with a random value.
- For **ACase 1** and the guessed OTS key is from an honest member, \mathcal{C} replaces the randomness (generated by the PRG scheme) for generating the OTS key with a random value. Before doing this, \mathcal{C} would play a series hybrid games to replace the seeds of the PRG which are ancestors of the target one with random values as well. There are at most $2^{\tilde{h}}$ nodes to replace, if the HHGS instance consists of only one party.
- For **ACase 2**, \mathcal{C} changes the PRF values for generating the encryption keys and encryption randomness of the target ciphertext with random values in hybrid games.

If there exists an adversary which can distinguish this game from the previous game with non-negligible probability, then it can be used to break one of the primitives in $\{\text{PPRF}, \text{PRF}, \text{PRG}\}$ determined by the guessed attacking cases. Considering the potential number of hybrid games, we have that

$$\Pr[\text{Win}_4^{\text{Trace}}] \leq \Pr[\text{Win}_5^{\text{Trace}}] + 2 \cdot \text{Adv}_{\mathcal{A}, \mathcal{F}_1}^{\text{IND-CMA}}(\kappa, 2) + 2^{\tilde{h}} \cdot \text{Adv}_{\mathcal{A}, \mathcal{G}}^{\text{IND}}(\kappa)$$

Game 6. Now we can reduce the security of HHGS to that of OTS. Specifically, the challenger \mathcal{C} aborts if the adversary outputs a valid OTS signature $\bar{\sigma}^*$ in **ACase 1**. If this abort event occurs with non-negligible probability, we build \mathcal{B} making use of \mathcal{A} to break the OTS's security. In the reduction, all other queries of \mathcal{A} can be honestly simulated by \mathcal{B} with her own secrets, whereas the signature of the target OTS key being attacked can be obtained from the OTS challenger when necessary. Thus, we have that

$$\Pr[\text{Win}_5^{\text{Trace}}] \leq \Pr[\text{Win}_7^{\text{Trace}}] + \text{Adv}_{\mathcal{A}, \mathcal{S}}^{\text{EUF-CMA}}(\kappa, 1).$$

Game 7. In this final game, the challenger \mathcal{C} aborts if the adversary \mathcal{A} successfully forges a signature containing a valid Identity-OPK key ciphertext that was not generated by \mathcal{C} before the corruption of the corresponding owner. Based on the modifications introduced in the previous game, the encryption key and randomness used for the target ciphertext in **ACase 2** are sampled independently as random values. This setup allows the security to be reduced to the integrity of ciphertexts under **E**. Additionally, the honest ciphertext utilizing the targeted key can be simulated through the encryption oracle provided by the AE challenger. Consequently, we establish the following bound:

$$\Pr[\text{Win}_7^{\text{Trace}}] \leq \Pr[\text{Win}_7^{\text{Trace}}] + \text{Adv}_{\mathcal{A}, \mathbf{E}}^{\text{CT-INT}}(\kappa).$$

As the adversary is unable to manipulate any components in the signature of an uncorrupted party (whether a member or GM), the adversary's advantage in the final game is reduced to 0, thereby concluding the proof.

C.2 Proof of Theorem 2

Let $\text{Win}_i^{\text{Anony}}$ denote the event that the **Anony** experiment (game) outputs 1 in the i -th game.

Game 0. The original game, G_0 , corresponds to the real **Anony** experiment (game). We have:

$$\Pr[\text{Win}_0^{\text{Anony}}] = \frac{1}{2} + \text{Adv}_{\mathcal{A}, \text{HHGS}}^{\text{Anony}}(\kappa, \text{SP}).$$

Game 1. In this game, the challenger \mathcal{C} introduces a new abort condition: it aborts if it fails to guess which two honest group members' OTS keys (i.e., $(\text{osk}_{\text{ID}_0^*}$ and $\text{osk}_{\text{ID}_1^*})$ in the member sub-layer of the obfuscated OTS layer) are involved in the $\mathcal{O}_{\text{PriTest}}(\text{ID}_0^*, \text{ID}_1^*, m)$ query. Since there are at most U group members, and E OTS keys in total, such that $U \cdot E = 2^{\tilde{h}}$. The probability of a correct guess is bounded by $\frac{1}{2^{\tilde{h}+2}}$. Consequently, we have:

$$\Pr[\text{Win}_0^{\text{Anony}}] = 2^{\tilde{h}+2} \cdot \Pr[\text{Win}_1^{\text{Anony}}].$$

Note that the correct guess of \mathcal{C} implies that it knows the indices of the target OTS keys on the corresponding tree, where the target OTS keys refer to the guessed OTS keys $(\text{osk}_{\text{ID}_0^*}, \text{osk}_{\text{ID}_1^*})$ and the OTS keys $(\text{osk}'_{\text{ID}_0^*}, \text{osk}'_{\text{ID}_1^*})$ used to sign them in the obfuscated OTS layers associated with the guessed OTS keys of the two honest members. These target OTS keys would associate with four Identity-OPK ciphertexts, referred to as $(c_{\text{ID}_0^*}, c_{\text{ID}_1^*}, c'_{\text{ID}_0^*}, c'_{\text{ID}_1^*})$, respectively.

Game 2. The challenger \mathcal{C} proceeds as in the previous game, but it replaces the random selection key k_s , and all the PRF values for generating the target OTS keys and the corresponding IoCs with random values. We may play similar hybrid games to replace them one-by one analogously in the Game 5 of the proof of Theorem 1. Due to the security of PRG and PRF, we have that

$$\Pr[\text{Win}_1^{\text{Anony}}] \leq \Pr[\text{Win}_2^{\text{Anony}}] + 3 \cdot \text{Adv}_{\mathcal{A}, \text{F}_1}^{\text{IND-CMA}}(\kappa, 3) + 2^{\tilde{h}} \cdot \text{Adv}_{\mathcal{A}, \mathbf{G}}^{\text{IND}}(\kappa).$$

As a result, the OTS keys of GM and uncorrupted members are generated with the same distributions defined by those key generation random values. Hence, the adversary cannot determine between the difference between the GM's OTS key and the member's OTS key. The OTS mounting scenarios of challenged parties' OTS keys (which are uncorrupted) in the obfuscated layer are hidden from the adversary.

Game 3. The game proceeds identically to the previous one, except that the challenge always uses $\hat{\text{ID}}_0^*$ to generate the target ciphertexts in the privacy test query $\mathcal{O}_{\text{PriTest}}(\text{ID}_0^*, \text{ID}_1^*, m)$, regardless of the challenge bit b .

To transition between these modifications, we employ hybrid games that sequentially modify the target ciphertexts one-by-one in each subgame. If an adversary \mathcal{A} exists that can distinguish between two adjacent games, it can be leveraged to construct an algorithm \mathcal{B} capable of breaking the security of **E**. Specifically, \mathcal{B} queries the AE challenger with identities (ID_0^*, X) to simulate the modified target ciphertext, where $X = \text{GM}$ for ciphertexts intended for GM, and $X = \text{ID}_1^*$ otherwise. All other Identity-OPK ciphertexts are generated either by invoking the encryption oracle provided by the AE challenger or through secret keys produced directly by \mathcal{B} itself. Consequently, we derive:

$$\Pr[\text{Win}_2^{\text{Anony}}] \leq \Pr[\text{Win}_3^{\text{Anony}}] + 4 \cdot \text{Adv}_{\mathcal{A}, \mathbf{E}}^{\text{IND-CCA}}(\kappa, 4).$$

This game modification ensures that the adversary cannot deduce any information about the identity from the IOCs.

Game 4. This game proceeds as before but the challenger randomly chooses an index for each target OTS key from the remaining leaves that can mount without using the output of the PRF F_2 . Due to the security of F_2 , we have that

$$\Pr[\text{Win}_3^{\text{Anony}}] \leq \Pr[\text{Win}_4^{\text{Anony}}] + 2 \cdot \text{Adv}_{\mathcal{A}, F_2}^{\text{IND-CMA}}(\kappa, 2).$$

In the worst-case scenario, which is most advantageous for the adversary, only the two indices of the target OTS keys remain unknown. All other indices of honest OTS keys can be obtained from HMSign queries. Consequently, two unrevealed indices remain—one belonging to ID_0^* and the other to ID_1^* . The attacker's goal is to determine which of these two indices belongs to which party, after obtaining the challenge signature. Despite differing probabilities in earlier selections, the final two unrevealed indices are equally likely to belong to either entity. Since the selection process is independent and random, the attacker has no additional information to distinguish ownership. Therefore, the probability of correctly identifying the owner of the given element is $1/2$. This is because the attacker is left with only two remaining choices, one for each entity. Namely, the adversary gains no advantage in this game.

Putting together the advantages in all the above games can only result in negligible advantage, which proves Theorem 2.

D Pseudo codes of HHGS

Here, we present the pseudocodes of the main algorithms of HHGS.

Algorithm 1: GMin(1^κ, SP): Initialization algorithm for the group manager.

Input: Security parameter 1^κ, setup parameter SP = (h, d).
Output: System parameters Pms, group secret key sk_{GM}, and group public key gpk.

- 1 Initialize PPRF schemes: $\{pms_{F_p^i} := F_p^i.\text{Setup}(1^\kappa)\}_{0 \leq i \leq d}$. Initialize OTS scheme: $pms_S := S.\text{Setup}(1^\kappa)$.
Initialize MT scheme: $pms_M := M.\text{Setup}(1^\kappa)$. Initialize AE scheme: $pms_E := E.\text{Setup}(1^\kappa)$. Initialize PRG scheme: $pms_G := G.\text{Setup}(1^\kappa)$. Initialize CRH scheme: $(hk, pms_H) := H.\text{Setup}(1^\kappa)$. Initialize PRF schemes: $\{pms_{F_i} := F_i.\text{Setup}(1^\kappa)\}_{i \in [2]}$. Set system parameters:
 $Pms := (\{pms_{F_p^i}\}_{0 \leq i \leq d}, pms_S, pms_M, pms_E, pms_G, pms_H, pms_{F_1}, pms_{F_2}, hk)$.
- 2 **for** $i := 0$ **to** d **do**
- 3 $k_p^i := F_p^i.\text{KGen}(pms_{F_p^i})$.
- 4 **end for**
- 5 Generate random secret key for F_1 : $k_{GM} := F_1.\text{KGen}(pms_{F_1})$.
- 6 Set group secret key: $sk_{GM} := (\{k_p^i\}_{0 \leq i \leq d}, k_{GM})$.
- 7 **for** $v \in [2^{h'}]$ **do**
- 8 $rs_{GM}^{d,1,v} := F_p^d(k_p^d, v)$. $(osk_{GM}^{d,1,v}, opk_{GM}^{d,1,v}) := S.\text{KGen}(rs_{GM}^{d,1,v})$. $lf_{d,1,v} := opk_{GM}^{d,1,v}$.
- 9 **end for**
- 10 Build the top-most tree: $(Tr_{d,1}, st_{BDS}^{d,1}) := M.\text{Build}(\{lf_{d,1,v}\}_{v \in [2^{h'}]})$.
- 11 Set group public key: $gpk := Tr_{d,1}.\text{Rt}$.
- 12 Set registered identity list: $RegIDL := \emptyset$. Set revocation list: $RL := \emptyset$. Set group management state:
 $st_{GM} := (RegIDL, RL)$.

Algorithm 2: MRegGen(sk_{GM}): Generate registration file of OTS keys

Input: Member secret key sk_{GM}, number of OTS keys ℓ_r .
Output: Registration file RF_{ID}

// Initialize PRG seed from the member's secret key

- 2 $sd_{ID}^0 := sd_{ID}$ extracted from sk_{GM};
- 3 **// Generate** ℓ_r **OTS key pairs**
- 4 **for** $v := 1$ **to** ℓ_r **do**
- 5 $(sd_{ID}^v, rs_{ID}^v) := G.\text{Eval}(sd_{ID}^{v-1})$; $(rs_{ID}^{v,1}, rs_{ID}^{v,2}) := G.\text{Eval}(rs_{ID}^v)$; $(osk_{ID}^v, opk_{ID}^v) := S.\text{KGen}(rs_{ID}^{v,1})$;
- 6 **end for**
- 7 **// Assemble the registration file and update the SK**
- 8 $RF_{ID} := \{opk_{ID}^v\}_{v \in [\ell_r]}$;
- 9 $sk_{ID} := sd_{ID}^{\ell_r}$;

Algorithm 3: Join($sk_{GM}, st_{GM}, ID, RF_{ID}$): Join protocol for the group manager.**Input:** Group secret key sk_{GM} , group management state st_{GM} , identity ID , registration file RF_{ID} .**Output:** Registration result RR_{ID} and updated states st_{GM}, st_{ID} .

```

1 // Process each OTS public key in the registration file
2 foreach  $opk_{ID} \in RF_{ID}$  do
3    $MI := \text{RandSelect}(sk_{GM}, ID, opk_{ID})$ ; Decompose  $MI$  into  $sTI := MI(h)$  and  $sMI$ ; Add  $(ID, opk_{ID}^{MI}, MI)$  to
    $RFCa_{MMT}^{sTI}$ ;
4   // Determine if the sub-tree needs initialization or can be reused
5   if  $F_p^1(k_p^1, sTI) \neq \perp$  then
6     // Case 1: Sub-tree initialization with new OTS leaves
7     Prepare  $L_{mI}$  and  $\bar{L}_{mI}$  for  $MMT[sTI]$ ;
8     foreach  $sMI \in L_{mI}$  do
9        $ke_{MI}^1 := F_1(k_{GM}, \text{"K-Mix"} || MI)$ ;  $re_{MI}^1 := F_1(k_{GM}, \text{"R-Mix"} || MI)$ ;  $c_{MI}^1 := E.\text{Enc}(ke_{MI}^1, ID; re_{MI}^1)$ ;
        $lf_{sMI} := H(opk_{ID}^{MI} || c_{MI}^1)$ ; Set  $c_{MI}^2 := \emptyset$ ,  $opk_{GM}^{MI} := \emptyset$ ,  $\sigma_{GM}^{sMI} := \emptyset$ ;  $k_p^{0'} := F_p^0(k_p^0, MI)$ ;
10    end foreach
11    foreach  $sMI' \in \bar{L}_{mI}$  do
12       $MI' := sTI || sMI'$ ;  $(osk_{GM}^{MI'}, opk_{GM}^{MI'}) := S.\text{KGen}(rs_{GM}^{MI'})$ , where  $rs_{GM}^{MI'} := F_p^0(k_p^0, MI')$ ;
       $ke_{MI'}^1 := F_1(k_{GM}, \text{"K-Mix"} || MI')$ ;  $re_{MI'}^1 := F_1(k_{GM}, \text{"R-Mix"} || MI')$ ;
       $c_{MI'}^1 := E.\text{Enc}(ke_{MI'}^1, GM; re_{MI'}^1)$ ;
13    end foreach
14    // Build and authenticate the new Merkle sub-tree
15     $(MMT[sTI], st_{BDS}^{MMT[sTI]}) := M.\text{Build}(\{lf_v\}_{v \in [2^{|sMI|}]})$ ;
     $\text{Auth}[MMT[sTI].Rt] := \text{GetAuth}(\{k_p^i\}_{i \in [d]}, MMT[sTI].Rt, sTI)$ ; foreach  $sMI \in L_{mI}$  do
16       $pf_{lf_{sMI}} := M.\text{GetPrf}(st_{BDS}^{MMT[sTI]}, lf_{sMI})$ ; for  $i := 1$  to  $d$  do
17         $w := \frac{(d+1-i) \cdot h}{d}$ ;  $k_p^{i'} := F_p^i(k_p^i, MI(w))$ ;
18      end for
19    end foreach
20    Store  $st_{BDS}^{lf_{sMI'}}$  and  $\text{Auth}[MMT[sTI].Rt]$  in  $st_{GM}$ ;
21  end if
22  else
23    // Case 2: Reuse existing sub-tree and authenticate member OTS key
24     $(osk_{GM}^{MI}, opk_{GM}^{MI}) := S.\text{KGen}(rs_{GM}^{MI})$ , where  $rs_{GM}^{MI} := F_p^0(k_p^0, MI)$ ;  $ke_{MI}^1 := F_1(k_{GM}, \text{"K-Mix"} || MI)$ ;
     $re_{MI}^1 := F_1(k_{GM}, \text{"R-Mix"} || MI)$ ;  $c_{MI}^1 := E.\text{Enc}(ke_{MI}^1, GM; re_{MI}^1)$ ;  $lf_{sMI} := H(opk_{GM}^{MI} || c_{MI}^1)$ ;
     $pf_{lf_{sMI}} := M.\text{GetPrf}(st_{BDS}^{lf_{sMI}}, lf_{sMI})$ ; Retrieve  $\text{Auth}[MMT[sTI].Rt]$  from  $st_{GM}$ ;
     $ke_{MI}^2 := F_1(k_{GM}, \text{"K-Mem"} || MI)$ ;  $re_{MI}^2 := F_1(k_{GM}, \text{"R-Mem"} || MI)$ ;  $c_{MI}^2 := E.\text{Enc}(ke_{MI}^2, ID; re_{MI}^2)$ ;
     $\sigma_{GM}^{sMI} := S.\text{Sign}(osk_{GM}^{sMI}, opk_{ID}^{MI} || c_{MI}^2)$ ; Remove  $st_{BDS}^{lf_{sMI}}$  and puncture  $MI$ ;
25  end if
26   $\text{Auth}[opk_{ID}^{MI}] := (opk_{GM}^{MI}, c_{MI}^1, c_{MI}^2, pf_{lf_{sMI}}, \sigma_{GM}^{sMI}, MMT[sTI].Rt, \text{Auth}[MMT[sTI].Rt])$ ;
27  Append  $\text{Auth}[opk_{ID}^{MI}]$  to  $RR_{ID}$ ;
28 end foreach
29 Update  $st_{GM}$  and  $st_{ID}$ ;

```

Algorithm 4: $\text{Sign}(sk_{iD}, st_{iD}, m)$: Signature generation by a member

Input: Member secret key sk_{iD} , member state st_{iD} , message m
Output: Group signature σ and updated states sk'_{iD} , st'_{iD}

```

1 // Generate fresh randomness and keys
2  $(sd_{iD}^v, rs_{iD}^v) := \text{G.Eval}(sd_{iD}^{v-1}); (rs_{iD}^{v,1}, rs_{iD}^{v,2}) := \text{G.Eval}(rs_{iD}^v); (osk_{iD}, opk_{iD}) := \text{S.KGen}(rs_{iD}^{v,1});$ 
3 // Retrieve latest authentication path
4  $\text{Auth}[opk_{iD}] :=$  first entry from  $st_{iD}$ ;
5 // Check which sub-layer the OTS key belongs to
6 if  $\sigma_{GM}^{sMI} \neq \emptyset$  in  $\text{Auth}[opk_{iD}]$  then
7    $\sigma_{iD} := \text{S.Sign}(osk_{iD}, m);$ 
8    $\sigma := (opk_{iD}, \text{Auth}[opk_{iD}], \sigma_{iD});$ 
9 end if
10 else
11    $(osk'_{iD}, opk'_{iD}) := \text{S.KGen}(rs_{iD}^{v,2});$ 
12    $\sigma_{iD} := \text{S.Sign}(osk'_{iD}, m);$ 
13   Generate  $ke' \xleftarrow{\$} \mathcal{K}_{AE}$  and  $re' \xleftarrow{\$} \mathcal{ER}_{AE};$ 
14    $c'_{iD} := \text{E.Enc}(ke', opk'_{iD}; re');$ 
15    $\sigma'_{iD} := \text{S.Sign}(osk_{iD}, opk'_{iD} || c'_{iD});$ 
16   Fill  $\text{Auth}[opk_{iD}]$  with  $opk_{iD}$ ,  $\sigma'_{iD}$ ,  $c'_{iD}$ ;
17    $\sigma := (opk'_{iD}, \text{Auth}[opk_{iD}], \sigma_{iD});$ 
18 end if
19 // Update member state and key
20  $sk'_{iD} := sd_{iD}^v;$ 
21  $st'_{iD} := st_{iD} \setminus (opk_{iD}, \text{Auth}[opk_{iD}]);$ 

```

Algorithm 5: $\text{Verify}(\text{gpk}, m, \sigma, \text{RL})$: Signature verification algorithm

Input: Group public key gpk , message m , signature σ , revocation list RL
Output: 1 if valid, 0 otherwise

```

1 // Parse the input signature and authentication path
2 Parse  $\sigma := (opk'_{iD}, \text{Auth}[opk_{iD}], \sigma_{iD})$ ; Parse
    $\text{Auth}[opk_{iD}] := (opk_{GM}^{MI}, c_{MI}^1, c_{MI}^2, pf_{lf_{sMI}}, \sigma_{GM}^{sMI}, \text{MMT}[\text{sTI}].\text{Rt}, \text{Auth}[\text{MMT}[\text{sTI}].\text{Rt}]);$ 
3 // Check revocation status
4 if  $c_{MI}^1 \in \text{RL}$  then
5   return 0;
6 end if
7 // Verify the authentication path in the hypertree
8 if  $\text{AuthVrfy}(\text{gpk}, \text{MMT}[\text{sTI}].\text{Rt}, \text{Auth}[\text{MMT}[\text{sTI}].\text{Rt}]) = 0$  then
9   return 0;
10 end if
11 // Verify Merkle proof of GM's leaf in subtree
12  $lf_{sMI} := \text{H}(opk_{GM}^{MI} || c_{MI}^1)$ ; if  $\text{M.Verify}(\text{MMT}[\text{sTI}].\text{Rt}, lf_{sMI}, pf_{lf_{sMI}}) = 0$  then
13   return 0;
14 end if
15 // Verify GM's signature on member's OTS key and ciphertext
16 if  $\text{S.Verify}(opk_{GM}^{MI}, opk'_{iD} || c_{MI}^2, \sigma_{GM}^{sMI}) = 0$  then
17   return 0;
18 end if
19 // Verify member's OTS signature on the message
20 if  $\text{S.Verify}(opk'_{iD}, m, \sigma_{iD}) = 0$  then
21   return 0;
22 end if
23 return 1;

```
