

Efficient Blind Signatures and Threshold Signatures from Pairing-Free Groups and Lattices

Chenzhi Zhu

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Stefano Tessaro, Chair

Huijia Lin, Chair

Nirvan Tyagi

Program Authorized to Offer Degree:

Computer Science & Engineering

©Copyright 2025

Chenzhi Zhu

University of Washington

Abstract

Efficient Blind Signatures and Threshold Signatures from Pairing-Free Groups and Lattices

Chenzhi Zhu

Co-Chairs of the Supervisory Committee:

Stefano Tessaro

Department of Computer Science & Engineering

Huijia Lin

Department of Computer Science & Engineering

Blind signatures and threshold signatures are two fundamental cryptographic primitives that have attracted significant real-world interest, largely due to the development of practically efficient schemes. This thesis focuses on designing efficient constructions in pairing-free groups and lattices, along with rigorous security analyses. Pairing-free constructions are attractive because of their compact key/signature sizes and well-established library support, while lattice-based constructions are of particular interest as they are conjectured to remain secure against quantum adversaries. The thesis makes progress in the following three directions:

Pairing-free blind signatures: The thesis proposes the most efficient pairing-free construction that is provably secure in the concurrent setting, under the discrete logarithm (DL) assumption in the algebraic group model (AGM) and the random oracle model (ROM).

Pairing-free threshold signatures: This thesis introduces a new syntax and security hierarchy for analyzing FROST, the state-of-the-art pairing-free threshold signature scheme. It provides the first security proof of FROST in the ROM under the one-more discrete logarithm (OMDL) assumption, and further proposes variants of FROST that is provably secure under the standard DL assumption. Those are the first partially non-interactive constructions based on the DL assumption.

Lattice-based threshold signatures: This thesis develops new techniques that establish the provable security of one of the state-of-the-art two-round lattice-based threshold signature schemes under standard lattice assumptions, in contrast to prior work which relied on a new non-standard assumption. These techniques also yield a new security analysis for another state-of-the-art two-round scheme with a simpler setup, significantly improving its efficiency.

TABLE OF CONTENTS

	Page
List of Figures	iii
Chapter 1: Introduction	1
Chapter 2: Notation	9
Part I: Blind Signatures in Pairing-Free Groups	11
Chapter 3: Blind Signatures: Introduction	12
3.1 A Scheme in the GGM	14
3.2 AGM Security and Partial Blindness	18
3.3 Preliminaries	19
Chapter 4: The Weighted Fractional ROS Problem	22
4.1 Problem Description and Lower Bound	22
4.2 Proof of the Lower Bound for WFROS (Theorem 4.1.1)	23
Chapter 5: Efficient Blind Signatures in the GGM	34
5.1 Construction and Security	34
5.2 Proof of OMUF (Theorem 5.1.2)	38
Chapter 6: Efficient Blind Signatures in the AGM	59
6.1 Construction and Security	59
6.2 Proof of OMUF (Theorem 6.1.2)	62
6.3 Partially Blind Signatures	68
Part II: Threshold Signatures in Pairing-Free Groups	81
Chapter 7: Threshold Signatures: Introduction	82
7.1 Stronger Security for Non-Interactive Threshold Signatures	83
7.2 Partially Non-Interactive Threshold Signatures from Weaker Assumptions	88

Chapter 8:	A Framework for Non-Interactive Threshold Signatures	92
8.1	Syntax and Correctness	92
8.2	Unforgeability and Strong Unforgeability	96
8.3	Relations and Transformations	98
Chapter 9:	Security of the FROST Schemes	112
9.1	The FROST Schemes	112
9.2	Security of FROST2	115
9.3	Security of FROST1	122
9.4	Security of FROST3	129
9.5	Attacks for the FROST Schemes	133
Chapter 10:	Partially Non-Interactive Threshold Signatures from Linear Hash Functions	137
10.1	Linear Hash Functions	137
10.2	Algebraic One-more Preimage Resistance	138
10.3	Instantiations From the Discrete Logarithm Problem	141
10.4	Threshold Signatures from LHF	143
Part III:	Lattice-based Threshold Signatures	156
Chapter 11:	The Algebraic One-More MISIS Problem and Lattice-based Threshold Signatures: Introduction	157
11.1	Technical Overview	162
11.2	Preliminaries	169
Chapter 12:	The Algebraic One-More MISIS Problem	174
12.1	Definition of AOM-MISIS	174
12.2	Comparison with prior work	175
12.3	Reduction from MSIS and MLWE	178
Chapter 13:	Two-Round Lattice-based Threshold Signatures	184
13.1	Analysis of the CATZ Construction	184
13.2	Analysis of the EKT Construction	194
Bibliography	210

LIST OF FIGURES

Figure Number	Page
2.1 The DLog game.	10
3.1 The OMUF security game for a blind signature scheme BS.	20
3.2 The Blind security game for a blind signature scheme BS.	21
4.1 The WFROS problem. Here, $\alpha, \beta \in \mathbb{Z}_p^{2\ell+1}$, which is indexed as $\alpha = (\alpha^{(0)}, \dots, \alpha^{(2\ell)})$ and $\beta = (\beta^{(0)}, \dots, \beta^{(2\ell)})$	23
5.1 The blind signature scheme $BS_1 = BS_1[GGen]$	35
5.2 The OMUF security game in GGM for the blind signature scheme $BS_1[GGen]$	37
5.3 The definition of $Game_4$. The symbols P and P' denote polynomials over variables $X, \{A_i, Y_i\}_{i \in [sid]}$. Also, a new equality notation, “ $=_L$ ”, is used. We say $P_1 =_L P_2$ if and only if $P_1 - P_2$ can be represented as a linear combination of polynomials in L	39
5.4 The definition for $Game_1^A, Game_2^A$, and $Game_2'^A$, where $Game_1^A$ only contains the dashed box, $Game_2^A$ contains all but the gray box, and $Game_2'^A$ contains all but the dashed box.	41
5.5 The definition for $Game_3^A$ and its difference from $Game_2^A$. $Game_2^A$ contains all but the solid boxes and $Game_3^A$ contains all but the dashed boxes. We also define an intermediate game $Game_3'^A$ which contains both dashed and solid boxes.	47
5.6 The definition for $Game_4^A$ and its difference from $Game_3^A$. $Game_3^A$ contains all but the solid box and $Game_4^A$ contains all but the dashed box.	49
6.1 The blind signature scheme $BS_2 = BS_2[GGen]$	60
6.2 The OMUF security game for the blind signature scheme $BS_2[GGen]$	63
6.3 The partially blind signature scheme $PBS = PBS[GGen]$	69
6.4 The PBlind security game for a partially blind signature scheme PBS.	70
6.5 The OMUF security game for the partially blind signature scheme $PBS[GGen]$	71
6.6 The rel-DLog game.	76
8.1 The TS-COR game for a threshold signature scheme TS for n signers and threshold t	93

8.2	The games TS-UF- i and TS-SUF- i of a threshold signature scheme TS. In the game TS-UF- i (resp. TS-SUF- i), the set TF_M (resp. TF_L) records the set of messages (resp. leader requests) that are considered trivial under the trivial forgery predicate tf_i (Figure 8.3). The oracle INIT must be invoked exactly once, and only before any other oracle is queried. In particular, the predicate tf_3 and, thus, TS-UF-3 and TS-SUF-3 are defined only if TS is an echo scheme.	95
8.3	Top: Trivial-forgery conditions $\text{tf}_i(lr)$ ($i = 0, 1, 2, 3, 4$) used to define games TS-UF- i and TS-SUF- i (Figure 8.2). Bottom: Relations between notions of security.	97
8.4	Reference threshold signature schemes $\text{RTS}_\ell[\text{DS}]$ associated to signature scheme DS for $\ell = 1, 2, 3, 4$. The verification and stronger verification algorithms are provided in Figure 8.5. 100	
8.5	The verification and strong verification algorithms of the reference threshold signature schemes $\text{RTS}_\ell[\text{DS}]$ associated to signature scheme DS for $\ell = 1, 2, 3, 4$	101
8.6	The threshold signature $\text{ATS}[\text{TS}, \text{DS}]$ constructed from an echo scheme TS and a digital signature scheme DS such that $\text{ATS}.n = \text{TS}.n$ and $\text{ATS}.t = \text{TS}.t$. The algorithm OriginLR transforms a well-formed leader request lr for ATS to a well-formed leader request in TS. $\text{st}_0.\text{SigMap}$ is a table that stores the signature corresponding to each token generated by honest signers, which is initially set to empty. PS denotes a set of partial signatures.	106
8.7	The game $\text{SUF-CMA}_{\text{DS}}$, where DS is a digital signature scheme.	107
8.8	Adversary \mathcal{B} for the proof of Lemma 8.3.4. \mathcal{B} also compute the sets L, PP, and $\text{curSSL}(lr)$ for each $lr \in L$ following the same logic as in the game $\text{TS-SUF-4}_{\text{ATS}}$ and thus can check whether the event $\text{WIN} \wedge (\neg \text{BadLR})$ occurs.	109
9.1	The common components of $\text{FROST1}[\text{GGen}]$, $\text{FROST2}[\text{GGen}]$ and $\text{FROST3}[\text{GGen}]$, where GGen is a group generation algorithm. The algorithms LR and CompPar are defined in Figure 9.2 Further, n is the number of parties, and t is the threshold of the schemes. PS denotes a set of partial signatures. $\lambda_i^{lr, \text{SS}}$ denotes the Lagrange coefficients, which is defined in Equation (9.1).	113
9.2	The algorithms LR and CompPar of $\text{FROST1}[\text{GGen}]$, $\text{FROST2}[\text{GGen}]$ and $\text{FROST3}[\text{GGen}]$, where GGen is a group generation algorithm. In particular, FROST1 contains all the dashed boxes, FROST2 contains all the highlighted boxes, and FROST3 contains all the solid boxes. The function $\text{H}_i(\cdot)$ is computed as $\text{H}(i, \cdot)$ for $i = 1, 2$. PS denotes a set of partial signatures.	114
9.3	The AOMDL game, where \mathbb{G} is a cyclic group with prime order p and generator g	115
9.4	The forking algorithm build from \mathcal{A} used in Lemma 9.2.3..	117
9.5	The forking algorithm build from \mathcal{A} used in Lemma 9.3.4.	125
9.6	Adversary \mathcal{A} that wins the game $\text{TS-UF-4}_{\text{FROST1}}$, where M is a fixed message.	134
9.7	Adversary \mathcal{A} that wins the game $\text{TS-UF-3}_{\text{FROST2}}$, where M is a fixed message.	135
9.8	Adversary \mathcal{A} that wins the game $\text{TS-UF-2}_{\text{FROST3}}$, where M is a fixed message.	136
10.1	The CR security game for a linear hash family $\text{LHF} = (\text{PGen}, \text{F})$	138

10.2	The AOMPR game for a linear hash function family $\text{LHF} = (\text{PGen}, \text{F})$. For the inputs of PI, X is in \mathcal{R} , α is in \mathcal{D} , and each β_i is in \mathcal{S}	139
10.3	The protocol $\text{FROST1-H}[\text{LHF}]$ and $\text{FROST1-H}[\text{LHF}]$, where $\text{LHF} = (\text{PGen}, \text{F})$ is a linear hash function family. The protocol FROST1-H contains all but the dashed box, and the protocol FROST2-H contains all but the solid box. Further, n is the number of parties, and t is the threshold of the schemes. $\times_{(\cdot)}$ is an injection from $[n]$ to \mathcal{S} and $\lambda_i^{lr,SS}$ denotes the Lagrange coefficient which is computed as $\lambda_i^{lr,SS} := \prod_{j \in \mathcal{S} \setminus \{i\}} \frac{x_j}{x_j - x_i}$. \mathcal{D}_{key} is a subset of \mathcal{D} such that F is a bijection between \mathcal{D}_{key} and \mathcal{S}	144
11.1	The module-SIS and module-LWE problems, where $R := \mathbb{Z}[X]/(X^N + 1)$, $R_q := R/qR$ and $\mathcal{B}_{\beta_{\text{lwe}}}^m := \{\mathbf{x} \in R^m \mid \ \mathbf{x}\ _\infty \leq \beta_{\text{lwe}}\}$	172
12.1	The AOM-MISIS game, where $R := \mathbb{Z}[X]/(X^N + 1)$ and $R_q := R/qR$	175
12.2	The AOM-MLWE game.	176
13.1	Lattice-based t -out-of- n threshold signatures $\text{CATZ}[\text{SecSha}]$, where $\text{SecSha} = \text{SecSha}_{t,n,B_{\text{ss}}}$ is t -out-of- n a linear secret sharing scheme with small coefficients (see Definition 13.1.1). In particular, $\lambda_j^{lr,SS}$ denotes the reconstruction coefficient. Also, $\text{H}_1 : \{0, 1\}^* \rightarrow \mathcal{S}_b^\ell$ and $\text{H}_2 : \{0, 1\}^* \rightarrow \mathcal{S}_c$. The algorithms LPP and LR are defined the same as in Figure 10.3.	186
13.2	Parameters for CATZ and EKT. Some parameters only apply to EKT.	187
13.3	The concrete parameters and estimated efficiency for $\kappa = 128$ and $n = 32$ in [41] and this work. We set $(N, \ell, \beta_c) = (512, 26, 64)$. The last second column denotes the communication complexity per signer.	189
13.4	Lattice-based t -out-of- n threshold signatures $\text{EKT}[\text{PRF}]$, where PRF is a pseudorandom function. Here, $\text{H}_1 : \{0, 1\}^* \rightarrow \mathcal{S}_b^\ell$ and $\text{H}_2 : \{0, 1\}^* \rightarrow \mathcal{S}_c$. Also, $\lambda_{lr,SS,j}$ denotes the Lagrange coefficient, and $\overline{\text{pk}}, \overline{\mathbf{h}} \in R^k$ denote the lift (see Section 11.2.1 for more details) of pk and \mathbf{h} respectively. The algorithms LPP and LR are defined the same as in Figure 10.3.	195
13.5	The concrete parameters and estimated efficiency of the EKT scheme for $\kappa = 128, 192, 256$ and $n = 1024$. We set $(N, \ell, \beta_c) = (512, 26, 64)$. The last second column denotes the communication complexity per signer.	198
13.6	The Ideal-TUF game, where the algorithms CompPar and Vf are defined in Figure 13.4.	199
13.7	The \mathbf{G}_0 game, where the algorithms CompPar and Vf are defined in Figure 13.4.	201
13.8	The PSIGNO oracle of the games $\mathbf{G}_1, \mathbf{G}_2$, and \mathbf{G}_3 , where \mathbf{G}_1 only contains dashed boxes, \mathbf{G}_2 only contains highlighted boxes, and \mathbf{G}_3 only contains solid boxes. In addition, each entry of the tables curSum and curSumR is initialized to 0. The rest of each game is identical to \mathbf{G}_0	202

ACKNOWLEDGMENTS

I would like to express my deepest gratitude to my advisor, Stefano Tessaro, for his invaluable guidance throughout my Ph.D. I am especially grateful for his advice on conducting research, writing papers, and giving talks, all of which have been central to my growth as a researcher.

I would also like to thank my co-advisor, Huijia Lin, for her support and feedback, even though we do not have direct research collaborations. I am grateful to Nirvan Tyagi and Gaku Liu for serving on my Ph.D. exam committee.

I would like to thank Ilan Komargodski, who supervised me during my internship at NTT and opened a door into a new research area, even though that work is not included in this thesis. I am also thankful to Melissa Chase and Esha Ghosh for supervising me during my internship at Microsoft Research. I am further grateful to Vipul Goyal, who hosted me during my undergraduate visit to CMU and first introduced me to research in cryptography.

I would like to thank my co-author Mihir Bellare for the opportunity to collaborate and for our insightful discussions on threshold signatures included in this thesis. I am also grateful to my co-author Rutchathon Chairattana-Apirom (Champ) for many useful discussions on lattice-based constructions and secret sharing schemes. I would also like to thank all my other co-authors for their effort, stimulating conversations, and support in preparing talks.

Special thanks go to the other members of the UW cryptography group—Binyi Chen, Pratik Soni, Ashrujit Ghoshal, Ji Luo, Xihu Zhang, Anna Kornfeld Simpson, Tianren Liu, Hanjun Li, Marshall Ball, Yao-Ching Hsieh, Marian Dietz, and Andrea Coladangelo—for creating a positive and collaborative atmosphere, organizing engaging events, and supporting me through the difficult times of my Ph.D. journey.

I am thankful to the UW CSE advising team for their assistance with countless matters that made my Ph.D. journey much smoother.

Finally, I owe heartfelt thanks to my family and friends for their unwavering support and encouragement.

Chapter 1

INTRODUCTION

Digital signatures are fundamental cryptographic primitives widely used across various applications. They play a crucial role in authenticating transactions and communications. While many efficient signature schemes exist with fast signing and verification, emerging scenarios demand more advanced protocols for generating digital signatures. Among these, two key areas of focus are blind signatures and threshold signatures.

Blind signature protocols, first introduced by David Chaum in 1982 [43], allow users to obtain signatures from a signer without revealing either the messages being signed or the resulting signatures to the signer. More precisely, their security consists of two key properties: *blindness*, which ensures that the issued messages and signatures remain unlinkable to the signing interaction—even if the signer is malicious, and *one-more unforgeability*, which prevents a malicious user from forging more signatures than the number of signing interactions they engage in. Blind signatures are highly versatile and essential for applications requiring anonymity guarantees. Initially proposed for anonymous cash systems [43, 46, 108, 35], they were later adapted for other privacy-preserving applications, such as anonymous e-voting [45, 66] and anonymous credentials [35, 36, 37, 13]. Furthermore, thanks to the existence of efficient constructions [44, 27, 57], blind signatures have been deployed in various real-world applications, including Private Click Measurement [1], Google’s VPN service, and Privacy Pass [52].

Threshold signatures [58, 59] enable the signing key to be distributed among a group of n signers, allowing a valid signature to be issued only if at least a threshold number t of signers are involved. Crucially, any subset of fewer than t signers should be unable to produce a valid signature, even if they collude and deviate from the protocol. The primary motivation for developing practically efficient threshold signatures comes from blockchain systems (e.g., digital wallets [70]), where they are used to mitigate the risk of a single point of failure in key management. Additionally, threshold signatures have recently gained significant attention due to ongoing standardization efforts by NIST [102].

MAIN FOCUSES. The primary focus of this thesis is on constructing efficient cryptographic schemes in pairing-free groups. These constructions treat prime-order cyclic groups as black boxes and do not rely on bilinear pairings. In practice, such groups are typically instantiated using standard elliptic curves. By “efficient,” we refer to constructions that avoid heavy cryptographic tools, such as general-purpose non-interactive zero-knowledge proofs.

While efficient constructions based on pairings or RSA exist for blind signatures [44, 27, 57] and threshold signatures [53, 75, 119, 27], pairing-free constructions are preferable for practical applications for the following reasons. Pairing-based constructions suffer from limited library support, making them difficult to deploy in browser applications. RSA-based constructions are generally less efficient than pairing-free alternatives in terms of signature size, signing time, verification time, and communication complexity. This inefficiency arises because RSA-based schemes require significantly larger modulus sizes to ensure the hardness of the RSA problem.

Another motivation for studying pairing-free constructions is that they pave the way for building efficient post-quantum schemes, which is the second focus of this thesis. As the threat of quantum computers looms on the horizon, there is growing interest in developing cryptosystems based on quantum-resistant assumptions. Lattice-based assumptions are among the most promising candidates. The selection of Dilithium [99] and Falcon [111] by NIST for standardization reinforces this belief. A common approach for designing efficient lattice-based schemes is to adapt efficient pairing-free constructions. For example, Dilithium follows the “Fiat–Shamir with Aborts” paradigm [97], which is based on this approach. This paradigm has also led to practical lattice-based schemes in multi-signatures [34] and threshold signatures [54].

SCHNORR SIGNATURES. Most of pairing-free constructions for blind signatures and threshold signatures are all built around Schnorr signatures [115], which is one of the most efficient digital signature schemes in pairing-free groups. Denote a group as \mathbb{G} with generator g and order p . The signing key of Schnorr signature is a random scalar $sk \in \mathbb{Z}_p$ and the corresponding public key is $pk \leftarrow g^{sk}$. A Schnorr signature for a message m consists of one group element and one scalar $(R, z) \in \mathbb{G} \times \mathbb{Z}_p$ satisfying

$$g^z = R \cdot pk^c, \tag{1.1}$$

with $c = H(pk, R, m)$ and H is a cryptographic hash function. We can also use (c, z) as a Schnorr signature since one can compute c from R and also compute $R (= g^z pk^{-c})$ from c . Then, a Schnorr signature consists of two scalars in \mathbb{Z}_p . Schnorr signatures are now supported by major blockchain

systems, including Bitcoin. Pointcheval and Stern [110] shows the security of Schnorr signatures under the discrete logarithm (DL) assumption in the random oracle model (ROM) [18], where H is modeled as a random oracle.

PAIRING-FREE BLIND SIGNATURES. Chaum and Pedersen [47] proposed a highly efficient three-move blind signature scheme for issuing Schnorr signatures, commonly referred to as blind Schnorr. Here, "three-move" means that each signing session consists of three messages exchanged between the signer and the user, with the signer sending the initial message and the user receiving the final message.

However, Schnorr [117] identified a hard problem, called the ROS problem, in the security analysis of blind Schnorr. Specifically, if an adversary can solve the ROS problem, they can break the one-more unforgeability of blind Schnorr in the concurrent setting. The concurrent setting means that there can be arbitrarily many signing sessions happening concurrently with communication from different sessions interleaved arbitrarily. In contrast, the sequential setting requires the signer to complete one signing session before starting another. Consider that the concurrent setting is crucial in practice, as restricting a scheme to sequential signing makes it vulnerable to denial-of-service (DoS) attacks.

A recent break-through result by Benhamouda et al. [24] shows that there exists a polynomial-time algorithm that solves the ROS problem. Concretely, the attack shows that one can forge $\lceil \log_2(p) \rceil + 1$ valid signatures by engaging in only $\lceil \log_2(p) \rceil$ number of signing sessions, where p denotes the group size. Several attempts [3, 65, 88, 39] have been made to circumvent the ROS problem. However, these solutions either rely on non-standard assumptions [65] or are significantly less efficient compared to blind Schnorr [3, 88, 39]. Therefore, it remains a major open problem to construct an efficient, pairing-free blind signature scheme that is provably secure under standard assumptions.

PAIRING-FREE THRESHOLD SIGNATURES. Currently, the most efficient pairing-free threshold signature scheme is FROST, proposed by Komlo and Goldberg [91]. FROST issues Schnorr signatures and has a two-round signing protocol. More precisely, each signer communicates with a coordinator (which can be one of the signers). In the first round, each signer generates a random token and sends it to the coordinator. In the second round, the coordinator distributes the collected tokens to all signers, who then compute and send partial signatures back to the coordinator. Finally, the coordinator combines them to produce a valid signature.

FROST improves upon prior constructions [120, 72] in terms of both communication and round

complexity at the price of giving up robustness. Additionally, it offers significant usability advantages. The first-round messages are generated independently of both the message being signed and the set of participating signers. This allows the first round to be preprocessed, where the coordinator can collect a batch of tokens in advance. When receiving a signing request for a message m with a designated subset of signers, the coordinator simply selects a token for each designated signer and proceeds with the second-round protocol to generate the final signature.

However, there is no formal security analysis for FROST. The original paper does not provide a formal security definition or proof for the scheme. Moreover, we found that existing syntax and security definitions fail to capture the aforementioned offline-online signing pattern. Prior security definitions, such as the one given in [70], only consider signing protocols where the message to be signed is determined at the start of the execution.

Another issue with prior definitions is that they assume all honest designated signers participate in every signing session. However, in FROST, a malicious coordinator may choose to interact with only a subset of the honest designated signers, and it is unclear whether a signature is considered issued or not in this case. Definitions that account for these more fine-grained signing patterns exist only for non-interactive threshold signatures [119, 124, 29], where the signing protocol consists of a single round of communication between the user and each designated signer. It is unclear how we can define similar security notions for FROST. Also, these definitions are inconsistent on whether a signature is considered issued. The definitions from [124, 29] consider a signature is issued for a message m as long as the adversary receives a partial signature from at least one honest signers, while the definition given by Shoup [119] is stronger, where they consider a signature is issued only if the adversary receives partial signatures from at least $t - k$ honest signers, where k denotes the number of corrupted parties.

LATTICE-BASED CONSTRUCTIONS. For lattice-based constructions, we only focus on threshold signatures in this thesis. The existing lattice-based threshold signatures are not yet satisfactory. Early constructions were based on fully homomorphic encryption (FHE) [29, 8], a heavy cryptographic tool that renders these schemes impractical. A more recent line of work [54, 87, 61, 60] has made significant progress by adapting constructions from pairing-free groups to the lattice setting. In particular, these schemes produce lattice analogs of Schnorr signatures. The high-level idea is to replace the group exponentiation g^x with matrix multiplication $A\mathbf{x}$, where A is a random matrix over \mathbb{Z}_p for some integer p and \mathbf{x} is a small-norm vector in \mathbb{Z}_p . The secret key sk of a lattice-based signature scheme is a random small-norm vector in \mathbb{Z}_p^m , and its corresponding public key is com-

puted as $\text{pk} \leftarrow \text{Ask}$, with $A \in \mathbb{Z}_p^{k \times m}$ being part of the public parameter. A signature of a message μ consists of $(\mathbf{R} \in \mathbb{Z}_p^k, \mathbf{z} \in \mathbb{Z}_p^m)$ satisfying $A\mathbf{z} = \mathbf{R} + \text{H}(\text{pk}, \mathbf{R}, \mu) \cdot \text{pk}$ and \mathbf{z} has small norm, where H is a cryptographic hash function.

Despite this progress, it remains an open problem whether we can construct an efficient 2-round lattice-based analog of FROST that is provably secure under standard lattice assumptions. Existing approaches either require a higher number of rounds [54, 87, 61], relies on non-standard assumptions [60] or has significantly larger signature size and communication complexity [41].

OUR CONTRIBUTIONS.

Blind signatures in pairing-free groups. We present the first blind signature schemes in pairing-free groups that simultaneously achieve: (1) concurrent security, (2) provable security in either the generic group model (GGM) [104, 118] or the algebraic group model (AGM) [64] under the DL assumption additionally assuming random oracles, and (3) efficiency comparable to blind Schnorr signatures.

Our GGM-based construction (BS_1) produces signatures of size 3 scalars—just one more than a standard Schnorr signature—and requires 2 group elements and 3 scalars in communication. Our AGM-based construction (BS_2) has the same signature size and adds one additional scalar to the communication.

Previously, the only scheme satisfying the first two properties was due to Abe [3, 85], but it is significantly less efficient: Abe’s signatures consist of 2 group elements and 6 scalars, with a communication cost of 3 group elements, 6 scalars, and an additional κ bits, where κ denotes the security parameter.

Moreover, we show that BS_2 can be naturally extended to support partially blind signing without incurring any additional cost in signature size or communication complexity. In this setting, a portion of the message to be signed is publicly known to both the signer and the user. This feature is particularly useful in real-world applications, as it allows binding public information, such as the date of issuance, to the signature. This part of the contributions is presented in Part I.

Threshold signatures in pairing-free groups. In this thesis, we formalize a notion called partially non-interactive threshold signatures, which captures schemes like FROST, and develop a security hierarchy for partially non-interactive schemes. These schemes consist of a message-

independent pre-processing round followed by a non-interactive signing round. In the pre-processing round, each signer generates a token and sends it to a leader (coordinator). In the signing round, the leader constructs a leader request lr , which includes the message m to be signed, the set of participating signers S , and an aggregated message derived from the signers' tokens. For FROST, the aggregated message is $\{(i, R_i, S_i)\}_{i \in S}$. The leader then sends lr to the involved signers, who respond with their partial signatures. Finally, the leader aggregates these partial signatures to produce the final signature.

In our unforgeability game, the adversary can corrupt the leader and up to $t-1$ signers. It also has access to two oracles: a pre-processing oracle, which allows the adversary to request pre-processing tokens from any honest party, and a signing oracle, which allows the adversary to request partial signatures from any honest party on any leader request. The adversary wins if it can forge a signature that is not "considered issued". We propose different ways of defining "considered issued", leading to a security hierarchy TS-UF-0 to TS-UF-4, where TS-UF- $(i+1)$ is stronger than TS-UF- i .

The weakest security, TS-UF-0, considers a signature for a message m to be issued if and only if the adversary has obtained any partial signature from an honest party on a leader request for signing message m . This is exactly the security notion considered by the prior work on FROST [50]. TS-UF-1 considers a signature for a message m to be issued if and only if the adversary has obtained partial signature from at least $t - |CS|$ honest signers for signing message m , where CS denotes the set of corrupted signers. A similar notion is considered in the case of non-interactive threshold signatures [119, 94].

For partially non-interactive schemes, it is possible to consider stronger security notions. The stronger notions in the hierarchy aim to guarantee that the partial signatures obtained for different leader requests cannot be combined. TS-UF-2 considers a signature for m to be issued if and only if there exists a leader request lr for signing m such that the adversary has obtained partial signatures for lr from at least $t - |CS|$ signers. The condition of an issued signature in TS-UF-3 additionally requires the adversary to obtain partial signatures from each honest signer $i \in lr.S$ that the token (R_i, S_i) in lr is generated by i . The strongest security in our hierarchy, TS-UF-4, considers a signature for a message m to be issued if and only if there exists a leader request lr for signing m such that the adversary has obtained partial signatures from all honest signers involved in $lr.S$.

We also analyze FROST1 (the original FROST [91]) and two variants FROST2 [50] and FROST3 [113], using our security hierarchy. In particular, we show TS-UF-3 of FROST1, TS-UF-2 of FROST2, and TS-UF-1 of FROST3 under the AOMDL assumption [105] in the ROM. We also give counterexamples showing that they do not guarantee stronger notions in the hierarchy.

Our second main contribution here is a new approach that can transform FROST into threshold signature schemes whose security relies on weaker assumptions. Specifically, while all three FROST variants rely on the AOMDL assumption in the ROM, our new schemes are provably secure under the plain DL assumption in the ROM. The resulting schemes produce Okamoto signatures [107], which are just one scalar larger than Schnorr signatures. This part of the contributions is presented in Part II.

Lattice-based threshold signatures. We establish the security of the 2-round lattice-based threshold signatures proposed by Espitau et al. [60] (referred to as the EKT scheme) under the MSIS and MLWE assumptions in the ROM. Previously, the security of EKT relied on the hardness of the algebraic one-more MLWE (AOM-MLWE) problem, a new assumption introduced in their work.

To achieve this, we propose a modification of the AOM-MLWE problem, referred to as the algebraic one-more MISIS problem (AOM-MISIS) problem. The new problem is no harder than the original problem, meaning that the hardness of AOM-MISIS implies the hardness of AOM-MLWE. We then show the hardness of the AOM-MISIS problem is implied by the MSIS and MLWE assumptions. The reduction is inspired by our techniques for removing the AOMDL assumption.

We also present a new security analysis of EKT directly based on the hardness of the AOM-MISIS problem, which yields a tighter reduction. The concrete parameters derived from our reduction is slightly weaker than those in their work, as their estimates are based their cryptanalysis of the AOM-MLWE problem rather than security reductions.

Our techniques also lead to a new security analysis of another efficient 2-round lattice-based threshold scheme [41] (referred to as the CATZ scheme), which significantly improves its concrete parameters. Compared to EKT, the main advantage of CATZ is that it avoids the need for setting up pairwise secret keys between signers during distributed key generation.

Instead, similar to FROST, each signer only holds its own share of the secret key. This part of the contributions is presented in Part [III](#).

Chapter 2

NOTATION

INTEGERS AND VECTORS. If $b \geq a \geq 1$ are positive integers, then \mathbb{Z}_a denotes the set $\{0, \dots, a-1\}$ and $[a..b]$ denotes the set $\{a, \dots, b\}$. We write $\mathbb{Z}_p^* = \mathbb{Z}_p \setminus \{0\}$. If \mathbf{x} is a vector then $|\mathbf{x}|$ is its length (the number of its coordinates), x_i is its i -th coordinate. For a sequence of variables x_1, \dots, x_ℓ , we use $x_{[i]}$ to denote (x_1, \dots, x_j) and $x_{[i..j]}$ to denote (x_i, \dots, x_j) .

SETS AND DISTRIBUTIONS. Let S be a finite set. The size of a set S is denoted $|S|$. We let $x \leftarrow_{\$} S$ denote sampling an element uniformly at random from S and assigning it to x . For a distribution \mathcal{D} , denote $\text{Supp}(\mathcal{D})$ as the support of \mathcal{D} , and $x \leftarrow_{\$} \mathcal{D}$ denotes sampling x according to \mathcal{D} . For any $y \in \text{Supp}(\mathcal{D})$, denote $\mathcal{D}(y) := \Pr_{x \leftarrow_{\$} \mathcal{D}}[x = y]$, and for $S \subseteq \text{Supp}(\mathcal{D})$, denote $\mathcal{D}(S) := \Pr_{x \leftarrow_{\$} \mathcal{D}}[x \in S]$. For any set T and any function $F : \text{Supp}(\mathcal{D}) \rightarrow T$, denote $F(\mathcal{D})$ as the distribution of $F(x)$ for x sampled from \mathcal{D} .

VECTOR SPACES. For any vector space V over a field F and a set $S \subseteq V$, we denote $\text{Span}_F(S)$ as the F -span of S , which is the smallest F -subspace of V that contains S . In particular, we omit F from the subscript if $F = \mathbb{R}$. For a finite set $S = \{v_1, \dots, v_n\} \subseteq V$, we say S is F -linearly independent if and only if for any non-zero $(a_1, \dots, a_n) \in F^n$, $\sum_{i \in [n]} a_i v_i \neq 0$. We say S is a F -basis of V if and only if S is F -linearly independent and $\text{Span}_F(S) = V$. When F is not specified, we assume $F = \mathbb{R}$. The dimension of V is equal to the size of S .

ALGORITHMS. We let $y \leftarrow A^{O_1, \dots}(x_1, \dots; r)$ denote executing algorithm A on inputs x_1, \dots and coins r with access to oracles O_1, \dots and letting y be the result. We let $y \leftarrow_{\$} A^{O_1, \dots}(x_1, \dots)$ be the result of picking r at random and letting $y \leftarrow A^{O_1, \dots}(x_1, \dots; r)$. Algorithms are randomized unless otherwise indicated. The running time is the worst case. We use κ to denote the security parameter.

GROUPS. The group generation algorithm $\text{GGen}(1^\kappa)$ takes the security parameter as input and outputs (\mathbb{G}, p, g) , where \mathbb{G} is a cyclic group with prime size $p \geq 2^\kappa$ and generator g . We let $1_{\mathbb{G}}$ denote the identity element of \mathbb{G} . We let $\mathbb{G}^* = \mathbb{G} \setminus \{1_{\mathbb{G}}\}$ denote the set of non-identity elements, which is the set of generators of \mathbb{G} if the latter has prime order. If $g \in \mathbb{G}^*$ is a generator and $X \in \mathbb{G}$, the discrete logarithm base g of X is denoted $\text{DLog}_{\mathbb{G}, g}(X)$, and it is in the set $\mathbb{Z}_{|\mathbb{G}|}$. The

Game $\text{DLog}_{\text{GGen}}^{\mathcal{A}}(\kappa)$:

$(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$; $g \leftarrow g(\mathbb{G}_\lambda)$; $X \leftarrow_{\mathcal{S}} \mathbb{G}_\lambda$

$y \leftarrow \mathcal{A}(p, g, \mathbb{G}, X)$

If $g^y = X$ then return 1

Return 0

Figure 2.1: The DLog game.

discrete logarithm game is defined in Figure 2.1. For a DL adversary \mathcal{A} , its advantage is denoted as $\text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{A}, 1^\kappa) := \Pr[\text{DLog}_{\text{GGen}}^{\mathcal{A}}(\kappa) = 1]$.

MODULES. For any ring R with multiplicative identity 1 and any abelian group $(M, +)$, we say M is an R -module if there exists an operation $\cdot : R \times M \rightarrow M$ such that for any $a, b \in R$ and any $x, y \in M$, (i) $a \cdot (x + y) = a \cdot x + a \cdot y$, (ii) $(a + b) \cdot x = a \cdot x + b \cdot x$, (iii) $(ab) \cdot x = a \cdot (b \cdot x)$, (iv) $1 \cdot x = x$. Also, we use 0 to denote the identity of M .

MODULE HOMOMORPHISMS. For any R -modules M and N , a map $f : M \rightarrow N$ is a homomorphism of R -modules if for any $r \in R$ and $x, y \in M$, $f(x + r \cdot y) = f(x) + r \cdot f(y)$. We say a homomorphism f is an epimorphism if f is a surjection. We say a homomorphism f is a monomorphism if f is an injection.

CHARACTERISTIC OF A FIELD. For any field \mathbb{F} , the characteristic of \mathbb{F} , denoted by $\text{char}(\mathbb{F})$, is the smallest positive number k such that $k \cdot \mathbf{1} = \sum_{i=1}^k \mathbf{1} = \mathbf{0}$, where $\mathbf{1}$ denotes the multiplicative identity of \mathbb{F} and $\mathbf{0}$ denotes the additive identity of \mathbb{F} . If k does not exist, we say the characteristic of F is 0.

Part I

BLIND SIGNATURES IN PAIRING-FREE GROUPS

Chapter 3

BLIND SIGNATURES: INTRODUCTION

Blind signatures [42] allow a *user* to interact with a *signer* to produce a valid signature that cannot be linked back by the signer to the interaction that produced it. Blind signatures are used in several applications, such as e-cash systems [42, 46], anonymous credentials (e.g., [37]), privacy-preserving ad-click measurement [1], and various forms of anonymous tokens [82, 2]. They are also covered by an RFC draft [56].

This thesis develops the first practical pairing-free three-move blind signature schemes that (1) are concurrently secure, (2) produce short signatures (i.e., *three* or *four* group elements/scalars), and (3) are provably secure either in the *generic group model* (GGM) [118, 101] or in the *algebraic group model* (AGM) [64] under the discrete logarithm (DL) or the one-more discrete logarithm (OMDL) assumption (in addition to assuming *random oracles* [18]). Our DL-based scheme also admits a *partially blind* version [4], roughly following a paradigm by Abe and Okamoto [5], that targets applications where signatures need to depend on some public input (e.g., an issuing date) known to the signer. An overview of our schemes is given in Table 3.1.

Unlike blind Schnorr [47], Okamoto-Schnorr [110], and other generic constructions based on identification schemes [80], we do not rely on the hardness of the ROS problem, for which a polynomial-time attack has recently been presented [24]. Also, unlike Clause Blind Schnorr (CBS) signatures [65], we do not rely on the assumed hardness of the mROS problem, which is subject to (mildly) sub-exponential attacks and we can thus support smaller group sizes.¹ In fact, our schemes all admit tight bounds, and this suggests that they can achieve $(\lambda/2)$ -bit of security on λ -bit elliptic curves, supporting an instantiation with 256-bit curves. Our security proofs rely on a reduction to a new variant of the ROS problem, called *weighted fractional ROS* (WFROS), for which we prove an exponential, unconditional lower bound. Therefore, another benefit over CBS, beyond concrete parameters, is that we do not need to rely on an additional assumption.

¹The best known attack against mROS [65] runs in time $2^{\ell + \log(\ell+1) + \lambda/(1+\log(\ell+1))}$, where λ is the security parameter and ℓ corresponds to the number of concurrent sessions. The worst ℓ gives a $2^{O(\lambda/\log \lambda)}$ attack, and in practice, this suggests a choice of $\lambda = 512$ to achieve 128-bit security for all ℓ 's.

Perhaps as a testament of the unsatisfactory status of pairing-free schemes, the *only* other scheme known to achieve exponential, concurrent, security is Abe’s scheme [3]. Although its original (standard-model) proof was found to be flawed, proofs were then given both in the GGM [106] and the AGM [85], along with a proof for the restricted setting of sequential security [11]. Still, it produces longer signatures and public keys, and is overall less efficient. Also, it only offers computational blindness (under DDH), whereas our scheme provides perfect blindness.

DISCRETE-LOGARITHM BASED BLIND SIGNATURES. We stress that our focus here is making pairing-free schemes as practical and as secure as possible. Indeed, very simple pairing-based blind signature schemes in the ROM can be obtained from BLS signatures [31, 27]. Blind BLS offers a different trade-off: signatures are short (i.e., one group element) and signing requires only *two* moves, but signature verification requires a more expensive (and more complex) pairing evaluation. Indeed, the current blind signature RFC draft [56] favors RSA over BLS, also due to lesser availability of pairings implementations. In particular, several envisioned applications of blind signatures are inherently browser-based, and the available cryptographic libraries (e.g., NSS for Firefox and BoringSSL for Chrome) do not yet offer pairing-friendly curve implementations.

In contrast, (non-blind) Schnorr signatures [114, 116] (such as EdDSA [25]) are short, can rely on standard libraries, and outperform RSA. Though their blind evaluation requires three rounds, this may be less concerning in applications where verification cost is the dominating factor and the signing application can easily keep state. Indeed, [56] identifies CBS as the only plausible alternative to RSA, and our schemes improve upon CBS by avoiding the mROS assumption. Once the group order is adjusted to resist sub-exponential attacks, we achieve comparable signature size, more efficient signing, and accommodate for partial blindness. (No partially blind version of CBS is known to the best of our knowledge.)

Finally, note that it is easier to prove security of pairing-free schemes under sequential access to the signer. For example, Kastner et al. [85] prove that plain blind Schnorr signatures are secure in this case, in the AGM, assuming the hardness of OMDL. Also, Baldimtsi and Lysyanskaya [11] (implicitly) prove sequential security of Abe’s scheme. However, many applications, like PCM, easily enable concurrent attacks.

ON IDEAL MODELS. The use of the AGM or the GGM, along with the ROM, still appears necessary for the most practical pairing-free schemes with concurrent security. As of now, solutions solely assuming the ROM can only handle bounded concurrency [80] or, alternatively, their communication and computation costs grow with the number of signing sessions [88, 40, 123].

Scheme	PK size	Sig. size	Assumption	Communication
BS ₁ (Section 5.1)	1 \mathbb{G}	3 \mathbb{Z}_p	GGM	2 \mathbb{G} + 3 \mathbb{Z}_p
BS ₂ (Section 6.1)	2 \mathbb{G}	4 \mathbb{Z}_p	DL	2 \mathbb{G} + 4 \mathbb{Z}_p
PBS (Section 6.3)	1 \mathbb{G}	4 \mathbb{Z}_p	DL	2 \mathbb{G} + 4 \mathbb{Z}_p
Blind Schnorr [65]	1 \mathbb{G}	2 \mathbb{Z}_p	OMDL + ROS	1 \mathbb{G} + 2 \mathbb{Z}_p
Clause Blind Schnorr [65]	1 \mathbb{G}	2 \mathbb{Z}_p	OMDL + mROS	2 \mathbb{G} + 4 \mathbb{Z}_p
Abe [3, 85]	3 \mathbb{G}	2 \mathbb{G} + 6 \mathbb{Z}_p	DL	λ bits + 3 \mathbb{G} + 6 \mathbb{Z}_p

Table 3.1: **Overview of our results.** The four schemes proposed in this paper compared to pairing-free schemes that admit GGM/AGM security proofs in the literature. All schemes are three-move and secure assuming the ROM; All schemes except BS₁ admit AGM security proofs; further $p = |\mathbb{G}|$. As in plain Schnorr signatures, most schemes allow replacing one element in \mathbb{Z}_p with a group element in the signature. The ROS assumption can be broken in polynomial time unless the scheme is restricted to tolerate only a very small number of sessions. Also, the mROS assumption admits sub-exponential attacks, which require the choice of a larger order p over all schemes (roughly 512-bit for 128-bit security [65]).

A number of other schemes [68, 26, 67, 63, 62, 77, 86] partially or completely avoid ideal models, some of which are fairly practical. However, they do not yet appear suitable for at-scale deployment.

3.1 A Scheme in the GGM

Our simplest scheme only admits a proof in the generic-group model (GGM) but best illustrates our ideas, in particular, how we bypass ROS-style attacks. It is slightly less efficient than Schnorr signatures, i.e., a signature that consists of *three* scalars mod p (or alternatively, two scalars and a group element). Nonetheless, it has a very similar flavor (in particular, signature verification can be built on top of a suitable implementation of Schnorr signatures in a black-box way).

PREFACE: BLIND SCHNORR SIGNATURES AND ROS. Recall that we seek an interactive scheme (1) that is one-more unforgeable (i.e., no adversary should be able to generate $\ell + 1$ signatures by interacting only ℓ times with the signer), and (2) for which interaction can be blinded. It is helpful to illustrate the main technical barrier behind proving (1) for *interactive* Schnorr signatures. Recall

that the verification key is $X = g^x$ for a generator g of a cyclic group \mathbb{G} of prime order p , and a signing key x . The signer starts the session by sending $A = g^a$, for a random $a \in \mathbb{Z}_p$. Then, the user sends a challenge $c = H(A, m)$ for a hash function H and a message m to be signed. Finally, the signer responds with $s = a + c \cdot x$, and the signature is $\sigma = (c, s)$.

Let us now consider an adversary that obtains ℓ initial messages A_1, \dots, A_ℓ from the signer, where $A_i = g^{a_i}$. By solving the so-called *ROS problem* [117, 80, 65], the attacker can find $\ell + 1$ vectors $\alpha_1, \dots, \alpha_{\ell+1} \in \mathbb{Z}_p^\ell$ and a vector $(c_1, \dots, c_\ell) \in \mathbb{Z}_p^\ell$ such that

$$\sum_{j=1}^{\ell} \alpha_i^{(j)} \cdot c_j = c_i^* \quad (3.1)$$

for all $i \in [\ell + 1]$, where $c_i^* = H(\prod_{j=1}^{\ell} A_j^{\alpha_i^{(j)}}, m_i^*)$, for some message $m_i^* \in \{0, 1\}^*$. (Here, $\alpha_i^{(j)}$ is the j -th component of α_i .) Then, the attacker can obtain $s_j = a_j + c_j x$ from the signer for all $j \in [\ell]$ by completing the ℓ signing sessions. It is now easy to verify that (c_i^*, s_i^*) is a valid signature for m_i^* for all $i \in [\ell + 1]$, where $s_i^* = \sum_{j=1}^{\ell} \alpha_i^{(j)} \cdot s_j$. Benhamouda et al. [24] recently gave a simple polynomial-time algorithm to solve the ROS problem for the case $\ell > \log(p)$, which thus breaks one-more unforgeability.²

Fuchsbauer et al. [65] propose a different interactive signing process for Schnorr signatures that is one-more unforgeable (in the AGM + ROM) assuming that a variant of the ROS problem, called mROS, is hard. The mROS problem, however, admits sub-exponential attacks, and as it gives approximately only 70 bits of security from an implementation on a 256-bit curve, it effectively forces the use of 512-bit curves.³

OUR FIRST SCHEME. We take a different path which completely avoids the ROS and mROS problems to obtain our first scheme, BS₁. Again, we present a non-blind version – the scheme can be made blind via fairly standard tricks, as we explain in the body of the paper below. Again, the public key is $X = g^x$ for a secret key x . Then, the signer and the user engage in the following protocol to sign $m \in \{0, 1\}^*$:

1. The signer sends $A = g^a$ and $Y = X^y$ for random $a, y \in \mathbb{Z}_p$.

²Many envisioned implementations allow for $\ell > \log(p)$. Still, it is worth noting that the scheme retains some security for $\ell < \log(p)$ even in the standard model [80].

³mROS depends on a parameter ℓ , with a similar role as in ROS – sub-exponential attacks require $\ell < \log(p)$, but a one-more unforgeability attack for a small ℓ implies one for any $\ell' > \ell$ simply by generating $(\ell' - \ell)$ additional valid signatures.

2. The user responds with $c = H(A, Y, m)$
3. The signer returns a pair (s, y) , where $s = a + cxy$.
4. The user accepts the signature $\sigma = (c, s, y)$ iff $g^s = A \cdot Y^c$ and $Y = X^y$.

Verification simply checks that $H(g^s X^{-yc}, X^y, M) = c$. In particular, note that (c, s) is a valid Schnorr signature with respect to the public-key X^y – this can be leveraged to implement the verification algorithm on top of an existing implementation of basic Schnorr signatures that also hash the public key (EdDSA does exactly this).⁴ Further, as in Schnorr signatures, we could replace c with A in σ , and our results would be unaffected.

SECURITY INTUITION. To gather initial insights about the security of BS_1 , it is instructive to *attempt* an ROS-style attack. The attacker opens ℓ sessions and obtains pairs $(A_1, Y_1), \dots, (A_\ell, Y_\ell)$, where $A_i = g^{a_i}$ and $Y_i = X^{y_i} = g^{xy_i}$ for all $i \in [\ell]$. One natural extension of the ROS attack is to find $\ell + 1$ vectors $\alpha_i \in \mathbb{Z}_p^\ell$ along with messages $m_1^*, m_2^*, \dots \in \{0, 1\}^*$ such that

$$c_i^* = H \left(\prod_{j=1}^{\ell} A_j^{\alpha_i^{(j)}}, \prod_{j=1}^{\ell} Y_j^{\alpha_i^{(j)}}, m_i^* \right)$$

for all $i \in [\ell + 1]$ and then find $(c_1, \dots, c_\ell) \in \mathbb{Z}_p^\ell$ such that

$$\sum_{j=1}^{\ell} \alpha_i^{(j)} \cdot y_j \cdot c_j = c_i^* \cdot \sum_{j=1}^{\ell} \alpha_i^{(j)} \cdot y_j, \quad (3.2)$$

for all $i \in [\ell + 1]$. Indeed, if this succeeded, the adversary could complete the ℓ sessions to learn (s_j, y_j) by inputting c_j , where y_j is random and $s_j = a_j + c_j \cdot x \cdot y_j$. One could generate $\ell + 1$ signatures (c_i^*, s_i^*, y_i^*) for $i \in [\ell + 1]$ by setting $s_i^* = \sum_{j=1}^{\ell} \alpha_i^{(j)} s_j$ and $y_i^* = \sum_{j=1}^{\ell} \alpha_i^{(j)} \cdot y_j$. These would be valid because

$$\begin{aligned} g^{s_i^*} &= g^{\sum_{j=1}^{\ell} \alpha_i^{(j)} (a_j + c_j x y_j)} \\ &= \prod_{j=1}^{\ell} A_j^{\alpha_i^{(j)}} \cdot X^{\sum_{j=1}^{\ell} \alpha_i^{(j)} c_j y_j} \stackrel{(3.2)}{=} \prod_{j=1}^{\ell} A_j^{\alpha_i^{(j)}} \cdot \left(\prod_{j=1}^{\ell} Y_j^{\alpha_i^{(j)}} \right)^{c_i^*}. \end{aligned}$$

⁴Note that this only superficially resembles key-blinding for Schnorr signatures [83]. Here, the “blinding” y is actually public and part of the signature.

However, finding (c_1, \dots, c_ℓ) that satisfy (3.2) for $\ell + 1$ i 's simultaneously is *much harder* than ROS. An initial intuition here is that X^y *completely hides* y to the point where y is revealed later in the session, where it appears like a random and fresh weight in the sum, *independent of* c_i . This intuition is however not correct, as an attacker can use the group element X^y and can try to gain information about y , but our proof will show (among other things) that in the GGM no useful information is obtained about y , and y is (close to) uniform when it is later revealed.

THE WFROS PROBLEM. The above attack paradigm is in fact generalized in terms of a new ROS-like problem that we call WFROS (this stands for *Weighted Fractional ROS*), for which we prove an unconditional lower bound. WFROS considers a game with two oracles that can be invoked adaptively in an interleaved way:

- The first oracle, H, accepts as input a pair of vectors $\alpha, \beta \in \mathbb{Z}_p^{2\ell+1}$, which are then associated with a random $\delta \in \mathbb{Z}_p^*$.
- The second oracle, S, allows to bind, for some $i \in [\ell]$, chosen input $c_i \in \mathbb{Z}_p$ with a random *weight* $y_i \in \mathbb{Z}_p^*$. During the course of the game, this latter oracle must be called *exactly* once for each $i \in [\ell]$.

The adversary finally commits to a subset of $\ell + 1$ prior H queries and wins if for each query in the subset, which has defined a pair of vector α, β and returned δ , we have $A/B = \delta$, where

$$A = \alpha^{(0)} + \sum_{i \in [\ell]} y_i (\alpha^{(2i-1)} + c_i \cdot \alpha^{(2i)}), \quad B = \beta^{(0)} + \sum_{i \in [\ell]} y_i (\beta^{(2i-1)} + c_i \cdot \beta^{(2i)}).$$

Here, $v^{(i)}$ denotes the i -th component of vector v . Our main result (Theorem 4.1.1) says that no adversary making Q_H queries to H can win this game with probability better than $(Q_H^2 + 2\ell Q_H)/(p-1)$, or, in other words, $Q_H \geq \min\{\sqrt{p}, p/\ell\}$ is needed to win with constant probability. Note that $\ell \ll \sqrt{p}$ is generally true, as for our usage, ℓ is bounded by the number of signing sessions.

Our GGM proof for BS_1 transforms any generic attacker into one breaking the WFROS problem. This transformation is actually not immediate because a one-more unforgeability attacker can learn functions of the secret key x when obtaining the second message from the signer. A similar challenge occurs in proving hardness of the OMDL problem in the GGM, which was recently resolved by Bauer et al. [12], and we rely on their techniques.

3.2 AGM Security and Partial Blindness

The Algebraic Group Model (AGM) [64] can be seen as a weaker idealization than the GGM. In particular, AGM proofs deal with actual groups (as opposed to representing group elements with random labels) and proceed via *reductions* that apply only to “algebraic adversaries”, which provide representation of the group elements they output to the reduction. AGM has become a very popular model for validating security of a number of practical group-based protocols.

The main barrier to proving one-more unforgeability of BS_1 in the AGM is that the representation of X^y could leak some information about y that would not be available in the GGM, and thus we would not be able to apply our argument showing that y is still (close to) random looking when it is later revealed – our reduction in the GGM security proof crucially relies on this. To overcome this issue, for BS_2 , we replace X^y with a *hiding* commitment to y :

Scheme BS_2 . Here, $g^t X^y$ is replaced by $g^t Z^y$, where Z is an extra random group element included in the verification key.

We can prove security of BS_2 in the AGM based *solely* on the DL assumption. While BS_2 requires a longer key than BS_1 , one could mitigate this by obtaining Z as the output of a hash function (assumed to be a random oracle) evaluated on some public input. The proof of security for BS_2 consists of showing that any adversary breaking one-more unforgeability can be transformed into one breaking either OMDL or DL (depending on the scheme) *or* into one breaking the WFROS problem. For the latter, however, we can resort to our unconditional hardness lower bound (Theorem 4.1.1).

ADDING PARTIAL BLINDNESS. Finally, we note that it is not too hard to add partial blindness to BS_3 , which is another reason to consider this scheme. In particular, to obtain the resulting PBS scheme, we can adopt a framework by Abe and Okamoto [5]. The main idea is simply to use a hash function (modeled as a random oracle) to generate the extra group element Z in a way that is dependent on a public input upon which the signature depends. We target in particular a stronger notion of one-more unforgeability, which shows that if the protocol is run ℓ times for a public input, then no $\ell + 1$ signatures can be generated for that public input regardless of how many signatures have been generated for *different* public inputs. We defer more details to Section 6.3.

Outline of Part I

Section 3.3 will introduce some basic preliminaries. Chapter 4 will then introduce the WFROS problem, and prove a lower bound for it. We will then discuss our GGM-based scheme in Chapter 5, whereas variants secure in the AGM are presented in Chapter 6. In particular, we give a partially blind instantiation of our AGM scheme in Section 6.3.

3.3 Preliminaries

BLIND SIGNATURES. This paper focuses on *three-move* blind signature schemes, and our notation is similar to that of prior works (e.g., [80, 65]). Formally, a (three-move) *blind signature scheme* BS is a tuple of efficient (randomized) algorithms

$$\text{BS} = (\text{Setup}, \text{KeyGen}, S_1, S_2, U_1, U_2, \text{Ver}) ,$$

with the following behavior:

- The *parameter generation* algorithm $\text{Setup}(1^\lambda)$ outputs a string of parameters par , whereas the *key generation* algorithm $\text{KeyGen}(par)$ outputs a key-pair (sk, pk) , where sk is the *secret* (or *signing*) key and pk is the *public* (or *verification*) key.
- The interaction between the user and the signer to sign a message $m \in \{0, 1\}^*$ with key-pair (pk, sk) is defined by the following experiment:

$$\begin{aligned} (\text{st}^s, \text{msg}_1) &\leftarrow \text{BS}.S_1(sk) , & (\text{st}^u, \text{chl}) &\leftarrow \text{BS}.U_1(pk, \text{msg}_1, m) , \\ \text{msg}_2 &\leftarrow \text{BS}.S_2(\text{st}^s, \text{chl}) , & \sigma &\leftarrow \text{BS}.U_2(\text{st}^u, \text{msg}_2) . \end{aligned} \tag{3.3}$$

Here, σ is either the resulting *signature* or an *error message* \perp .

- The (deterministic) *verification algorithm* outputs a bit $\text{BS}.Ver(pk, \sigma, m)$.

We say that BS is (perfectly) *correct* if for every message $m \in \{0, 1\}^*$, with probability one over the sampling of parameters and the key pair (pk, sk) , the experiment in (3.3) returns σ such that $\text{BS}.Ver(pk, \sigma, m) = 1$. All of our schemes are going to be perfectly correct.

<p><u>Game $\text{OMUF}_{\text{BS}}^{\mathcal{A}}(\lambda)$:</u></p> <p>$par \leftarrow \text{BS.Setup}(1^\lambda)$</p> <p>$(sk, pk) \leftarrow \text{BS.KeyGen}(par)$</p> <p>$sid \leftarrow 0 ; \ell \leftarrow 0 ; \mathcal{I}_{\text{fin}} \leftarrow \emptyset$</p> <p>$\{(m_k^*, \sigma_k^*)\}_{k \in [\ell+1]} \leftarrow_{\\$} \mathcal{A}^{\text{S}_1, \text{S}_2}(pk)$</p> <p>If $\exists k_1 \neq k_2$ such that $(m_{k_1}^*, \sigma_{k_1}^*) = (m_{k_2}^*, \sigma_{k_2}^*)$</p> <p style="padding-left: 20px;">then return 0</p> <p>If $\exists k \in [\ell + 1]$ such that $\text{BS.Ver}(pk, \sigma_k^*, m_k^*) = 0$</p> <p style="padding-left: 20px;">then return 0</p> <p>Return 1</p>	<p><u>Oracle S_1 :</u></p> <p>$sid \leftarrow sid + 1$</p> <p>$(\text{st}_{\text{sid}}^s, \text{msg}_1) \leftarrow \text{BS.S}_1(sk)$</p> <p>Return (sid, msg_1)</p> <p><u>Oracle $\text{S}_2(i, c_i)$:</u></p> <p>If $i \notin [sid] \setminus \mathcal{I}_{\text{fin}}$ then return \perp</p> <p>$\text{msg}_2 \leftarrow \text{BS.S}_2(\text{st}_i^s, c_i)$</p> <p>$\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$</p> <p>$\ell \leftarrow \ell + 1$</p> <p>Return msg_2</p>
---	--

Figure 3.1: The OMUF security game for a blind signature scheme BS.

ONE-MORE UNFORGEABILITY. The standard notion of security for blind signatures is *one-more unforgeability* (OMUF). OMUF ensures that no adversary playing the role of a user interacting with the signer ℓ times, in an arbitrarily concurrent fashion, can issue $\ell + 1$ signatures (or more, of course). The $\text{OMUF}_{\text{BS}}^{\mathcal{A}}$ game for a blind signature scheme BS is defined in Figure 3.1. The corresponding advantage of \mathcal{A} is defined as $\text{Adv}_{\text{BS}}^{\text{omuf}}(\mathcal{A}, \lambda) := \Pr[\text{OMUF}_{\text{BS}}^{\mathcal{A}}(\lambda) = 1]$. All of our analyses will further assume one or more random oracles, which are modeled as an additional oracle to which the adversary \mathcal{A} is given access.

BLINDNESS. We also consider the standard notion of blindness against a malicious server that can, in particular, attempt to publish a malformed public key. The corresponding game $\text{Blind}_{\text{BS}}^{\mathcal{A}}$ is defined in Figure 3.2, and for any adversary \mathcal{A} , we define its advantage as $\text{Adv}_{\text{BS}}^{\text{blind}}(\mathcal{A}, \lambda) := |\Pr[\text{Blind}_{\text{BS}}^{\mathcal{A}}(\lambda) = 1] - \frac{1}{2}|$. We say the scheme is perfectly blind if and only if $\text{Adv}_{\text{BS}}^{\text{blind}}(\mathcal{A}, \lambda) = 0$ for any \mathcal{A} and all λ .

GAME-PLAYING PROOFS. Several of our proofs adopt a lightweight variant of the standard “Game-Playing Framework” by Bellare and Rogaway [19].

<p><u>Game Blind_{BS}^A(λ) :</u> $par \leftarrow \text{BS.Setup}(1^\lambda)$ $b \leftarrow_{\\$} \{0, 1\}; b_0 \leftarrow b; b_1 \leftarrow 1-b$ $b' \leftarrow_{\\$} \mathcal{A}^{\text{INIT}, U_1, U_2}(par)$ If $b' = b$ then return 1 Return 0</p> <p><u>Oracle INIT($\tilde{pk}, \tilde{m}_0, \tilde{m}_1$) :</u> $sess_0 \leftarrow \text{init}$ $sess_1 \leftarrow \text{init}$ $pk \leftarrow \tilde{pk}$ $m_0 \leftarrow \tilde{m}_0; m_1 \leftarrow \tilde{m}_1$</p>	<p><u>Oracle U₁($i, \text{msg}_1^{(i)}$) :</u> If $i \notin \{0, 1\}$ or $sess_i \neq \text{init}$ then return \perp $sess_i \leftarrow \text{open}$ $(st_i^u, \text{chl}^{(i)}) \leftarrow \text{BS.U}_1(pk, \text{msg}_1^{(i)}, m_{b_i})$ Return $\text{chl}^{(i)}$</p> <p><u>Oracle U₂($i, \text{msg}_2^{(i)}$) :</u> If $i \notin \{0, 1\}$ or $sess_i \neq \text{open}$ then return \perp $sess_i \leftarrow \text{closed}$ $\sigma_{b_i} \leftarrow \text{BS.U}_2(st_i^u, \text{msg}_2^{(i)})$ If $sess_0 = sess_1 = \text{closed}$ then If $\sigma_0 = \perp$ or $\sigma_1 = \perp$ then return (\perp, \perp) Return (σ_0, σ_1) Return (i, closed)</p>
---	---

Figure 3.2: The Blind security game for a blind signature scheme BS.

Chapter 4

THE WEIGHTED FRACTIONAL ROS PROBLEM

4.1 Problem Description and Lower Bound

This section introduces and analyzes an unconditionally hard problem underlying all of our proofs, which we call the *Weighted Fractional ROS* problem (WFROS). It is a variant of the original ROS problem [117, 80, 65], which, in turn, stands for *Random inhomogeneities in a Overdetermined Solvable system of linear equations*. While ROS can be solved in polynomial time [24] and its mROS variant can be solved in sub-exponential time [65], we are going to prove an *exponential* lower bound for WFROS.

THE WFROS PROBLEM. The problem is defined via the game $\text{WFROS}_{\ell,p}^{\mathcal{A}}$, described in Figure 4.1, which involves an adversary \mathcal{A} and depends on two integer parameters ℓ and p , where p is a prime. The adversary here interacts with two oracles, H and S. The first oracle allows the adversary to link a vector pair $\alpha, \beta \in \mathbb{Z}_p^{2\ell+1}$ with a random inhomogeneous part $\delta \in \mathbb{Z}_p^*$ – each such query defines implicitly an equation $A/B = \delta$ in the unknowns C_1, \dots, C_ℓ and Y_1, \dots, Y_ℓ . A call to $S(i, c_i)$ lets us set the value of C_i to c_i and set Y_i to a random value y_i . The second oracle $S(i, \cdot)$ must be called once for every $i \in [\ell]$. It is noteworthy to stress that the c_i 's can be chosen arbitrarily, whereas the corresponding y_i 's are random and independent.

In the end, the adversary wins the game if a subset of $\ell + 1$ equations defined by the H queries is satisfied by the assignment defined by querying S. In particular, we define $\text{Adv}_{\ell,p}^{\text{wfros}}(\mathcal{A}) = \Pr[\text{WFROS}_{\ell,p}^{\mathcal{A}} = 1]$. Note that it would be possible to carry out some of the following security proofs using restricted versions of the WFROS game, but the above formulation lets us handle all schemes via a single notion.

A LOWER BOUND FOR WFROS. The following theorem, our main result on WFROS, shows that any adversary winning WFROS with constant probability requires $Q_H = \Omega(\min\{\sqrt{p}, p/\ell\})$ queries. (Also, note that all applications of interest assume $\ell \ll \sqrt{p}$.)

Theorem 4.1.1 (Lower bound for WFROS). *For any $\ell > 0$, any prime number p , and any adver-*

<p><u>Game WFROS$_{\ell,p}^A$:</u> hid $\leftarrow 0$; $\mathcal{I}_{\text{fin}} \leftarrow \emptyset$ $\mathcal{J} \leftarrow \mathcal{A}^{\text{H,S}}(p)$ If $\mathcal{J} \not\subseteq [\text{hid}]$ or $\mathcal{J} \leq \ell$ or $\mathcal{I}_{\text{fin}} \neq [\ell]$ then Return 0 For each $j \in \mathcal{J}$, $A_j \leftarrow \alpha_j^{(0)} + \sum_{i \in [\ell]} y_i (\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)})$ $B_j \leftarrow \beta_j^{(0)} + \sum_{i \in [\ell]} y_i (\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)})$ If $\forall j \in \mathcal{J} : (A_j = \delta_j B_j \wedge B_j \neq 0)$ then Return 1 Return 0</p>	<p><u>Oracle H(α, β) :</u> hid $\leftarrow \text{hid} + 1$ $\alpha_{\text{hid}} \leftarrow \alpha$; $\beta_{\text{hid}} \leftarrow \beta$ $\delta_{\text{hid}} \leftarrow \\$_{p}^*$ Return $\delta_{\text{hid}}, \text{hid}$</p> <p><u>Oracle S($i, c_i$) :</u> If $i \notin [\ell] \setminus \mathcal{I}_{\text{fin}}$ then return \perp $y_i \leftarrow \\$_{p}^*$ $\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$ Return y_i</p>
---	--

Figure 4.1: The WFROS problem. Here, $\alpha, \beta \in \mathbb{Z}_p^{2\ell+1}$, which is indexed as $\alpha = (\alpha^{(0)}, \dots, \alpha^{(2\ell)})$ and $\beta = (\beta^{(0)}, \dots, \beta^{(2\ell)})$.

sary \mathcal{A} playing the WFROS $_{\ell,p}^A$ game that makes at most Q_{H} queries to H , we have

$$\text{Adv}_{\ell,p}^{\text{wfros}}(\mathcal{A}) \leq \frac{Q_{\text{H}}(2\ell + Q_{\text{H}})}{p - 1}.$$

The proof is given in the next section. To gain some very high-level intuition, we observe that a key contributor to the hardness of WFROS are values y_i , which are defined *after* the c_i 's are fixed and hence randomize the A_j and B_j 's. Therefore, to satisfy $A_j = \delta_j \cdot B_j$, the adversary is restricted in the way it plays. For example, to satisfy an equation defined by an H query (α_j, β_j) , the adversary can pick c_i 's such that $(\alpha_j^{(2i-1)} + c_i \alpha_j^{(2i)}) = \delta_j \cdot (\beta_j^{(2i-1)} + c_i \beta_j^{(2i)})$ for all $i \in [\ell]$. Then, the equation $A_j = \delta_j B_j$ is satisfied no matter what the y_i 's are. Our proof shows that the adversary *has* to pick c_i 's this way – and in fact, it has to follow even more restrictions. Finally, we show that under these restrictions, no set of $\ell + 1$ equations can be satisfied simultaneously.

4.2 Proof of the Lower Bound for WFROS (Theorem 4.1.1)

Let \mathcal{A} be an adversary for the WFROS game that makes at most Q_{H} queries to H . Without loss of generality, we assume that \mathcal{A} makes exactly one query (i, c_i) to S for each $i \in [\ell]$ and that \mathcal{A} always outputs $\mathcal{J} \subseteq [Q_{\text{H}}]$.

In the $\text{WFROS}_{\ell,p}^{\mathcal{A}}$ game, for each $j \in [Q_H]$, denote the event W_j as

$$\alpha_j^{(0)} + \sum_{i \in [\ell]} y_i (\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)}) = \delta_j \left(\beta_j^{(0)} + \sum_{i \in [\ell]} y_i (\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)}) \right) \quad (\text{W1})$$

$$\wedge \beta_j^{(0)} + \sum_{i \in [\ell]} y_i (\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)}) \neq 0. \quad (\text{W2})$$

In other words, W_j is the event that the equation defined by the j -th H query is satisfied. Then, \mathcal{A} wins if and only if $|\mathcal{J}| > \ell$ and W_j occur for each $j \in \mathcal{J}$. Denote $W := (|\mathcal{J}| > \ell) \wedge \left(\bigwedge_{j \in \mathcal{J}} W_j \right)$ and we have $\text{Adv}_{\ell,p}^{\text{wfros}}(\mathcal{A}) = \Pr[W]$.

To bound $\Pr[W]$, we need notation to refer to some values (formally, random variables) defined in the execution of the $\text{WFROS}_{\ell,p}^{\mathcal{A}}$ game. First, denote as $\mathcal{I}_{\text{fin}}^{(j)}$ the contents of the set \mathcal{I}_{fin} when the adversary makes the j -th query to H, and let (α_j, β_j) be the input of this query to H, which is answered with δ_j . Also, let $\mathcal{I}_{\text{unk}}^{(j)} := [\ell] \setminus \mathcal{I}_{\text{fin}}^{(j)}$, i.e., the set of indices $i \in [\ell]$ for which \mathcal{A} has not yet made any query (i, \cdot) to S when the j -th query to H is made. Further, c_1, \dots, c_ℓ and y_1, \dots, y_ℓ are the values defined by querying S.

Now, for each $j \in [Q_H]$, we define the following events:

Event $E_j^{(1)}$. First, let $E_{1,j}^{(1)}$ be the event that $\beta_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(j)}} y_i (\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)}) \neq 0$. For each $i \in \mathcal{I}_{\text{unk}}^{(j)}$, also let $E_{2,(j,i)}^{(1)}$ be the event that $\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} \neq \delta_j (\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)})$. Finally, let $E_j^{(1)} := E_{1,j}^{(1)} \vee \left(\bigvee_{i \in \mathcal{I}_{\text{unk}}^{(j)}} E_{2,(j,i)}^{(1)} \right)$.

Event $E_j^{(2)}$. We denote the event $E_j^{(2)}$ as the event where

$$\forall i \in \mathcal{I}_{\text{unk}}^{(j)} : \alpha_j^{(2i)} \cdot \beta_j^{(2i-1)} = \alpha_j^{(2i-1)} \cdot \beta_j^{(2i)}. \quad (4.1)$$

Note that events $E_j^{(1)}$ and $E_j^{(2)}$ are, by themselves, not necessarily unlikely – the adversary can certainly provoke them. However, we intend to show that this has implications on the ability to satisfy the j -th equation. In particular, we prove the following two lemmas in Sections 4.2.1 and 4.2.2 below, respectively.

Lemma 4.2.1. $\Pr[W_j \wedge E_j^{(1)}] \leq \frac{\ell+1}{p-1}$.

Lemma 4.2.2. $\Pr[W_j \wedge (\neg E_j^{(1)}) \wedge E_j^{(2)}] \leq \frac{\ell}{p-1}$.

Now, if we denote $E^{(1)} := \bigvee_{j \in [Q_H]} (W_j \wedge E_j^{(1)})$ and $E^{(2)} := \bigvee_{j \in [Q_H]} (W_j \wedge (\neg E_j^{(1)}) \wedge E_j^{(2)})$, the union bound yields $\Pr[E^{(1)}] \leq \frac{Q_H(\ell+1)}{p-1}$ and $\Pr[E^{(2)}] \leq \frac{Q_H \cdot \ell}{p-1}$. Our final lemma (proved in Section 4.2.3) is then the following:

Lemma 4.2.3. $\Pr[W \wedge (\neg E^{(1)}) \wedge (\neg E^{(2)})] \leq \frac{Q_H(Q_H-1)}{p-1}$.

The three lemmas can be combined to obtain

$$\Pr[W] \leq \Pr[E^{(1)}] + \Pr[E^{(2)}] + \Pr[W \wedge (\neg E^{(1)}) \wedge (\neg E^{(2)})] \leq \frac{Q_H(2\ell + Q_H)}{p-1}.$$

which concludes the proof. In the next three sections, we prove the three preceding lemmas.

4.2.1 Proof of Lemma 4.2.1

Throughout this proof, let us fix $j \in [Q_H]$. We first define a sequence of random variables $(D_0, D_1, \dots, D_n, X_1, \dots, X_n)$, where $n = \ell + 1$, such that $E_j^{(1)}$ implies one of D_0, \dots, D_n is not equal to 0 and $D_0 + \sum_{k \in [n]} D_k X_k = 0$. Further, we also ensure that X_k is uniformly distributed over \mathbb{Z}_p^* independent of $(D_0, D_1, \dots, D_k, X_1, \dots, X_{k-1})$ for each $k \in [n]$ and use this to bound $\Pr[E_j^{(1)}]$. More concretely:

- Let

$$D_0 := \alpha_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(j)}} y_i \left(\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i-1)} \right),$$

$$X_1 := -\delta_j, \quad D_1 := \beta_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(j)}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right),$$

and note that $E_{1,j}^{(1)}$ is equivalent to $D_1 \neq 0$.

- Further, for $1 \leq k \leq |\mathcal{I}_{\text{unk}}^{(j)}|$, denote $i_k \in \mathcal{I}_{\text{unk}}^{(j)}$ as the index such that (i_k, c_{i_k}) is the k -th query made to S among the indexes in $\mathcal{I}_{\text{unk}}^{(j)}$ and let

$$X_{k+1} = y_{i_k}, \quad D_{k+1} := \alpha_j^{(2i_k-1)} + c_{i_k} \cdot \alpha_j^{(2i_k)} - \delta_j \left(\beta_j^{(2i_k-1)} + c_{i_k} \cdot \beta_j^{(2i_k)} \right),$$

we have $E_{2,(j,i_k)}^{(1)}$ occurs is equivalent to $D_{k+1} \neq 0$.

- For $|\mathcal{I}_{\text{unk}}^{(j)}| + 1 < k \leq n$, let $D_k = 0$ and X_k be a random variable uniformly distributed in \mathbb{Z}_p^* independent of $(D_0, D_1, \dots, D_k, X_1, \dots, X_{k-1})$.¹

¹For $|\mathcal{I}_{\text{unk}}^{(j)}| + 1 < k \leq n$, D_k, X_k act as placeholders so that we can apply Lemma 4.2.5 for an a priori fixed value n instead of a random variable $|\mathcal{I}_{\text{unk}}^{(j)}| + 1$.

Note that

$$\begin{aligned}
D_0 + \sum_{k=1}^n D_k X_k &= \alpha_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(j)}} y_i \left(\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} \right) \\
&\quad - \delta_j \left(\beta_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(j)}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) \right) \\
&\quad + \sum_{i \in \mathcal{I}_{\text{unk}}^{(j)}} y_i \left(\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} - \delta_j \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) \right) \\
&= \alpha_j^{(0)} + \sum_{i \in [\ell]} y_i \left(\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} \right) \\
&\quad - \delta_j \left(\beta_j^{(0)} + \sum_{i \in [\ell]} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) \right).
\end{aligned}$$

Therefore, by (W1), we know W_j occurs implies $D_0 + \sum_{i=1}^n D_i X_i = 0$. Thus, the event $W_j \wedge E_j^{(1)}$ implies, in addition, that one of D_0, \dots, D_n is not equal to 0. Also, we prove the following claim.

Claim 4.2.4. *For each $k \in [n]$, X_k is uniformly distributed over \mathbb{Z}_p^* independent of $(D_0, \dots, D_k, X_1, \dots, X_{k-1})$.*

proof of Claim 4.2.4. For $k = 1$, we have $X_1 = -\delta_j$. Consider the step when δ_j is generated. Since \mathcal{A} has made the j -th query to H, we know $\mathcal{I}_{\text{unk}}^{(j)}$, β_j , α_j , and $\{y_i, c_i\}_{i \in \mathcal{I}_{\text{fin}}^{(j)}}$ are already determined, which implies D_0 and D_1 are also determined. Since δ_j is picked uniformly at random from \mathbb{Z}_p^* , we know $X_1 = -\delta_j$ is uniformly distributed over \mathbb{Z}_p^* independent of (D_0, D_1) .

For $2 \leq k \leq |\mathcal{I}_{\text{unk}}^{(j)}| + 1$, we have $X_k = y_{i_{k-1}}$. Consider the step when $y_{i_{k-1}}$ is generated. We know \mathcal{A} has made the query $(i_{k-1}, c_{i_{k-1}})$ to S and the values $i_{k-1}, c_{i_{k-1}}$ are determined. Since $i_{k-1} \in \mathcal{I}_{\text{unk}}^{(j)}$, we know \mathcal{A} has made the j -th query to H, and thus the values $\beta_j, \alpha_j, \delta_j$, and (D_0, D_1) are determined. For $1 \leq k' < k - 1$, since the query $(i_{k'}, c_{i_{k'}})$ to S has returned, we know the values $i_{k'}, c_{i_{k'}}, y_{i_{k'}}$ are determined, which implies $D_{k'+1}$ and $X_{k'+1}$ are determined. Also, since $i_{k-1}, c_{i_{k-1}}$ are determined, we know D_k is determined. Therefore, since $y_{i_{k-1}}$ is picked uniformly at random from \mathbb{Z}_p^* , we know $X_k = y_{i_{k-1}}$ is uniformly distributed over \mathbb{Z}_p^* independent of $(D_0, \dots, D_k, X_1, \dots, X_{k-1})$.

For $|\mathcal{I}_{\text{unk}}^{(j)}| + 1 < k \leq n$, by the definition of X_k , we know X_k is uniformly distributed over \mathbb{Z}_p^* independent of $(D_0, \dots, D_k, X_1, \dots, X_{k-1})$. Therefore, the claim holds. \square

Now, we can show the upper bound $\Pr[W_j \wedge E_j^{(1)}] \leq \frac{\ell+1}{p-1}$ by the following lemma.²

Lemma 4.2.5. *Let p be prime. Let $D_0, D_1, \dots, D_n, X_1, \dots, X_n \in \mathbb{Z}_p$ be random variables such that for all $k \in [n]$, X_k is uniform over $U_k \subseteq \mathbb{Z}_p$ and independent of $(D_0, \dots, D_k, X_1, \dots, X_{k-1})$. Then,*

$$\Pr \left[\exists i \in \{0, \dots, n\} : D_i \neq 0 \wedge D_0 + \sum_{j=1}^n D_j X_j = 0 \right] \leq \sum_{i=1}^n \frac{1}{|U_i|}.$$

Proof. For $k \in \{0, \dots, n\}$, define E_k as

$$\exists i \in \{0, \dots, k\} \text{ such that } D_i \neq 0 \wedge D_0 + \sum_{j=1}^k D_j X_j = 0.$$

We will prove the theorem using induction. It is clear that $\Pr[E_0] = 0$. For $k \geq 1$, assume $\Pr[E_{k-1}] \leq \sum_{i=1}^{k-1} \frac{1}{|U_i|}$. It holds that

$$\begin{aligned} \Pr[E_k] &= \Pr[E_k | E_{k-1}] \Pr[E_{k-1}] + \Pr[E_k | \neg E_{k-1}] \Pr[\neg E_{k-1}] \\ &\leq \Pr[E_{k-1}] + \Pr[E_k | \neg E_{k-1}] \\ &= \Pr[E_{k-1}] + \Pr[E_k | (\neg E_{k-1}) \wedge D_k \neq 0] \Pr[D_k \neq 0 | \neg E_{k-1}] \\ &\quad + \Pr[E_k | (\neg E_{k-1}) \wedge D_k = 0] \Pr[D_k = 0 | \neg E_{k-1}] \\ &\leq \Pr[E_{k-1}] + \Pr[E_k | (\neg E_{k-1}) \wedge D_k \neq 0] + \Pr[E_k | (\neg E_{k-1}) \wedge D_k = 0]. \end{aligned} \tag{4.2}$$

It is left to bound $\Pr[E_k | (\neg E_{k-1}) \wedge D_k \neq 0]$ and $\Pr[E_k | (\neg E_{k-1}) \wedge D_k = 0]$.

Suppose E_{k-1} does not occur and then we have either $D_i = 0$ for all $0 \leq i < k$ or $D_0 + \sum_{j=1}^{k-1} D_j X_j = 0$.

If $D_k = 0$, we have either $D_i = 0$ for all $0 \leq i \leq k$, or $D_0 + \sum_{j=1}^k D_j X_j = D_0 + \sum_{j=1}^{k-1} D_j X_j \neq 0$, which means E_k does not occur. Therefore, we have

$$\Pr[E_k | (\neg E_{k-1}) \wedge D_k = 0] = 0. \tag{4.3}$$

Otherwise, if $D_k \neq 0$, we know E_k occurs if and only if $D_0 + \sum_{j=1}^k D_j X_j \neq 0$. Since X_k is uniformly distributed over U_k independent of $(D_0, \dots, D_k, X_1, \dots, X_{k-1})$ given $D_k \neq 0$ and E_{k-1}

²Note that this lemma cannot be directly derived from the Schwartz-Zippel lemma by viewing $D_0 + \sum_{j=1}^n D_j X_j = 0$ as a polynomial of X_1, \dots, X_n , since we cover for example the case where D_0, D_1, \dots, D_n are adaptively chosen, i.e., each D_i can depend on X_1, \dots, X_{i-1} .

does not occur, it holds that

$$\begin{aligned}
\Pr[E_k \mid (\neg E_{k-1}) \wedge D_k \neq 0] &= \Pr \left[D_0 + \sum_{j=1}^k D_j X_j = 0 \mid (\neg E_{k-1}) \wedge D_k \neq 0 \right] \\
&= \Pr \left[X_k = \frac{D_0 + \sum_{j=1}^{k-1} D_j X_j}{D_k} \mid (\neg E_{k-1}) \wedge D_k \neq 0 \right] \\
&\leq \frac{1}{|U_i|}.
\end{aligned} \tag{4.4}$$

Therefore, from (4.2), (4.3), and (4.4), we have

$$\Pr[E_k] \leq \Pr[E_{k-1}] + \frac{1}{|U_i|} \leq \sum_{i=1}^k \frac{1}{|U_i|}.$$

Therefore, by induction, we have

$$\Pr \left[\exists i \in \{0, \dots, n\} : D_i \neq 0 \wedge D_0 + \sum_{j=1}^n D_j X_j = 0 \right] = \Pr[E_n] \leq \sum_{i=1}^n \frac{1}{|U_i|}.$$

□

4.2.2 Proof of Lemma 4.2.2

It is easier to introduce a new event F_j and show that $W_j \wedge (\neg E_j^{(1)})$ implies F_j . We will then bound $\Pr[F_j \wedge E_j^{(2)}]$. In particular, define the event F_j as

$$\forall i \in \mathcal{I}_{\text{unk}}^{(j)} : \alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} - \delta_j \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) = 0 \tag{F1}$$

$$\wedge \sum_{i \in \mathcal{I}_{\text{unk}}^{(j)}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) \neq 0, \tag{F2}$$

and we have the following lemma.

Lemma 4.2.6. *If $W_j \wedge (\neg E_j^{(1)})$ occurs, then the event F_j occurs.*

of Lemma 4.2.6. By the definition of F_j , we need only show that if $W_j \wedge (\neg E_j^{(1)})$ occurs, then (F1) and (F2) hold for j .

Suppose W_j occurs but $E_j^{(1)}$ does not occur. Since $E_j^{(1)} = E_{1,j}^{(1)} \vee \left(\bigvee_{i \in [\mathcal{I}_{\text{unk}}^{(j)}]} E_{2,(j,i)}^{(1)} \right)$, we know all of $E_{1,j}^{(1)}$ and $\{E_{2,(j,i)}^{(1)}\}_{i \in [\mathcal{I}_{\text{unk}}^{(j)}]}$ do not occur. Since the event $E_{2,(j,i)}^{(1)}$ does not occur implies

$$\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} - \delta_j \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) = 0,$$

we know (F1) holds for j .

Also, since the event $E_{1,j}^{(1)}$ does not occur, we have

$$\beta_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(j)}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) = 0.$$

Since W_j occurs, we know (W2) holds and, by the above equation, we have

$$\sum_{i \in \mathcal{I}_{\text{unk}}^{(j)}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) = \beta_j^{(0)} + \sum_{i \in [\ell]} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) \neq 0.$$

Therefore, we know (F2) holds for j . □

We also denote

$$\mathcal{D}_j := \left\{ \frac{\alpha_j^{(2i)}}{\beta_j^{(2i)}} \mid i \in \mathcal{I}_{\text{unk}}^{(j)}, \beta_j^{(2i)} \neq 0 \right\} \cup \left\{ \frac{\alpha_j^{(2i-1)}}{\beta_j^{(2i-1)}} \mid i \in \mathcal{I}_{\text{unk}}^{(j)}, \beta_j^{(2i)} = 0, \beta_j^{(2i-1)} \neq 0 \right\}.$$

We have $|\mathcal{D}_j| \leq |\{i \in \mathcal{I}_{\text{unk}}^{(j)} \mid \beta_j^{(2i)} \neq 0\} \cup \{i \in \mathcal{I}_{\text{unk}}^{(j)} \mid \beta_j^{(2i)} = 0\}| = |\mathcal{I}_{\text{unk}}^{(j)}|$.

Claim 4.2.7. *The event $F_j \wedge E_j^{(2)}$ implies $\delta_j \in \mathcal{D}_j$.*

Proof. Suppose $F_j \wedge E_j^{(2)}$ occurs but $\delta_j \notin \mathcal{D}_j$. We are going to show that $\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} = 0$ for each $i \in \mathcal{I}_{\text{unk}}^{(j)}$. Then, since F_j occurs, we know (F2) holds, which yields a contradiction, and thus the claim holds.

For $i \in \mathcal{I}_{\text{unk}}^{(j)}$, if $\beta_j^{(2i)} \neq 0$, since $\delta_j \notin \mathcal{D}_j$, we have $\delta_j \neq \frac{\alpha_j^{(2i)}}{\beta_j^{(2i)}}$, which implies $\alpha_j^{(2i)} - \delta_j \cdot \beta_j^{(2i)} \neq 0$.

Since F_j occurs, by (F1), we have $c_i = -\frac{\alpha_j^{(2i-1)} - \delta_j \cdot \beta_j^{(2i-1)}}{\alpha_j^{(2i)} - \delta_j \cdot \beta_j^{(2i)}}$. Since $E_j^{(2)}$ occurs, by (4.1), we have $\alpha_j^{(2i)} \cdot \beta_j^{(2i-1)} = \alpha_j^{(2i-1)} \cdot \beta_j^{(2i)}$, and thus

$$\begin{aligned} \beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} &= \beta_j^{(2i-1)} - \frac{\alpha_j^{(2i-1)} \cdot \beta_j^{(2i)} - \delta_j \cdot \beta_j^{(2i-1)} \cdot \beta_j^{(2i)}}{\alpha_j^{(2i)} - \delta_j \cdot \beta_j^{(2i)}} \\ &= \beta_j^{(2i-1)} - \frac{\alpha_j^{(2i)} \cdot \beta_j^{(2i-1)} - \delta_j \cdot \beta_j^{(2i-1)} \cdot \beta_j^{(2i)}}{\alpha_j^{(2i)} - \delta_j \cdot \beta_j^{(2i)}} \\ &= \beta_j^{(2i-1)} - \beta_j^{(2i-1)} = 0. \end{aligned}$$

Otherwise, suppose $\beta_j^{(2i)} = 0$. Then, if $\beta_j^{(2i-1)} = 0$, we also have $\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} = 0$. If $\beta_j^{(2i-1)} \neq 0$, since $\alpha_j^{(2i)} \cdot \beta_j^{(2i-1)} = \alpha_j^{(2i-1)} \cdot \beta_j^{(2i)} = 0$, we have $\alpha_j^{(2i)} = 0$. Since $\beta_j^{(2i)} = 0$, $\beta_j^{(2i-1)} \neq 0$, and $\delta_j \notin \mathcal{D}_j$, we have $\delta_j \neq \frac{\alpha_j^{(2i-1)}}{\beta_j^{(2i-1)}}$ and thus we have

$$\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} - \delta_j \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) = \alpha_j^{(2i-1)} - \delta_j \cdot \beta_j^{(2i-1)} \neq 0,$$

which contradicts (F1). Therefore, it is impossible that $\beta_j^{(2i)} = 0$ and $\beta_j^{(2i-1)} \neq 0$.

Therefore, from the above arguments, we have $\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} = 0$ for any $i \in \mathcal{I}_{\text{unk}}^{(j)}$, and thus $\sum_{i \in \mathcal{I}_{\text{unk}}^{(j)}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) = 0$. However, since F_j occurs, we know (F2) holds, which yields a contradiction, and thus the claim holds. \square

Note that δ_j is generated uniformly at random, independently of \mathcal{D}_j , since the latter is defined by the j -th H query. Therefore, Lemma 4.2.6 and Claim 4.2.7 yield

$$\begin{aligned} \Pr[W_j \wedge (\neg E_j^{(1)}) \wedge E_j^{(2)}] &\leq \Pr[F_j \wedge E_j^{(2)}] \\ &\leq \Pr[\delta_j \in \mathcal{D}_j] \leq \frac{|\mathcal{I}_{\text{unk}}^{(j)}|}{p-1} \leq \frac{\ell}{p-1}. \end{aligned}$$

4.2.3 Proof of Lemma 4.2.3

To conclude the analysis, we introduce yet another event, $E^{(3)}$. We will show below that $W \wedge (\neg E^{(1)}) \wedge (\neg E^{(2)})$ implies $E^{(3)}$, and thus it is enough to upper bound the probability of $E^{(3)}$ occurring. Concretely, $E^{(3)}$ is defined as follows (the definition of the following events $F_{j'}$ is given in Section 4.2.2).

Event $E^{(3)}$. For each $j_1, j_2 \in [Q_H]$ and $j_1 < j_2$, denote the event $E_{(j_1, j_2)}^{(3)}$ as

$$\exists i \in \mathcal{I}_{\text{unk}}^{(j_1)} \cap \mathcal{I}_{\text{unk}}^{(j_2)} : \alpha_{j_1}^{(2i)} \cdot \beta_{j_1}^{(2i-1)} \neq \alpha_{j_1}^{(2i-1)} \cdot \beta_{j_1}^{(2i)} \wedge \alpha_{j_2}^{(2i)} \cdot \beta_{j_2}^{(2i-1)} \neq \alpha_{j_2}^{(2i-1)} \cdot \beta_{j_2}^{(2i)}.$$

Denote $E'_{(j_1, j_2)}{}^{(3)} := E_{(j_1, j_2)}^{(3)} \wedge F_{j_1} \wedge F_{j_2}$ and $E^{(3)} := \bigvee_{j_1, j_2 \in [Q_H], j_1 < j_2} E'_{(j_1, j_2)}{}^{(3)}$.

To see why the above implication is true, assume that W indeed occurs, but both $E^{(1)}$ and $E^{(2)}$ do not occur. We now fix some $j \in \mathcal{J}$. We know W_j occurs, but both $E_j^{(1)}$ and $E_j^{(2)}$ do not occur. In particular, by the definition of $E_j^{(2)}$, we know there exists $i \in \mathcal{I}_{\text{unk}}^{(j)}$ such that $\alpha_j^{(2i)} \cdot \beta_j^{(2i-1)} \neq \alpha_j^{(2i-1)} \cdot \beta_j^{(2i)}$.

Let $i_{\min}^{(j)}$ be the smallest index in $\mathcal{I}_{\text{unk}}^{(j)}$ such that $\alpha_j^{(2i_{\min}^{(j)})} \cdot \beta_j^{(2i_{\min}^{(j)}-1)} \neq \alpha_j^{(2i_{\min}^{(j)}-1)} \cdot \beta_j^{(2i_{\min}^{(j)})}$. Since W occurs, we know $|\mathcal{J}| > \ell$. Then, since $i_{\min}^{(j)} \in \mathcal{I}_{\text{unk}}^{(j)} \subseteq [\ell]$ for each $j \in \mathcal{J}$ and $|\mathcal{J}| > \ell$, by the pigeonhole principle, we know there exists $j_1, j_2 \in \mathcal{J}$ such that $j_1 < j_2$ and $i_{\min}^{(j_1)} = i_{\min}^{(j_2)}$, which implies $E_{(j_1, j_2)}^{(3)}$ occurs. Also, since we know both $W_{j_1} \wedge (\neg E_{j_1}^{(1)})$ and $W_{j_2} \wedge (\neg E_{j_2}^{(1)})$ occur, by Lemma 4.2.6, we have F_{j_1} and F_{j_2} both occur. Therefore, we know $E'_{(j_1, j_2)}^{(3)} = E_{(j_1, j_2)}^{(3)} \wedge F_{j_1} \wedge F_{j_2}$ occurs, which implies $E^{(3)}$ occurs.

Therefore, we have

$$\Pr[W \wedge (\neg E^{(1)}) \wedge (\neg E^{(2)})] \leq \Pr[E^{(3)}] \leq \sum_{j_1, j_2 \in [Q_H], j_1 < j_2} \Pr[E'_{(j_1, j_2)}^{(3)}].$$

We now just need to bound $\Pr[E'_{(j_1, j_2)}^{(3)}]$ for any $j_1 < j_2$.

To gain insight, suppose $E'_{(j_1, j_2)}^{(3)}$ occurs. We can show that there exists $i \in \mathcal{I}_{\text{unk}}^{(j_1)} \cap \mathcal{I}_{\text{unk}}^{(j_2)}$ such that $\alpha_{j_1}^{(2i)} - \delta_{j_1} \beta_{j_1}^{(2i)} \neq 0$ and $\alpha_{j_2}^{(2i)} - \delta_{j_2} \beta_{j_2}^{(2i)} \neq 0$. Then, since F_{j_1} and F_{j_2} occur, by (F1), it holds that

$$\frac{\alpha_{j_1}^{(2i-1)} - \delta_{j_1} \cdot \beta_{j_1}^{(2i-1)}}{\alpha_{j_1}^{(2i)} - \delta_{j_1} \cdot \beta_{j_1}^{(2i)}} = c_i = \frac{\alpha_{j_2}^{(2i-1)} - \delta_{j_2} \cdot \beta_{j_2}^{(2i-1)}}{\alpha_{j_2}^{(2i)} - \delta_{j_2} \cdot \beta_{j_2}^{(2i)}}.$$

However, this can occur with only small probability since δ_{j_1} and δ_{j_2} are sampled independently. The following claim, proved in Section 4.2.4, makes this formal.

Claim 4.2.8. *For any $j_1, j_2 \in [Q_H]$ such that $j_1 < j_2$, suppose $E'_{(j_1, j_2)}^{(3)}$ occurs. Let i_{dif} be the smallest index in $\mathcal{I}_{\text{unk}}^{(j_1)} \cap \mathcal{I}_{\text{unk}}^{(j_2)}$ such that $\alpha_{j_1}^{(2i_{\text{dif}})} \cdot \beta_{j_1}^{(2i_{\text{dif}}-1)} \neq \alpha_{j_1}^{(2i_{\text{dif}}-1)} \cdot \beta_{j_1}^{(2i_{\text{dif}})}$ and $\alpha_{j_2}^{(2i_{\text{dif}})} \cdot \beta_{j_2}^{(2i_{\text{dif}}-1)} \neq \alpha_{j_2}^{(2i_{\text{dif}}-1)} \cdot \beta_{j_2}^{(2i_{\text{dif}})}$. Then, we have $\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}})} \neq 0$. Moreover, let $T = \frac{\alpha_{j_1}^{(2i_{\text{dif}}-1)} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}}-1)}}{\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}})}}$, and we have $\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})} \neq 0$ and $\delta_{j_2} = \frac{\alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})}}{\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})}}$.*

Let T and i_{dif} be the values defined in the above claim. Consider the step when δ_{j_2} is generated. We know the j_2 -th query to H has been made, and thus α_{j_2} and β_{j_2} are determined. Also, since $j_1 < j_2$, the j_1 -th query to H has returned, and thus α_{j_1} , α_{j_2} , and δ_{j_1} are determined. Therefore, we know i_{dif} and T are also determined. Thus, we know δ_{j_2} is picked uniformly at random from

\mathbb{Z}_p^* independent of i_{dif} , α_{j_1} , α_{j_2} , β_{j_1} , β_{j_2} , δ_{j_1} , and T . Then, by the above claim,

$$\begin{aligned} \Pr[E'_{(j_1, j_2)}^{(3)}] &\leq \Pr \left[\begin{array}{l} \alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}})} \neq 0 \\ \wedge \beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})} \neq 0 \end{array} \wedge \delta_{j_2} = \frac{\alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})}}{\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})}} \right] \\ &\leq \Pr \left[\delta_{j_2} = \frac{\alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})}}{\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})}} \mid \begin{array}{l} \alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}})} \neq 0 \\ \wedge \beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})} \neq 0 \end{array} \right] \\ &\leq \frac{1}{p-1}. \end{aligned}$$

4.2.4 Proof of Claim 4.2.8

This proof relies on the following simple lemma, which we first state and prove.

Lemma 4.2.9. *Let p be a prime number. Let $a, b, c, d \in \mathbb{Z}_p$ be arbitrary values such that $a \cdot d \neq c \cdot b$. Then, for any $T \in \mathbb{Z}_p$ such that $a + T \cdot b = 0$, we have $c + T \cdot d \neq 0$.*

Proof. Since $a + T \cdot b = 0$ and $a \cdot d \neq c \cdot b$, we have

$$0 = d(a + T \cdot b) = a \cdot d + T \cdot b \cdot d \neq b \cdot c + T \cdot b \cdot d = b(c + T \cdot d),$$

which implies $c + T \cdot d \neq 0$. □

Proof of Claim 4.2.8. Consider $j_1, j_2 \in [Q_H]$ such that $j_1 < j_2$. Suppose $E'_{j_1, j_2}^{(3)}$ occurs. We know the events $E_{(j_1, j_2)}^{(3)}$, F_{j_1} , and F_{j_2} occur. Since $E_{j_1, j_2}^{(3)}$ occurs, let i_{dif} be the smallest index in $\mathcal{I}_{\text{unk}}^{(j_1)} \cap \mathcal{I}_{\text{unk}}^{(j_2)}$ such that $\alpha_{j_1}^{(2i_{\text{dif}})} \cdot \beta_{j_1}^{(2i_{\text{dif}}-1)} \neq \alpha_{j_1}^{(2i_{\text{dif}}-1)} \cdot \beta_{j_1}^{(2i_{\text{dif}})}$ and $\alpha_{j_2}^{(2i_{\text{dif}})} \cdot \beta_{j_2}^{(2i_{\text{dif}}-1)} \neq \alpha_{j_2}^{(2i_{\text{dif}}-1)} \cdot \beta_{j_2}^{(2i_{\text{dif}})}$.

We first show that $\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}})} \neq 0$. Suppose $\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}})} = 0$. Since $\alpha_{j_1}^{(2i_{\text{dif}})} \cdot \beta_{j_1}^{(2i_{\text{dif}}-1)} \neq \alpha_{j_1}^{(2i_{\text{dif}}-1)} \cdot \beta_{j_1}^{(2i_{\text{dif}})}$, by Lemma 4.2.9, we know

$$\alpha_{j_1}^{(2i_{\text{dif}}-1)} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}}-1)} \neq 0.$$

Therefore, we have

$$\begin{aligned} &\alpha_{j_1}^{(2i_{\text{dif}}-1)} + c_{i_{\text{dif}}} \cdot \alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \left(\beta_{j_1}^{(2i_{\text{dif}}-1)} + c_{i_{\text{dif}}} \cdot \beta_{j_1}^{(2i_{\text{dif}})} \right) \\ &= \alpha_{j_1}^{(2i_{\text{dif}}-1)} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}}-1)} + c_{i_{\text{dif}}} \left(\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}})} \right) \\ &= \alpha_{j_1}^{(2i_{\text{dif}}-1)} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}}-1)} \neq 0. \end{aligned}$$

However, since F_{j_1} occurs, we know (F1) holds for $j = j_1$, which yields a contradiction. Thus, we have $\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \beta_{j_1}^{(2i_{\text{dif}})} \neq 0$.

Similarly, we have $\alpha_{j_2}^{(2i_{\text{dif}})} - \delta_{j_2} \beta_{j_2}^{(2i_{\text{dif}})} \neq 0$. Then, since F_{j_1} and F_{j_2} both occur, we know (F1) holds for $j = j_1$ and $j = j_2$, and thus

$$\frac{\alpha_{j_1}^{(2i_{\text{dif}}-1)} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}}-1)}}{\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}})}} = c_{i_{\text{dif}}} = \frac{\alpha_{j_2}^{(2i_{\text{dif}}-1)} - \delta_{j_2} \cdot \beta_{j_2}^{(2i_{\text{dif}}-1)}}{\alpha_{j_2}^{(2i_{\text{dif}})} - \delta_{j_2} \cdot \beta_{j_2}^{(2i_{\text{dif}})}}.$$

Denote $T = \frac{\alpha_{j_1}^{(2i_{\text{dif}}-1)} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}}-1)}}{\alpha_{j_1}^{(2i_{\text{dif}})} - \delta_{j_1} \cdot \beta_{j_1}^{(2i_{\text{dif}})}}$ and we have

$$\alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})} - \delta_{j_2} (\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})}) = 0. \quad (4.5)$$

We now show that $\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})} \neq 0$. Suppose $\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})} = 0$. Since $\alpha_{j_2}^{(2i_{\text{dif}})} \cdot \beta_{j_2}^{(2i_{\text{dif}}-1)} \neq \alpha_{j_2}^{(2i_{\text{dif}}-1)} \cdot \beta_{j_2}^{(2i_{\text{dif}})}$, by Lemma 4.2.9, we know $\alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})} \neq 0$ and

$$\alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})} - \delta_{j_2} (\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})}) = \alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})} \neq 0,$$

which contradicts (4.5). Therefore, we have

$$\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})} \neq 0,$$

and from (4.5), it holds that

$$\delta_{j_2} = \frac{\alpha_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \alpha_{j_2}^{(2i_{\text{dif}})}}{\beta_{j_2}^{(2i_{\text{dif}}-1)} - T \cdot \beta_{j_2}^{(2i_{\text{dif}})}}.$$

□

Chapter 5

EFFICIENT BLIND SIGNATURES IN THE GGM

5.1 Construction and Security

This section introduces our first blind signature scheme, BS_1 , which relies on a prime-order cyclic group and a hash function H . We describe this scheme formally in Figure 5.1. Roughly, it extends (blind) Schnorr Signatures by sending an additional group element $Y = X^y$ in the first round. Then, the signer's final response to challenge c reveals y along with $s = a + cxy$. We also note that we could consider a variant of the scheme where the signature consists of $\sigma = (A', s', y')$, where A' replaces c' .

SECURITY ANALYSIS. First off, we observe that the protocol is blind.

Theorem 5.1.1. *For a group generation algorithm GGen , the blind signature scheme $\text{BS}_1[\text{GGen}]$ is perfectly blind.*

Proof of Theorem 5.1.1. Let \mathcal{A} be an adversary playing the $\text{Blind}_{\text{BS}_1[\text{GGen}]}^{\mathcal{A}}$ game. Without loss of generality, we can assume the randomness of \mathcal{A} is fixed and \mathcal{A} always finishes both signing sessions and receives valid signatures (σ_0, σ_1) .¹

Define the view of \mathcal{A} after its execution as $\pi = (X, m_0, m_1, T_0, T_1, \sigma_0, \sigma_1)$, where $T_i := (A_i, Y_i, c_i, s_i, y_i)$, denoting the transcripts learned from interactions with the i -th signing session and $\sigma_i = (c'_i, s'_i, y'_i)$. Since the randomness of \mathcal{A} is fixed, the only randomness left is the randomness in U_1 and U_2 . Denote $\eta := (r_1^{(0)}, r_2^{(0)}, \gamma^{(0)}, r_1^{(1)}, r_2^{(1)}, \gamma^{(1)})$ as the total randomness. To prove the theorem, we need only show that the distribution of π is identical in both the case $b = 0$ and $b = 1$. We prove this by showing that for any fixed view Δ such that $\Pr[\pi = \Delta | b = 1] > 0$, there exists a unique value of the randomness η that makes $\pi = \Delta$ for the cases $b = 0$ and $b = 1$.

¹Since the output of each query to U_1 that does not return \perp is uniformly random over \mathbb{Z}_p , we know the behavior of \mathcal{A} is identical in both the case $b = 0$ and $b = 1$ before \mathcal{A} receives the valid signature (σ_0, σ_1) . Therefore, we know the probability that \mathcal{A} returns before receiving (σ_0, σ_1) or receives (\perp, \perp) after finishing both signing sessions is equal in both the case $b = 0$ and $b = 1$, which means we consider only the case where \mathcal{A} receives valid signatures.

<p><u>Algorithm Setup(1^κ) :</u> $(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$ Select $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p$ Return $par \leftarrow (p, g, H)$</p> <p><u>Algorithm KeyGen(par) :</u> $(p, g, H) \leftarrow par$ $x \leftarrow_{\\$} \mathbb{Z}_p^*$; $X \leftarrow g^x$ $sk \leftarrow x$; $pk \leftarrow X$ Return (sk, pk)</p> <p><u>Algorithm $S_1(sk)$:</u> $x \leftarrow sk$; $X \leftarrow g^x$ $a \leftarrow_{\\$} \mathbb{Z}_p$; $y \leftarrow_{\\$} \mathbb{Z}_p^*$ $A \leftarrow g^a$; $Y \leftarrow X^y$ $st^s \leftarrow (a, y, x)$; $msg_1 \leftarrow (A, Y)$ Return (st^s, msg_1)</p> <p><u>Algorithm $S_2(st^s, c)$:</u> $(a, y, x) \leftarrow st^s$ $s \leftarrow a + c \cdot y \cdot x$ Return $msg_2 \leftarrow (s, y)$</p>	<p><u>Algorithm $U_1(pk, msg_1, m)$:</u> $X \leftarrow pk$; $(A, Y) \leftarrow msg_1$ $r_1, r_2 \leftarrow_{\\$} \mathbb{Z}_p$; $\gamma \leftarrow_{\\$} \mathbb{Z}_p^*$ $Y' \leftarrow Y^\gamma$; $A' \leftarrow g^{r_1} \cdot A^\gamma \cdot Y'^{r_2}$ $c' \leftarrow H(A' \ Y' \ m)$ $c \leftarrow c' + r_2$ $st^u \leftarrow (c, c', r_1, \gamma, X, Y, A)$ Return (st^u, c)</p> <p><u>Algorithm $U_2(st^u, msg_2)$:</u> $(c, c', r_1, \gamma, X, Y, A) \leftarrow st^u$ $(s, y) \leftarrow msg_2$ If $y = 0$ or $Y \neq X^y$ or $g^s \neq A \cdot Y^c$ then return \perp $s' \leftarrow \gamma \cdot s + r_1$; $y' \leftarrow \gamma \cdot y$ Return $\sigma \leftarrow (c', s', y')$</p> <p><u>Algorithm Ver(pk, σ, m) :</u> $(c, s, y) \leftarrow \sigma$ If $y = 0$ then return 0 $Y \leftarrow X^y$; $A \leftarrow g^s \cdot Y^{-c}$ If $c \neq H(A \ Y \ m)$ then return 0 Return 1</p>
---	---

Figure 5.1: The blind signature scheme $BS_1 = BS_1[\text{GGen}]$.

For both the cases $b = 0$ and $b = 1$, we now show that $\pi = \Delta$ if and only if for each $i \in \{0, 1\}$, it holds that

$$\begin{aligned}
 \gamma^{(i)} &= y'_{b_i}{}^\Delta / y_i^\Delta, \\
 r_1^{(i)} &= s'_{b_i}{}^\Delta - \gamma^{(i)} \cdot s_i^\Delta, \\
 r_2^{(i)} &= c_i^\Delta - c'_{b_i}{}^\Delta,
 \end{aligned} \tag{5.1}$$

where the superscript $(\cdot)^\Delta$ represents the corresponding value in Δ . From the algorithms U_1 and U_2 , it is clear that the “only if” part holds. For the “if” part, suppose (5.1) holds. Since the

randomness of \mathcal{A} is fixed, the view of \mathcal{A} can differ only on the outputs c_0, c_1 from the oracle U_1 or the output (σ_0, σ_1) from the oracle U_2 . Since both signatures in Δ are valid, we have

$$A_i^\Delta = g^{s_i^\Delta} X^{\Delta - c_i^\Delta \cdot y_i^\Delta}, \quad Y_i^\Delta = X^{\Delta y_i^\Delta}, \quad (5.2)$$

$$c'_{b_i}{}^\Delta = H(g^{s'_{b_i}{}^\Delta} X^{\Delta - y'_{b_i}{}^\Delta \cdot c'_{b_i}{}^\Delta} \parallel X^{\Delta y'_{b_i}{}^\Delta} \parallel m_{b_i}^\Delta). \quad (5.3)$$

For c_i where $i \in \{0, 1\}$, suppose the values in the view of \mathcal{A} that have already determined when c_i is generated, which must include (X, m_i, A_i, Y_i) , are consistent with Δ . By (5.1), we have

$$\begin{aligned} c_i &= r_2^{(i)} + H(g^{r_1^{(i)}} A_i^{\gamma^{(i)}} Y_i^{\gamma^{(i)} \cdot r_2^{(i)}} \parallel Y_i^{\gamma^{(i)}} \parallel m_{b_i}) \\ &= r_2^{(i)} + H(g^{r_1^{(i)}} A_i^{\Delta \gamma^{(i)}} Y_i^{\Delta \gamma^{(i)} \cdot r_2^{(i)}} \parallel Y_i^{\Delta \gamma^{(i)}} \parallel m_{b_i}^\Delta) \\ &= r_2^{(i)} + H(g^{\gamma^{(i)} \cdot s_i^\Delta + r_1^{(i)}} X^{\Delta - y_i^\Delta \cdot \gamma^{(i)} \cdot (c_i^\Delta - r_2^{(i)})} \parallel X^{\Delta y_i^\Delta \cdot \gamma^{(i)}} \parallel m_{b_i}^\Delta) \\ &= r_2^{(i)} + H(g^{s'_{b_i}{}^\Delta} X^{\Delta - y'_{b_i}{}^\Delta \cdot c'_{b_i}{}^\Delta} \parallel X^{\Delta y'_{b_i}{}^\Delta} \parallel m_{b_i}^\Delta) \\ &= r_2^{(i)} + c'_{b_i}{}^\Delta = c_i^\Delta, \end{aligned}$$

where the third equality is due to (5.2), the fourth equality is due to (5.1), and the final equality is due to (5.3). Then, consider the step when (σ_0, σ_1) is output. Suppose the current view, which contains T_i , is consistent with Δ . By (5.1), we have

$$\begin{aligned} y'_{b_i} &= \gamma^{(i)} \cdot y_i = \gamma^{(i)} \cdot y_i^\Delta = y'_{b_i}{}^\Delta, \\ s'_{b_i} &= r_1^{(i)} + \gamma^{(i)} \cdot s_i = r_1^{(i)} + \gamma^{(i)} \cdot s_i^\Delta = s'_{b_i}{}^\Delta, \\ c'_{b_i} &= c_i - r_2^{(i)} = c_i^\Delta - r_2^{(i)} = c'_{b_i}{}^\Delta. \end{aligned}$$

which implies $(\sigma_0, \sigma_1) = (\sigma_0^\Delta, \sigma_1^\Delta)$. Therefore, by induction, if (5.1) holds, we know $\pi = \Delta$. \square

Our main result shows OMUF security of BS_1 in the *generic-group model* (GGM) following Shoup's original formalization [118], which encodes every group element with a random label. To this end, we present in Figure 5.2 a game describing a GGM-version of OMUF security for BS_1 , adapting the one from Section 3.3. We also define a corresponding advantage $\text{Adv}_{BS_1[\text{GGen}]}^{\text{omuf-ggm}}(\mathcal{A}, \kappa)$ to measure the probability that \mathcal{A} wins the game. Note that to keep notation homogenous, it is convenient to allow the game to depend on \mathbb{G} , although the game itself only makes use of the order of the group. The game also models the hash function H as a random oracle, to which the adversary is given oracle access.

<p><u>Game OMUF-GGM^A_{BS₁[GGen]}(κ) :</u></p> <p>$(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$</p> <p>$\text{sid} \leftarrow 0 ; \ell \leftarrow 0 ; \mathcal{I}_{\text{fin}} \leftarrow \emptyset ; \text{Cur} \leftarrow \emptyset$</p> <p>$\Xi \leftarrow () ; T \leftarrow ()$</p> <p>$\{(m_k, \sigma_k)\}_{k \in [\ell+1]} \leftarrow_{\\$} \mathcal{A}^{\Pi, S_1, S_2, H}(p, \Phi(1), \Phi(x))$</p> <p>If $\exists k_1 \neq k_2$ such that $(m_{k_1}, \sigma_{k_1}) = (m_{k_2}, \sigma_{k_2})$ then</p> <p style="padding-left: 20px;">Return 0</p> <p>If $\exists k \in [\ell + 1]$ such that $y_k^* = 0$</p> <p style="padding-left: 20px;">or $c_k \neq H(\Phi(s_k - c_k \cdot y_k \cdot x) \parallel \Phi(y_k \cdot x) \parallel m_i)$</p> <p>where $(c_k, s_k, y_k) = \sigma_k$ then return 0</p> <p>Return 1</p> <p><u>Oracle $\Phi(v)$:</u></p> <p>If $v \in \text{Cur}$ then return $\Xi(v)$</p> <p>$\Xi(v) \leftarrow_{\\$} \{0, 1\}^{\log(p)} \setminus \Xi(\text{Cur})$</p> <p>$\text{Cur} \leftarrow \text{Cur} \cup \{v\}$</p> <p>Return $\Xi(v)$</p> <p><u>Oracle $\Pi(\xi, \xi', b)$:</u></p> <p>If $\exists v, v' \in \text{Cur}$ such that $\xi = \Xi(v)$ and $\xi' = \Xi(v')$ then</p> <p style="padding-left: 20px;">Return $\Phi(v + (-1)^b v')$</p> <p>Else return \perp</p>	<p><u>Oracle S_1 :</u></p> <p>$\text{sid} \leftarrow \text{sid} + 1$</p> <p>$a_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p ; y_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p^*$</p> <p>$\text{st}_{\text{sid}}^s \leftarrow (a_{\text{sid}}, y_{\text{sid}})$</p> <p>$\text{msg}_1 \leftarrow (\Phi(a_{\text{sid}}), \Phi(x y_{\text{sid}}))$</p> <p>Return $(\text{sid}, \text{msg}_1)$</p> <p><u>Oracle $S_2(i, c_i)$:</u></p> <p>If $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ then</p> <p style="padding-left: 20px;">Return \perp</p> <p>$(a_i, y_i) \leftarrow \text{st}_i^s$</p> <p>$s_i \leftarrow a_i + c_i \cdot y_i \cdot x$</p> <p>$\text{msg}_2 \leftarrow (s_i, y_i)$</p> <p>$\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$</p> <p>$\ell \leftarrow \ell + 1$</p> <p>Return msg_2</p> <p><u>Oracle $H(\text{str})$:</u></p> <p>If $T(\text{str}) = \perp$ then</p> <p style="padding-left: 20px;">$T(\text{str}) \leftarrow_{\\$} \mathbb{Z}_p$</p> <p>Return $T(\text{str})$</p>
---	--

Figure 5.2: The OMUF security game in GGM for the blind signature scheme $\text{BS}_1[\text{GGen}]$.

The following theorem states our main result in the form of a reduction to WFROS and is proved in Section 5.2.

Theorem 5.1.2 (OMUF Security of BS_1). *Let GGen be a group generation algorithm. For any adversary \mathcal{A} for the OMUF-GGM^{BS₁[GGen]}(κ) game making at most Q_Π queries to Π , Q_{S_1} queries to S_1 , and Q_H queries to the random oracle H , there exists an adversary \mathcal{B} for the WFROS $_{Q_{S_1}, p}$ problem, where p denotes the size of the group generated by $\text{GGen}(1^\kappa)$, making at most $Q_H + Q_{S_1} + 1$*

queries to the random oracle H such that

$$\text{Adv}_{\text{BS}_1[\text{GGen}]}^{\text{omuf-ggm}}(\mathcal{A}, \kappa) \leq \text{Adv}_{Q_{S_1}, p}^{\text{wfros}}(\mathcal{B}) + \frac{Q_\Phi(Q_\Phi + 2Q_H + 2Q_{S_1} + 2)}{p - (1 + Q_{S_1} + Q_\Phi^2)},$$

where Q_Φ is the maximum number of queries to Φ during the game OMUF-GGM, and we have $Q_\Phi = Q_\Pi + 4Q_{S_1} + 4$.

By Theorem 4.1.1, we have the following corollary.

Corollary 5.1.3. *Let GGen be a group generation algorithm. For any adversary \mathcal{A} playing game $\text{OMUF-GGM}^{\text{BS}_1[\text{GGen}]}(\lambda)$ making at most Q_Π queries to Π , Q_{S_1} queries to S_1 , and Q_H queries to the random oracle H , we have*

$$\text{Adv}_{\text{BS}_1[\text{GGen}]}^{\text{omuf-ggm}}(\mathcal{A}, \kappa) \leq \frac{2Q_\Phi(Q_\Phi + 2Q_H + 2Q_{S_1} + 2)}{2^\kappa - (1 + Q_{S_1} + Q_\Phi^2)},$$

where $Q_\Phi = Q_\Pi + 4Q_{S_1} + 4$.

We note in particular that the concrete security of BS_1 in the GGM is comparable to that of the discrete logarithm problem, in that $Q_\Phi = \Omega(\min\{\sqrt{p}, p/Q_H, p/Q_{S_1}\})$ is necessary to break security with constant probability.

5.2 Proof of OMUF (Theorem 5.1.2)

Let us fix an adversary \mathcal{A} that makes (without loss of generality) exactly Q_Π queries to Π , Q_{S_1} queries to S_1 , and Q_H queries to the random oracle H . Without loss of generality, assume it also makes exactly one query (i, c_i) to S_2 for each $i \in [Q_{S_1}]$. Then, after \mathcal{A} returns, we know $\ell = Q_{S_1}$ and $\mathcal{I}_{\text{fin}} = [Q_{S_1}]$. Also, it is clear that the overall number of queries to Φ in $\text{OMUF-GGM}_{\text{BS}_1}^{\mathcal{A}}$ is at most $Q_\Phi := Q_\Pi + 4Q_{S_1} + 4$.

We prove the theorem by going through a series of games, from Game_0 to Game_4 , where Game_0 is the $\text{OMUF-GGM}_{\text{BS}_1}^{\mathcal{A}}$ game and Game_4 is an intermediate game that enables an easier reduction to WFROS. Here, however, we first introduce Game_4 and Lemma 5.2.1 and then discuss the reduction to WFROS, which is the core of the proof. We leave the definition of the intermediate games between Game_0 to Game_4 to the proof of Lemma 5.2.1. The game-hopping argument is non-trivial, but it follows the same blueprint as in [12] and is hence deferred to Section 5.2.1.

<p><u>Game Game₄:</u> $(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$ $\text{sid} \leftarrow 0; \ell \leftarrow 0; \mathcal{S} \leftarrow \emptyset; \text{Cur} \leftarrow \emptyset$ $\Xi \leftarrow (); T \leftarrow ()$ $\{(m_k, \sigma_k)\}_{k \in [\ell+1]} \leftarrow_{\mathcal{S}} \mathcal{A}^{\Pi, S_1, S_2, H}(p, \Phi(1), \Phi(X))$ If $\exists k_1 \neq k_2$ such that $(m_{k_1}, \sigma_{k_1}) = (m_{k_2}, \sigma_{k_2})$ then Return 0 If $\exists k \in [\ell + 1]$ such that $y_k^* = 0$ or $c_k \neq H(\Phi(s_k - c_k \cdot y_k \cdot X) \parallel \Phi(y_k \cdot X) \parallel m_i)$ where $(c_k, s_k, y_k) = \sigma_k$ then return 0 Return 1</p> <p><u>Oracle $\Phi(P)$:</u> If $\exists P' \in \text{Cur}$ such that $P =_L P'$ then Return $\Xi(P')$ $\Xi(P) \leftarrow_{\mathcal{S}} \{0, 1\}^{\lceil \log(p) \rceil} \setminus \Xi(\text{Cur})$ $\text{Cur} \leftarrow \text{Cur} \cup \{P\}$ Return $\Xi(P)$</p> <p><u>Oracle $\Pi(\xi, \xi', b)$:</u> If $\exists P, P' \in \text{Cur}$ such that $\xi = \Xi(P)$ and $\xi' = \Xi(P')$ then Return $\Phi(P + (-1)^b P')$ Else return \perp</p>	<p><u>Oracle S_1 :</u> $\text{sid} \leftarrow \text{sid} + 1$ $\text{msg}_1 \leftarrow (\Phi(A_{\text{sid}}), \Phi(Y_{\text{sid}}))$ Return $(\text{sid}, \text{msg}_1)$</p> <p><u>Oracle $S_2(i, c_i)$:</u> If $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ then return \perp $s_i \leftarrow_{\mathcal{S}} \mathbb{Z}_p; y_i \leftarrow_{\mathcal{S}} \mathbb{Z}_p^*$ $R_1 \leftarrow A_i + c_i Y_i - s_i$ $R_2 \leftarrow Y_i - y_i X$ $L \leftarrow L \cup \{R_1, R_2\}$ $\text{msg}_2 \leftarrow (s_i, y_i)$ If $\exists P_1, P_2 \in \text{Cur}$ such that $P_1 \neq P_2$ and $P_1 =_L P_2$ then abort game $\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$ $\ell \leftarrow \ell + 1$ Return msg_2</p> <p><u>Oracle $H(\text{str})$:</u> If $T(\text{str}) = \perp$ then $T(\text{str}) \leftarrow_{\mathcal{S}} \mathbb{Z}_p$ Return $T(\text{str})$</p>
---	--

Figure 5.3: The definition of Game₄. The symbols P and P' denote polynomials over variables X , $\{A_i, Y_i\}_{i \in [\text{sid}]}$. Also, a new equality notation, “ $=_L$ ”, is used. We say $P_1 =_L P_2$ if and only if $P_1 - P_2$ can be represented as a linear combination of polynomials in L .

DEFINITION OF Game₄. The pseudocode description of Game₄ is given in Figure 5.3. The main difference from OMUF-GGM_{BS₁}^A is that the encoding oracle Φ takes as input a polynomial instead of an integer in \mathbb{Z}_p . (Note that the adversary cannot query Φ directly, and thus this difference is not directly surfaced.) This essentially captures the algebraic core of our proof.

Also, for a valid query (i, c_i) to S_2 , the output values (s_i, y_i) are directly sampled uniformly from $\mathbb{Z}_p \times \mathbb{Z}_p^*$. Furthermore, when this happens, two polynomials, $R_1 = A_i + c_i \cdot Y_i - s_i$ and $R_2 = Y_i - y_i \cdot X$, are recorded in the set L . Then, in the encoding oracle Φ , two polynomials, P_1 and P_2 , are considered to differ if and only if $P_1 \neq_L P_2$, where $P_1 =_L P_2$ means that $P_1 - P_2$ can be generated as a linear combination of polynomials in L . Still, $P_1 \neq_L P_2$ could occur when queries P_1 and P_2 are made to Φ , but they becomes equal (in the sense of “ $=_L$ ”) after L is updated. The game aborts when this happens. Overall, we prove the following lemma in Section 5.2.1.

Lemma 5.2.1. $\text{Adv}_{\text{BS}_1[\mathbb{G}]}^{\text{omuf-ggm}}(\mathcal{A}, \lambda) \leq \Pr[\text{Game}_4^{\mathcal{A}} = 1] + \frac{Q_\Phi^2}{p - (1 + Q_{S_1} + Q_\Phi^2)}$.

REDUCTION TO WFROS. The core of the proof is to relate the probability of the adversary \mathcal{A} winning Game_4 with the advantage of an adversary \mathcal{B} winning the WFROS problem, as stated in the following lemma. The proof is given in Section 5.2.3.

Lemma 5.2.2. *For every λ , there exists an adversary \mathcal{B} for the $\text{WFROS}_{Q_{S_1}, p}$ problem, where $p = |\mathbb{G}_\lambda|$, making at most $Q_H + Q_{S_1} + 1$ queries to H such that*

$$\Pr[\text{Game}_4^{\mathcal{A}} = 1] \leq \text{Adv}_{Q_{S_1}, p}^{\text{wfros}}(\mathcal{B}) + \frac{(2Q_\Phi + 1)(Q_H + Q_{S_1} + 1)}{p - Q_\Phi}. \quad (5.4)$$

The statement of Theorem 5.1.2 follows by combining Lemmas 5.2.1 and 5.2.2.

5.2.1 Proof of Lemma 5.2.1

We prove the lemma by going through a series of games.

Game₀^A: This is OMUF-GGM_{BS₁}^A (Figure 5.2).

Game₁^A: This is defined in Figure 5.4 that only contains the dashed box. We introduce variables $X, A_1, Y_1, \dots, A_{Q_{S_1}}, Y_{Q_{S_1}}$ in $\text{Game}_1^{\mathcal{A}}$. Each variable is assigned a value, that is, X is assigned x , A_i is assigned a_i , and Y_i is assigned $y_i \cdot x$. The input to Φ is a polynomial P of variables $X, \{A_i, Y_i\}_{i \in [Q_{S_1}]}$ over \mathbb{Z}_p instead of a single value $v \in \mathbb{Z}_p$ and the set Cur is a set of polynomials. Also, in Φ we check the equality of two polynomials by its evaluation on the assigned values, which is denoted by $=_{\text{eval}}$ (see Definition 5.2.3).

Definition 5.2.3. *For two polynomial P and P' of the variables X_1, \dots, X_n over a field \mathcal{F} , suppose each X_i is assigned with a value $x_i \in \mathcal{F}$. We say $P =_{\text{eval}} P'$ if and only if $P(X_1 = x_1, \dots, X_n = x_n) = P'(X_1 = x_1, \dots, X_n = x_n)$.*

<p>Game $\text{Game}_1^{\mathcal{A}}$, $\text{Game}_2^{\mathcal{A}}$, $\text{Game}'_2^{\mathcal{A}}$:</p> <p>$(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$; $x \leftarrow_{\\$} \mathbb{Z}_p^*$; assign x to variable X</p> <p>$\text{sid} \leftarrow 0$; $\ell \leftarrow 0$; $\mathcal{I}_{\text{fin}} \leftarrow \emptyset$; $\Xi \leftarrow ()$; $T \leftarrow ()$</p> <p>$\text{Cur} \leftarrow \emptyset$; $L \leftarrow \emptyset$</p> <p>$\{(m_k, \sigma_k)\}_{k \in [\ell+1]} \leftarrow_{\\$} \mathcal{A}^{\Pi, S_1, S_2, H}(p, \Phi(1), \Phi(\mathsf{X}))$</p> <p>If $\exists k_1 \neq k_2$ such that $(m_{k_1}, \sigma_{k_1}) = (m_{k_2}, \sigma_{k_2})$ then</p> <p style="padding-left: 20px;">Return 0</p> <p>If $\exists k \in [\ell + 1]$ such that $y_k^* = 0$</p> <p style="padding-left: 20px;">or $c_k \neq H(\Phi(s_k - c_k \cdot y_k \cdot \mathsf{X}) \parallel \Phi(y_k \cdot \mathsf{X}) \parallel m_i)$</p> <p style="padding-left: 20px;">where $(c_k, s_k, y_k) = \sigma_k$ then return 0</p> <p>Return 1</p> <p><u>Oracle $\Phi(P)$:</u></p> <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p>If $\exists P' \in \text{Cur}$ such that $P =_{\text{eval}} P'$ and $P \neq_L P'$ then abort game</p> </div> <div style="border: 1px dashed black; padding: 5px; margin: 5px 0;"> <p>If $\exists P' \in \text{Cur}$ such that $P =_{\text{eval}} P'$ then</p> <p style="padding-left: 20px;">Return $\Xi(P')$</p> </div> <div style="border: 1px solid black; padding: 5px; margin: 5px 0; background-color: #f0f0f0;"> <p>If $\exists P' \in \text{Cur}$ such that $P =_L P'$ then</p> <p style="padding-left: 20px;">Return $\Xi(P')$</p> </div> <p>$\Xi(P) \leftarrow_{\\$} \{0, 1\}^{\log(p)} \setminus \Xi(\text{Cur})$</p> <p>$\text{Cur} \leftarrow \text{Cur} \cup \{P\}$</p> <p>Return $\Xi(P)$</p> <p><u>Oracle $\Pi(\xi, \xi', b)$:</u></p> <p>If $\exists P, P' \in \text{Cur}$ such that $\xi = \Xi(P)$ and $\xi' = \Xi(P')$ then</p> <p style="padding-left: 20px;">Return $\Phi(P + (-1)^b P')$</p> <p>Else return \perp</p>	<p><u>Oracle S_1 :</u></p> <p>$\text{sid} \leftarrow \text{sid} + 1$</p> <p>$a_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p$; $y_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p^*$</p> <p>$\text{st}_{\text{sid}}^s \leftarrow (a_{\text{sid}}, y_{\text{sid}})$</p> <p>Assign a_{sid} to variable A_{sid}</p> <p>Assign $y_{\text{sid}} \cdot x$ to variable Y_{sid}</p> <p>$\text{msg}_1 \leftarrow (\Phi(\mathsf{A}_{\text{sid}}), \Phi(\mathsf{Y}_{\text{sid}}))$</p> <p>Return $(\text{sid}, \text{msg}_1)$</p> <p><u>Oracle $S_2(i, c_i)$:</u></p> <p>If $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ then return \perp</p> <p>$(a_i, y_i) \leftarrow \text{st}_i^s$</p> <p>$s_i \leftarrow a_i + c_i \cdot y_i \cdot x$</p> <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p>$R_1 \leftarrow \mathsf{A}_i + c_i \mathsf{Y}_i - s_i$ $R_2 \leftarrow \mathsf{Y}_i - y_i \mathsf{X}$ $L \leftarrow L \cup \{R_1, R_2\}$</p> </div> <p>$\text{msg}_2 \leftarrow (s_i, y_i)$</p> <p>$\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$</p> <p>$\ell \leftarrow \ell + 1$</p> <p>Return msg_2</p> <p><u>Oracle $H(\text{str})$:</u></p> <p>If $T(\text{str}) = \perp$ then</p> <p style="padding-left: 20px;">$T(\text{str}) \leftarrow_{\\$} \mathbb{Z}_p$</p> <p>Return $T(\text{str})$</p>
--	--

Figure 5.4: The definition for $\text{Game}_1^{\mathcal{A}}$, $\text{Game}_2^{\mathcal{A}}$, and $\text{Game}'_2^{\mathcal{A}}$, where $\text{Game}_1^{\mathcal{A}}$ only contains the dashed box, $\text{Game}_2^{\mathcal{A}}$ contains all but the gray box, and $\text{Game}'_2^{\mathcal{A}}$ contains all but the dashed box.

For convenience, we also have $P =_{\text{eval}} P(X_1 = x_1, \dots, X_n = x_n)$.

It is easy to check that $=_{\text{eval}}$ is an equivalence relation over the polynomials of the variables X_1, \dots, X_n .

We first show that the oracle Φ in Game_1^A is well-defined, that is, for each query P to Φ , there exists at most one $P' \in \text{Cur}$ such that $P =_{\text{eval}} P'$. Suppose there exists $P', P'' \in \text{Cur}$ such that $P' \neq P''$, $P' =_{\text{eval}} P =_{\text{eval}} P''$. Suppose P'' is added to Cur after P' . Consider the query to Φ during which P'' is added to Cur . Since P' is already in Cur when P'' is added, we have $P' \neq_{\text{eval}} P''$, which yields a contradiction. Therefore, for each query to Φ , if there exists $P' \in \text{Cur}$ such that $P =_{\text{eval}} P'$, then P' is the unique polynomial in Cur such that $P =_{\text{eval}} P'$.

We now show that the views of the adversary in Game_0 and Game_1 are identical. Define an intermediate game Game'_1^A such that it is identical to Game_1^A except each the polynomial P appear in the game is replaced by its evaluation value $P(X = x, A_1 = a_1, Y_1 = y_1 \cdot x, \dots, A_{\text{sid}} = a_{\text{sid}}, Y_{\text{sid}} = y_{\text{sid}} \cdot x)$. It is clear that Game'_1^A is identical to Game_0^A . Also, since in the oracle Φ in Game_1^A , a polynomial P is considered equal or not equal to another polynomials by its evaluation value, the view of the adversary in Game_1 and Game'_1 are identical. Thus, we know the views of the adversary in Game_0 and Game_1 are identical, which implies

$$\Pr[\text{Game}_0^A = 1] = \Pr[\text{Game}_1^A = 1]. \quad (5.5)$$

Game_2^A : This is defined in Figure 5.4 by ignoring the graybox. A set L is introduced to record the information leaked to the adversary by S_2 . For the query (i, c_i) to S_2 , polynomials $R_1 = A_i + c_i Y_i - s_i$ and $R_2 = Y_i - y_i X$ are added to L . Suppose L is also recorded in Game_1^A . In Game_1^A , define the event E_1 as after an query P to Φ is made,

$$\exists P' \in \text{Cur} \text{ such that } P =_{\text{eval}} P' \text{ and } P \neq_L P'.$$

Then, Game_2^A is identical to Game_1^A except it aborts when E_1 occurs and we have

$$\Pr[\text{Game}_1^A = 1] \leq \Pr[\text{Game}_2^A = 1] + \Pr[E_1], \quad (5.6)$$

To bound $\Pr[E_1]$, for each $j \in [Q_\Phi]$, we denote the event $E_{1,j}$ in Game_1^A as during the j -th query to Φ

$$\exists P' \in \text{Cur} \text{ such that } P_j =_{\text{eval}} P' \text{ and } P_j \neq_L P'.$$

Then, we have $E_1 = \bigvee_{j \in [Q_\Phi]} E_{1,j}$. Denote $E'_{1,j} := E_{1,j} \wedge_{i \in [j]} (\neg E_{1,i})$. We now bound $\Pr[E'_{1,j}]$ for each $j \in [Q_\Phi]$.

We now fix a certain $j \in [Q_\Phi]$. Consider the step when the j -th query to Φ is made during Game_1^A . Denote the transcripts between the oracles and adversaries when the j -th query to Φ is made as π_j , which contains $\Phi(1), \Phi(X)$, and all the inputs and outputs of the queries to S_1, S_2, Π , and H made before the j -th query to Φ . For a certain transcript $\pi_j = \Delta$, for $1 \leq k \leq j$, denote the k -th query to Φ in Δ as P_k^Δ . From the transcript Δ , one can compute the set $\mathcal{I}_{\text{fin}}, \text{Cur}$, and L at the step when the j -th query to Φ is been made. Denote them by $\mathcal{I}_{\text{fin}}^\Delta, \text{Cur}^\Delta$, and L^Δ . For each $i \in \mathcal{I}_{\text{fin}}^\Delta$, denote the input and output of the query to the S_2 for the session i in the transcript Δ as c_i^Δ and (s_i^Δ, y_i^Δ) . Also, from the transcript π_j , one can tell whether $E_{1,k}$ occurs or not for $k \in [j-1]$, since the event $E_{1,k}$ occurs if and only if $P_k \neq_L P'$ for all $P' \in \text{Cur}$ but P_k is not added to Cur . Denote the value of sid when the j -th query to Φ is made as sid^Δ .

Denote \mathcal{T}_j as the set of all transcripts Δ such that $\Pr[\pi_j = \Delta] > 0$ and none of $\{E_{1,k}\}_{k \in [j]}$ occurs given $\pi_j = \Delta$. We just need to bound $\Pr[E'_{1,j} | \pi_j = \Delta]$ for each $\Delta \in \mathcal{T}_j$.

We now fix a certain $\Delta \in \mathcal{T}_j$. For any polynomial P , denote the event F_P as $P =_{\text{eval}} P_j$ and $P \neq_L P_j$. Then we know $E'_{1,j}$ implies one of $\{F_P\}_{P \in \text{Cur}^\Delta}$ occurs and we have

$$\Pr[E'_{1,j} | \pi_j = \Delta] \leq \Pr \left[\bigvee_{P \in \text{Cur}^\Delta} F_P | \pi_j = \Delta \right].$$

Therefore, it is left to bound $\Pr[F_P]$ for each $P \in \text{Cur}^\Delta$.

We now fix a certain $\hat{P} \in \text{Cur}^\Delta$. Since P_j^Δ and L^Δ are fixed in Δ , we can directly check whether $\hat{P} =_{L^\Delta} P_j^\Delta$ or not. If $\hat{P} =_{L^\Delta} P_j^\Delta$, then we have $\Pr[F_{\hat{P}}] = 0$. Therefore, we can assume $\hat{P} \neq_{L^\Delta} P_j^\Delta$. Then, we only need to bound the probability of $\hat{P} =_{\text{eval}} P_j^\Delta$. Since we fix $\pi_j = \Delta$, the only randomness here is the values assigned to the random variables $X, \{A_i, Y_i\}_{i \in [\text{sid}^\Delta]}$. Denote the values as $\boldsymbol{\eta} := (x, a_1, y_1 \cdot x, \dots, a_{\text{sid}^\Delta}, y_{\text{sid}^\Delta} \cdot x) \in \mathbb{Z}_p^{1+2\text{sid}^\Delta}$, where $x, \{a_i, y_i\}_{i \in [\text{sid}^\Delta]}$ are random variables sampled in the game, and we have $P =_{\text{eval}} P(X = \eta_1, \{A_i = \eta_{2i}, Y_i = \eta_{2i+1}\}_{i \in [\text{sid}^\Delta]})$.

To bound $\Pr[\hat{P} =_{\text{eval}} P_j^\Delta | \pi_j = \Delta]$, we first introduce Lemma 5.2.4 below. Then the proof structure can be described as follows. We first define a sequence of polynomials $D_0, D_1, \dots, D_m, B_1, \dots, B_{q+1}$ over variables $X, \{A_i, Y_i\}_{i \in [\text{sid}^\Delta]}$ such that $B_{q+1} := \hat{P} - P_j^\Delta$. Then, we try to apply Lemma 5.2.4 to bound the probability by showing $\boldsymbol{\eta}$ is uniformly distributed over $\mathcal{C}, \text{Zero}(B_{q+1}) \cap \mathcal{C} \neq \emptyset$, and $B_{q+1} \notin \text{Span}(\{1, B_1, \dots, B_q\})$ given $\pi_j = \Delta$, where \mathcal{C} is defined in Lemma 5.2.4.

Lemma 5.2.4 (Lemma 1 in [12]). *Let $D_1, \dots, D_m, B_1, \dots, B_{q+1}$ be polynomials in $\mathbb{Z}_p[X_1, \dots, X_n]$*

of degree 1. Let

$$\mathcal{C} := \left(\bigcap_{i \in [q]} \text{Zero}(B_i) \right) \setminus \left(\bigcup_{i \in [m]} \text{Zero}(D_i) \right),$$

where $\text{Zero}(P)$ means the zero set of P . Assume $\text{Zero}(B_{q+1}) \cap \mathcal{C} \neq \emptyset$ and $B_{q+1} \notin \text{Span}(\{1, B_1, \dots, B_q\})$. If \mathbf{x} is picked uniformly at random from \mathcal{C} then

$$\frac{p-m}{p^2} \leq \Pr[B_{q+1}(\mathbf{x}) = 0] \leq \frac{1}{p-m}.$$

Let $m := \text{sid}^\Delta + 1 + |\text{Cur}^\Delta|(|\text{Cur}^\Delta| - 1)$. Denote $D_1 := X$ and $D_{i+1} := Y_i$ for $i \in [\text{sid}^\Delta]$. For each $P, P' \in \text{Cur}^\Delta$ such that $P \neq P'$, denote $D_{P,P'} := P - P'$. We can relabel $\{D_{P,P'}\}_{P,P' \in \text{Cur}^\Delta, P \neq P'}$ to $D_{\text{sid}^\Delta+2}, \dots, D_m$.

Let $q := 2|\mathcal{I}_{\text{fin}}^\Delta|$. For each $i \in \mathcal{I}_{\text{fin}}^\Delta$, denote

$$B_{(i,1)} := A_i + c_i^\Delta Y_i - s_i^\Delta, \quad B_{(i,2)} := Y_i - y_i^\Delta X.$$

We can relabel $\{B_{(i,1)}, B_{(i,2)}\}_{i \in \mathcal{I}_{\text{fin}}^\Delta}$ to B_1, \dots, B_q and denote $B_{q+1} := \hat{P} - P_j^\Delta$. Here one thing to notice is that we have $L^\Delta = \{B_1, \dots, B_q\}$.

Denote $\mathcal{C} := \left(\bigcap_{i \in [q]} \text{Zero}(B_i) \right) \setminus \left(\bigcup_{i \in [m]} \text{Zero}(D_i) \right)$ and we have the following claim. The proof of the claim is deferred to Section 5.2.2.

Claim 5.2.5. *In Game_1^A , for any $\Delta \in \mathcal{T}_j$, given $\pi_j = \Delta$, we have $\boldsymbol{\eta}$ is uniformly distributed over \mathcal{C} .*

We now continue to show that $\text{Zero}(B_{q+1}) \cap \mathcal{C} \neq \emptyset$. If $\text{Zero}(B_{q+1}) \cap \mathcal{C} = \emptyset$, since by the above claim $\boldsymbol{\eta}$ must be in \mathcal{C} given $\pi_j = \Delta$, we know $B_{q+1} =_{\text{eval}} B_{q+1}(\boldsymbol{\eta}) \neq 0$, which implies $\Pr[\hat{P} =_{\text{eval}} P_j | \pi_j = \Delta] = 0$. Therefore, we only need to consider the case when $\text{Zero}(B_{q+1}) \cap \mathcal{C} \neq \emptyset$.

We then show that $B_{q+1} \notin \text{Span}(\{1, B_1, \dots, B_q\})$. Since $\hat{P} \neq_{L^{\text{Cur}}} P_j^\Delta$ and $L^\Delta = \{B_1, \dots, B_q\}$, we know $B_{q+1} \notin \text{Span}(\{B_1, \dots, B_q\})$. If $B_{q+1} \in \text{Span}(\{1, B_1, \dots, B_q\})$, we know there exists a constant $\delta \in \mathbb{Z}_p$ such that $\delta \neq 0$ and $B_{q+1} + \delta \in \text{Span}(\{B_1, \dots, B_q\})$. Let $B' = B_{q+1} + \delta$. Then, we have for any $\boldsymbol{\eta}_0 \in \mathcal{C}$, $B'(\boldsymbol{\eta}) = 0$ and thus $B_{q+1}(\boldsymbol{\eta}) = B'(\boldsymbol{\eta}) - \delta = -\delta \neq 0$, which means $\text{Zero}(B_{q+1}) \cap \mathcal{C} = \emptyset$. This contradicts with the above argument that $\text{Zero}(B_{q+1}) \cap \mathcal{C} \neq \emptyset$. Therefore, we have $B_{q+1} \notin \text{Span}(\{1, B_1, \dots, B_q\})$.

Then, by the above claim, we can apply Lemma 5.2.4 here and we have

$$\Pr[\hat{P} =_{\text{eval}} P_j^\Delta | \pi_j = \Delta] = \Pr[B_{q+1}(\boldsymbol{\eta}) = 0 | \pi_j = \Delta] \leq \frac{1}{p-m}.$$

Since $m = \text{sid}^\Delta + 1 + |\text{Cur}^\Delta|(|\text{Cur}^\Delta - 1|) \leq 1 + Q_{S_1} + Q_\Phi^2$, we have

$$\begin{aligned} \Pr[E'_{1,j}] &= \sum_{\Delta \in \mathcal{T}_j} \Pr[E'_{1,j} \wedge \pi_j = \Delta] \\ &= \sum_{\Delta \in \mathcal{T}_j} \Pr[\pi_j = \Delta] \sum_{\hat{P} \in \text{Cur}^\Delta} \Pr[F_{\hat{P}} \mid \pi_j = \Delta] \leq \frac{Q_\Phi}{p - (1 + Q_{S_1} + Q_\Phi^2)}. \end{aligned}$$

Therefore, we have $\Pr[E_1] = \sum_{j \in [Q_\Phi]} \Pr[E'_{1,j}] \leq \frac{Q_\Phi^2}{p - (1 + Q_{S_1} + Q_\Phi^2)}$ and by (5.6)

$$\Pr[\text{Game}_1^A = 1] \leq \Pr[\text{Game}_2^A = 1] + \frac{Q_\Phi^2}{p - (1 + Q_{S_1} + Q_\Phi^2)}. \quad (5.7)$$

Game₂^{'A}: This is defined in Figure 5.4 by ignoring the dashed box. The only difference between Game₂^A and Game₂^{'A} is that in the oracle Φ the condition “ $\exists P' \in \text{Cur}$ such that $P =_{\text{eval}} P'$ ” is changed to “ $\exists P' \in \text{Cur}$ such that $P =_L P'$ ”. We will show that $P =_{\text{eval}} P'$ is equivalent to $P =_L P'$ here in Game₂, and thus we know the view of adversary are identical in these two games.

In Game₂^A, consider an query P to the oracle Φ . Let P' be an arbitrary polynomial in Cur . Consider the step when the condition “ $\exists P' \in \text{Cur}$ such that $P =_{\text{eval}} P'$ ” is checked. We now show that $P =_{\text{eval}} P'$ is if and only if $P =_L P'$. Suppose $P =_{\text{eval}} P'$. Since the game does not abort, it must hold that $P =_L P'$. Therefore, we know $P =_{\text{eval}} P'$ implies $P =_L P'$.

On the other hand, we show the following lemma.

Lemma 5.2.6. *In Game₂^A, at any step of the execution, we have*

$$\forall P \in \text{Span}(L) : P =_{\text{eval}} 0, \quad (5.8)$$

which implies for any two polynomials P, P' of variables X and $\{A_i, Y_i\}_{i \in [\text{sid}]}$

$$P =_L P' \text{ implies } P =_{\text{eval}} P'. \quad (5.9)$$

Proof. We just need show that for each $R \in L$, we have $R =_{\text{eval}} 0$. From the description of S_2 , we know

$$L = \{A_i + c_i Y_i - s_i, Y_i - y_i X\}_{i \in \mathcal{I}_{\text{fin}}}.$$

For $R = A_i + c_i Y_i - s_i$, we have $R =_{\text{eval}} a_i + c_i \cdot y_i \cdot x_i - s_i = 0$, since $s_i = a_i + c_i \cdot y_i \cdot x_i$. For $R = Y_i - y_i X$, we have $R =_{\text{eval}} y_i \cdot x - y_i \cdot x = 0$. Therefore, we know the lemma holds. $\square \square$

From the above lemma, we know $P =_L P'$ is equivalent to $P =_{\text{eval}} P'$ at the step in Φ when the condition “ $\exists P' \in \text{Cur}$ such that $P =_{\text{eval}} P'$ ”. Therefore, we know the view of the adversary is identical in these two games, which implies

$$\Pr[\text{Game}_2^A = 1] = \Pr[\text{Game}_2'^A = 1] . \quad (5.10)$$

Game₃^A: Game₃^A is defined in Figure 5.5 by ignoring the dashed box, where the only difference from Game₂^{'A} is the original abort condition is removed from Φ and a new abort condition is added to S_2 . Also, in Game₃^A, since the new abort condition only uses the information L , we do not need to assign values to the variables anymore.

We first show that the oracle Φ in Game₃^A is well-defined, that is, for each query P to Φ , there exists at most one $P' \in \text{Cur}$ such that $P =_L P'$. Suppose during a query P to Φ in Game₂^A, the game does not abort and there exists $P', P'' \in \text{Cur}$ such that $P' =_L P =_L P''$. Without loss of generality assume P'' is added to Cur after P' . If L is not updated after P'' is added to Cur , then by the description of Φ , we know $P' \neq_L P''$, which yields a contradiction. Otherwise, L is updated after P'' is added to Cur . Consider the last time L is updated in S_2 . Since $P' =_L P''$ and $P', P'' \in \text{Cur}$, we know Game₃^A must abort in S_2 , which yields a contradiction. Therefore, we know the oracle Φ in Game₃^A is well-defined.

To show that the probability \mathcal{A} wins Game₂^{'A} is bounded by the probability \mathcal{A} wins Game₃^A, we introduce an intermediate game Game₃^{'A} which is defined in Figure 5.5 containing everything. We first show that the probability \mathcal{A} wins Game₂^{'A} is bounded by the probability \mathcal{A} wins Game₃^{'A}. Denote the event E_2 in Game₂^{'A} as during a query to S_2 after L is updated,

$$\exists P_1, P_2 \in \text{Cur} \text{ such that } P_1 \neq P_2 \text{ and } P_1 =_L P_2 .$$

Then, we have Game₃^{'A} is identical to Game₂^{'A} except it aborts when E_2 occurs, which implies

$$\Pr[\text{Game}_2'^A = 1] \leq \Pr[\text{Game}_3'^A = 1] + \Pr[E_2] , \quad (5.11)$$

We now show that $\Pr[E_2] = 0$. Suppose E_2 occurs. Then, we know at some timestep in Game₂^{'A} there exists $P_1, P_2 \in \text{Cur}$ such that $P_1 \neq P_2$ and $P_1 =_L P_2$. We first show that $P_1 \neq_{\text{eval}} P_2$. Suppose $P_1 =_{\text{eval}} P_2$. Without loss of generality, assume P_1 is added to Cur before P_2 . Consider the step when P_2 is added to Cur . Since P_1 is already in Cur , we know $P_1 \neq_L P_2$. However, since $P_2 \neq_L P_1$ but $P_2 =_{\text{eval}} P_1$, the game aborts, which yields a contradiction. Thus, we know

<p>Game $\boxed{\text{Game}_2'^A}, \boxed{\text{Game}_3^A}, \text{Game}_3'^A :$</p> <p>$(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa) ; \quad x \leftarrow_{\\$} \mathbb{Z}_p^* ;$</p> <p>$\boxed{\text{assign } x \text{ to variable } X}$</p> <p>$\text{sid} \leftarrow 0 ; \ell \leftarrow 0 ; \mathcal{S} \leftarrow \emptyset ; \text{Cur} \leftarrow \emptyset ; \Xi \leftarrow () ;$</p> <p>$T \leftarrow ()$</p> <p>$\{(m_k, \sigma_k)\}_{k \in [\ell+1]} \leftarrow_{\\$} \mathcal{A}^{\Pi, S_1, S_2, H}(p, \Phi(1), \Phi(X))$</p> <p>If $\exists k_1 \neq k_2$ such that $(m_{k_1}, \sigma_{k_1}) = (m_{k_2}, \sigma_{k_2})$ then</p> <p style="padding-left: 20px;">Return 0</p> <p>If $\exists k \in [\ell + 1]$ such that $y_k^* = 0$</p> <p style="padding-left: 20px;">or $c_k \neq H(\Phi(s_k - c_k \cdot y_k \cdot X) \parallel \Phi(y_k \cdot X) \parallel m_i)$</p> <p style="padding-left: 20px;">where $(c_k, s_k, y_k) = \sigma_k$ then return 0</p> <p>Return 1</p> <p><u>Oracle $\Phi(P) :$</u></p> <p>$\boxed{\text{If } \exists P' \in \text{Cur} \text{ such that } P =_{\text{eval}} P' \text{ and } P \neq_L P' \text{ then } \mathbf{abort \ game}}$</p> <p>If $\exists P' \in \text{Cur}$ such that $P =_L P'$ then</p> <p style="padding-left: 20px;">Return $\Xi(P')$</p> <p>$\Xi(P) \leftarrow_{\\$} \{0, 1\}^{\log(p)} \setminus \Xi(\text{Cur})$</p> <p>$\text{Cur} \leftarrow \text{Cur} \cup \{P\}$</p> <p>Return $\Xi(P)$</p> <p><u>Oracle $\Pi(\xi, \xi', b) :$</u></p> <p>If $\exists P, P' \in \text{Cur}$ such that $\xi = \Xi(P)$</p> <p style="padding-left: 20px;">and $\xi' = \Xi(P')$ then</p> <p style="padding-left: 40px;">Return $\Phi(P + (-1)^b P')$</p> <p>Else return \perp</p>	<p><u>Oracle $S_1 :$</u></p> <p>$\text{sid} \leftarrow \text{sid} + 1$</p> <p>$a_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p ; y_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p^*$</p> <p>$\text{st}_{\text{sid}}^s \leftarrow (a_{\text{sid}}, y_{\text{sid}})$</p> <p>$\boxed{\text{Assign } a_{\text{sid}} \text{ to variable } A_{\text{sid}}}$</p> <p>$\boxed{\text{Assign } y_{\text{sid}} \cdot x \text{ to variable } Y_{\text{sid}}}$</p> <p>$\text{msg}_1 \leftarrow (\Phi(A_{\text{sid}}), \Phi(Y_{\text{sid}}))$</p> <p>Return $(\text{sid}, \text{msg}_1)$</p> <p><u>Oracle $S_2(i, c_i) :$</u></p> <p>If $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ then return \perp</p> <p>$(a_i, y_i) \leftarrow \text{st}_i^s$</p> <p>$s_i \leftarrow a_i + c_i \cdot y_i \cdot x$</p> <p>$R_1 \leftarrow A_i + c_i Y_i - s_i$</p> <p>$R_2 \leftarrow Y_i - y_i X$</p> <p>$L \leftarrow L \cup \{R_1, R_2\}$</p> <p>$\text{msg}_2 \leftarrow (s_i, y_i)$</p> <p>$\boxed{\text{If } \exists P_1, P_2 \in \text{Cur} \text{ such that } P_1 \neq P_2 \text{ and } P_1 =_L P_2 \text{ then } \mathbf{abort \ game}}$</p> <p>$\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$</p> <p>$\ell \leftarrow \ell + 1$</p> <p>Return msg_2</p> <p><u>Oracle $H(\text{str}) :$</u></p> <p>If $T(\text{str}) = \perp$ then</p> <p style="padding-left: 20px;">$T(\text{str}) \leftarrow_{\\$} \mathbb{Z}_p$</p> <p>Return $T(\text{str})$</p>
---	--

Figure 5.5: The definition for Game_3^A and its difference from Game_2^A . Game_2^A contains all but the solid boxes and Game_3^A contains all but the dashed boxes. We also define an intermediate game $\text{Game}_3'^A$ which contains both dashed and solid boxes.

$P_1 \neq_{\text{eval}} P_2$. Then, by Lemma 5.2.6, we know $P_1 \neq_L P_2$ at any timestep in $\text{Game}_2'^A$, which yields

a contradiction. Therefore, we know E_2 never occurs in Game'_2^A , which implies

$$\Pr[\text{Game}'_2^A = 1] \leq \Pr[\text{Game}'_3^A = 1].$$

Also, since the only difference between Game'_3^A and Game_3^A is that Game'_3^A might abort in Φ while Game_3^A never abort in Φ , we have $\Pr[\text{Game}'_3^A = 1] \leq \Pr[\text{Game}_3^A = 1]$. Therefore, we have

$$\Pr[\text{Game}'_2^A = 1] \leq \Pr[\text{Game}'_3^A = 1] \leq \Pr[\text{Game}_3^A = 1]. \quad (5.12)$$

Game_4^A : This is defined in Figure 5.6 by ignoring the dashed box. Game_4^A is identical to Game_3^A , except the generation of x , $\{a_i, y_i, s_i\}_{i \in [\text{sid}]}$ are changed. More precisely, the sampling of x is removed from the main procedure, the sampling of $a_{\text{sid}}, y_{\text{sid}}$ is removed from S_1 , and in S_2 , y_i is sampled from \mathbb{Z}_p^* and s_i is sampled from \mathbb{Z}_p instead of computing from a_i and y_i . The oracle Φ in Game_4^A is well-defined, which can be shown using the same way as in Game_3^A .

We now show that the views of the adversary in Game_3 and Game_4 are identical. Since the value x and a_i are not used in Game_3^A except the dashed box, we just need to show that the distribution of (s_i, y_i) are identical in Game_3^A and Game_4^A for each query (i, c_i) to S_2 . Consider the step when the adversary makes a query (i, c_i) to S_2 in Game_3^A and assume $i \in [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$. The value y_i and a_i are not used anywhere in the game yet. Therefore, given the current transcript, the distribution of (s_i, y_i) is uniformly random in $\mathbb{Z}_p \times \mathbb{Z}_p^*$. Since $a_i \leftarrow s_i + c_i \cdot y_i \cdot x$ and s_i is uniformly in \mathbb{Z}_p given y_i , we know the distribution of a_i is uniformly random in \mathbb{Z}_p even given y_i . Therefore, the distribution of (a_i, y_i) is uniformly random in $\mathbb{Z}_p \times \mathbb{Z}_p^*$. Thus, we know the view of the adversary in Game_3^A and Game_4^A are identical, which implies

$$\Pr[\text{Game}_3^A = 1] = \Pr[\text{Game}_4^A = 1]. \quad (5.13)$$

5.2.2 Proof of Claim 5.2.5

Proof. Without loss of generality, assume the randomness used in Φ and the randomness of \mathcal{A} are fixed, and assume $\Pr[\pi_j = \Delta] > 0$ given the fixed randomness.

The claim is equivalent to show that

$$\forall \boldsymbol{\eta}_0 \in \mathcal{C} : \Pr_{x, \mathbf{a}, \mathbf{y}}[\boldsymbol{\eta} = \boldsymbol{\eta}_0 \mid \pi_j = \Delta] = \frac{1}{|\mathcal{C}|}.$$

The probability here is taken over the randomness $x, \mathbf{a}, \mathbf{y}$, where $\mathbf{a} = (a_1, \dots, a_{\text{sid}^\Delta})$, $\mathbf{y} = (y_1, \dots, y_{\text{sid}^\Delta})$. Also, $x, y_1, \dots, y_{\text{sid}^\Delta}$ are picked uniformly at random from \mathbb{Z}_p^* and $a_1, \dots, a_{\text{sid}^\Delta}$ are picked uniformly at random from \mathbb{Z}_p .

<p><u>Game</u> Game_3^A, Game_4^A :</p> <p>$(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$</p> <p>$x \leftarrow_{\\$} \mathbb{Z}_p^*$</p> <p>$\text{sid} \leftarrow 0$; $\ell \leftarrow 0$; $\mathcal{S} \leftarrow \emptyset$; $\text{Cur} \leftarrow \emptyset$; $\Xi \leftarrow ()$;</p> <p>$T \leftarrow ()$</p> <p>$\{(m_k, \sigma_k)\}_{k \in [\ell+1]} \leftarrow_{\\$} \mathcal{A}^{\Pi, S_1, S_2, H}(p, \Phi(1), \Phi(X))$</p> <p>If $\exists k_1 \neq k_2$ such that $(m_{k_1}, \sigma_{k_1}) = (m_{k_2}, \sigma_{k_2})$ then</p> <p style="padding-left: 20px;">Return 0</p> <p>If $\exists k \in [\ell + 1]$ such that $y_k^* = 0$</p> <p style="padding-left: 20px;">or $c_k \neq H(\Phi(s_k - c_k \cdot y_k \cdot X) \parallel \Phi(y_k \cdot X) \parallel m_i)$</p> <p>where $(c_k, s_k, y_k) = \sigma_k$ then return 0</p> <p>Return 1</p> <p><u>Oracle $\Phi(P)$</u> :</p> <p>If $\exists P' \in \text{Cur}$ such that $P =_L P'$ then</p> <p style="padding-left: 20px;">Return $\Xi(P')$</p> <p>$\Xi(P) \leftarrow_{\\$} \{0, 1\}^{\log(p)} \setminus \Xi(\text{Cur})$</p> <p>$\text{Cur} \leftarrow \text{Cur} \cup \{P\}$</p> <p>Return $\Xi(P)$</p> <p><u>Oracle $\Pi(\xi, \xi', b)$</u> :</p> <p>If $\exists P, P' \in \text{Cur}$ such that $\xi = \Xi(P)$</p> <p style="padding-left: 20px;">and $\xi' = \Xi(P')$ then</p> <p style="padding-left: 20px;">Return $\Phi(P + (-1)^b P')$</p> <p>Else return \perp</p>	<p><u>Oracle S_1</u> :</p> <p>$\text{sid} \leftarrow \text{sid} + 1$</p> <p>$a_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p$; $y_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p^*$</p> <p>$\text{st}_{\text{sid}}^s \leftarrow (a_{\text{sid}}, y_{\text{sid}})$</p> <p>$\text{msg}_1 \leftarrow (\Phi(A_{\text{sid}}), \Phi(Y_{\text{sid}}))$</p> <p>Return $(\text{sid}, \text{msg}_1)$</p> <p><u>Oracle $S_2(i, c_i)$</u> :</p> <p>If $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ then return \perp</p> <p>$(a_i, y_i) \leftarrow \text{st}_i^s$</p> <p>$s_i \leftarrow a_i + c_i \cdot y_i \cdot x$</p> <p>$s_i \leftarrow_{\\$} \mathbb{Z}_p$; $y_i \leftarrow_{\\$} \mathbb{Z}_p^*$</p> <p>$R_1 \leftarrow A_i + c_i Y_i - s_i$</p> <p>$R_2 \leftarrow Y_i - y_i X$</p> <p>$L \leftarrow L \cup \{R_1, R_2\}$</p> <p>$\text{msg}_2 \leftarrow (s_i, y_i)$</p> <p>If $\exists P_1, P_2 \in \text{Cur}$ such that</p> <p style="padding-left: 20px;">$P_1 \neq P_2$ and $P_1 =_L P_2$</p> <p>then abort game</p> <p>$\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$; $\ell \leftarrow \ell + 1$</p> <p>Return msg_2</p> <p><u>Oracle $H(\text{str})$</u> :</p> <p>If $T(\text{str}) = \perp$ then</p> <p style="padding-left: 20px;">$T(\text{str}) \leftarrow_{\\$} \mathbb{Z}_p$</p> <p>Return $T(\text{str})$</p>
--	--

Figure 5.6: The definition for Game_4^A and its difference from Game_3^A . Game_3^A contains all but the solid box and Game_4^A contains all but the dashed box.

We first show that

$$\pi_j = \Delta \quad \text{implies} \quad \eta \in \mathcal{C} .$$

Suppose $\pi_j = \Delta$ occurs. We just need to show $D_i(\eta) \neq 0$ for each $i \in [m]$ and $B_i(\eta) = 0$ for each

$i \in [q]$. For $D_1, \dots, D_{\text{sid}^\Delta+1}$, since $x \neq 0$ and $y_i \neq 0$ for each $i \in [\text{sid}^\Delta]$, we know $D_1(\eta) = x \neq 0$ and $D_{i+1}(\eta) = y_i \cdot x \neq 0$ for each $i \in [\text{sid}^\Delta]$. For $D_{\text{sid}^\Delta+1}, \dots, D_m$, we make the argument using the original label $\{D_{P,P'}\}_{P,P' \in \text{Cur}^\Delta, P \neq P'}$. For each $P, P' \in \text{Cur}^\Delta$ such that $P \neq P'$, assume without loss of generality P is added to Cur before P' . When P' is added to Cur , since P is already in Cur , we know $P' \neq_{\text{eval}} P$, which implies $D_{P,P'}(\eta) = P'(\eta) - P(\eta) \neq 0$.

Also, for each $i \in \mathcal{I}_{\text{fin}}^\Delta$, consider the query (i, c_i^Δ) made to S_2 . Since $\pi_j = \Delta$, we have $s_i^\Delta = a_i + c_i^\Delta \cdot y_i \cdot x$ and $y_i^\Delta = y_i$. Therefore, we have $B_{(i,1)}(\eta) = a_i + c_i^\Delta \cdot y_i \cdot x - s_i^\Delta = 0$ and $B_{(i,2)}(\eta) = y_i \cdot x - y_i^\Delta \cdot x = 0$.² Therefore, we have $\eta \in \mathcal{C}$.

We then show that

$$\eta \in \mathcal{C} \quad \text{implies} \quad \pi_j = \Delta .$$

Since $\Pr[\pi_j = \Delta] > 0$, we know there exists $(x_0, \mathbf{a}_0, \mathbf{y}_0) \in \mathbb{Z}_p^{1+2\text{sid}^\Delta}$ such that $\pi_j = \Delta$ when $(x, \mathbf{a}, \mathbf{y}) = (x_0, \mathbf{a}_0, \mathbf{y}_0)$. We now show that for any $(x_1, \mathbf{a}_1, \mathbf{y}_1) \in \mathbb{Z}_p^{1+2\text{sid}^\Delta}$, given $(x, \mathbf{a}, \mathbf{y}) = (x_1, \mathbf{a}_1, \mathbf{y}_1)$ and $\eta \in \mathcal{C}$, it must have $\pi_j = \Delta$.

Denote the case when $(x, \mathbf{a}, \mathbf{y}) = (x_0, \mathbf{a}_0, \mathbf{y}_0)$ as case 0 and the case when $(x, \mathbf{a}, \mathbf{y}) = (x_1, \mathbf{a}_1, \mathbf{y}_1)$ as case 1. We will show that the transcripts between the adversary and the oracles are exactly the same in these two cases, which implies $\pi_j = \Delta$ in case 1. We show this by induction. It is clear that the transcripts are the same at the beginning. For a time step T , suppose the transcripts are the same prior to this step and we have the following situations:

- Query to Φ, S_1, Π : Suppose the adversary receives $(\Phi(1), \Phi(X))$ or makes query to S_1 or Π at step T . For the case that the adversary makes a query to S_1 or Φ , the transcripts can only differ on the invocation of Φ in S_1 or Π . Therefore, we only need to consider the queries and outputs of each Φ .

For the k -th query to Φ where $k < j$, since the prior transcripts are the same in these two cases and the adversary is deterministic, we know the query P_k and the set Cur are the same in the two cases. If $P_k \neq_{\text{eval}} P'$ for any $P' \in \text{Cur}$ in case 0, then we know P_k is added to Cur in case 0. Since $\pi_j = \Delta$ occurs in case 0, we know $\{P_k\} \cup \text{Cur} \subseteq \text{Cur}^\Delta$. Since $\eta_1 \in \mathcal{C}$, we know $P_k(\eta_1) \neq P'(\eta_1)$ for any $P' \in \text{Cur}^\Delta$. Therefore, we have $P_k \neq_{\text{eval}} P'$ for any $P' \in \text{Cur}$ in case 1 too. Then, the outputs of Φ are the same in the two cases.

Otherwise, if $P_k =_{\text{eval}} P'$ for some $P' \in \text{Cur}$, we know such P' must be unique. Since $E_{1,k}$

²Note here the value $y_i \cdot x$ is assigned to Y_i

does not occur in case 0, we have $P_k =_L P'$ in case 0. Since the current L is the same in the two cases, we know $P_k =_L P'$ in case 1 too. Since $P_k =_L P'$ implies $P_k =_{\text{eval}} P'$, we have $P_k =_{\text{eval}} P'$ in case 1 too. Thus, the output of Φ must be the same in the two cases. Therefore, we know the transcripts in these two cases must be the same after the k -th query to Φ is finished.

- Query to S_2 : Suppose the adversary makes query (i, c_i) to S_2 at step T . Since $\pi_j = \Delta$ occurs in case 0, we know $i \in \mathcal{I}_{\text{fin}}^\Delta$, $c_i = c_i^\Delta$, $y_i = y_{0,i} = y_i^\Delta$, and $s_i = a_{0,i} + c_i \cdot y_{0,i} \cdot x_0 = s_i^\Delta$ in case 0. Since the transcripts are the same in the two cases prior to T and the adversary is deterministic, we know c_i is the same in both cases. Therefore, we know $c_i = c_i^\Delta$ in case 1. Since $\boldsymbol{\eta}_1 \in \mathcal{C}$, we have

$$B_{(i,1)}(\boldsymbol{\eta}_1) = a_{1,i} + c_i^\Delta \cdot y_{1,i} \cdot x_1 - s_i^\Delta = 0 ,$$

$$B_{(i,2)}(\boldsymbol{\eta}_1) = y_{1,i} \cdot x_1 - y_i^\Delta \cdot x_1 = 0 .$$

Therefore, we have $y_i = y_{i,1} = y_i^\Delta$ and $s_i = a_{1,i} + c_i \cdot y_{1,i} \cdot x_1 = a_{1,i} + c_i^\Delta \cdot y_{1,i} \cdot x_1 = s_i^\Delta$ in case 1. Since the output (y_i, s_i) is the same in the two cases, we know the transcripts must be the same in these two cases after the query to S_2 is finished.

- Query to H : Since H does not involve the randomness $x, \mathbf{a}, \mathbf{y}$, we know the transcripts are the same in the two cases after the query.

By induction, we know the transcript is the same by the step when the j -th query is made to Φ in the two cases. Therefore, we know $\pi_j = \Delta$ in case 1. Since it holds for any $(x_1, \mathbf{a}_1, \mathbf{y}_1) \in \mathbb{Z}_p^{\text{sid}^\Delta}$, we know $\boldsymbol{\eta} \in \mathcal{C}$ implies $\pi_j = \Delta$. Therefore, $\pi_j = \Delta$ is equivalent to $\boldsymbol{\eta} \in \mathcal{C}$, which implies for any $\boldsymbol{\eta}_0 \in \mathcal{C}$

$$\Pr_{x,\mathbf{a},\mathbf{y}}[\boldsymbol{\eta} = \boldsymbol{\eta}_0 | \pi_j = \Delta] = \Pr_{x,\mathbf{a},\mathbf{y}}[\boldsymbol{\eta} = \boldsymbol{\eta}_0 | \boldsymbol{\eta} \in \mathcal{C}] .$$

It is left to show $\Pr_{x,\mathbf{a},\mathbf{y}}[\boldsymbol{\eta} = \boldsymbol{\eta}_0 | \boldsymbol{\eta} \in \mathcal{C}] = \frac{1}{|\mathcal{C}|}$ for any $\boldsymbol{\eta}_0 \in \mathcal{C}$. Denote $\mathcal{E} := \mathbb{Z}_p^* \times (\mathbb{Z}_p \times \mathbb{Z}_p^*)^{\text{sid}^\Delta}$ and we know $(x, a_1, y_1, \dots, a_{\text{sid}^\Delta}, y_{\text{sid}^\Delta})$ is uniformly distributed over \mathcal{E} . Therefore, $\boldsymbol{\eta} = (x, a_1, y_1 \cdot x, \dots, a_{\text{sid}^\Delta} \cdot x, y_{\text{sid}^\Delta} \cdot x)$ is also uniformly distributed over \mathcal{E} , which implies for any $\boldsymbol{\eta}_0 \in \mathcal{E}$,

$$\Pr_{x,\mathbf{a},\mathbf{y}}[\boldsymbol{\eta} = \boldsymbol{\eta}_0] = \frac{1}{|\mathcal{E}|} .$$

Since $\mathcal{C} \subseteq \mathcal{E}$, we have for any $\eta_0 \in \mathcal{C}$

$$\Pr_{x,\alpha,\mathbf{y}}[\boldsymbol{\eta} = \eta_0 | \boldsymbol{\eta} \in \mathcal{C}] = \frac{1/|\mathcal{E}|}{|\mathcal{C}|/|\mathcal{E}|} = \frac{1}{|\mathcal{C}|}.$$

□

5.2.3 Proof of Lemma 5.2.2

We construct \mathcal{B} that interacts with \mathcal{A} by simulating the oracles from Game_4 using the two oracles S and H in WFROS. In particular, we extract suitable vectors α and β to query to H in WFROS, i.e., each RO query str is decomposed as $\text{str} = \xi^A \parallel \xi^Y \parallel m$, where ξ^A and ξ^Y are encodings of group elements. If both encodings are valid, there must exist P^A, P^Y such that $\Xi(P^A) = \xi^A$ and $\Xi(P^Y) = \xi^Y$; then, \mathcal{B} defines two vectors α and β to make a corresponding query to H in WFROS. The oracle S is also used to simulate the signer's second stage. Finally, when \mathcal{A} outputs $Q_{\text{S}_1} + 1$ different valid message-signature pairs in Game_4 , \mathcal{B} tries to map each valid message-signature pair to a query to H in WFROS. We show that this strategy succeeds with probability close to that of \mathcal{A} succeeding.

THE ADVERSARY \mathcal{B} . Specifically, \mathcal{B} initializes the variables sid , Cur , \mathcal{I}_{fin} , Ξ , and T as in Game_4 . In addition, \mathcal{B} initializes an empty table Hid , used later in the simulation of $\tilde{\text{H}}$.

Then, \mathcal{B} runs \mathcal{A} on input $(p, \tilde{\Phi}(1), \tilde{\Phi}(X))$ and with access to the oracles $\tilde{\Pi}$, $\tilde{\text{S}}_1$, $\tilde{\text{S}}_2$, and $\tilde{\text{H}}$. These oracles, along with $\tilde{\Phi}$, operate as follows:

Oracles $\tilde{\Phi}$, $\tilde{\Pi}$: Same as in Game_4 . In particular, L is updated by calls to $\tilde{\text{S}}_2$.

Oracle $\tilde{\text{S}}_1$: Same as in Game_4 .

Oracle $\tilde{\text{S}}_2$: Same as Game_4 except that instead of sampling y_i randomly, if $i \in [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$, \mathcal{B} makes a query (i, c_i) to S and uses its output as the value y_i .

Oracle $\tilde{\text{H}}$: After receiving a query str , if $T(\text{str}) \neq \perp$, the value $T(\text{str})$ is returned. Otherwise, str is decomposed as $\text{str} = \xi^A \parallel \xi^Y \parallel m$ such that the length of ξ^A and ξ^Y is $\lceil \log(p) \rceil$.

- If there exist $P^A, P^Y \in \text{Cur}$ such that $\Xi(P^A) = \xi^A$ and $\Xi(P^Y) = \xi^Y$, denote the coefficients of P^A, P^Y as

$$P^A = \hat{\alpha}^g + \hat{\alpha}^X X + \sum_{i \in [\text{sid}]} \hat{\alpha}^{A_i} A_i + \sum_{i \in [\text{sid}]} \hat{\alpha}^{Y_i} Y_i, \quad (5.14)$$

$$P^Y = \hat{\beta}^g + \hat{\beta}^X \mathbf{X} + \sum_{i \in [\text{sid}]} \hat{\beta}^{\mathbf{A}_i} \mathbf{A}_i + \sum_{i \in [\text{sid}]} \hat{\beta}^{\mathbf{Y}_i} \mathbf{Y}_i . \quad (5.15)$$

Then, \mathcal{B} issues the query $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ to \mathbf{H} , where $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{Z}_p^{2Q_{S_1}+1}$ are such that

$$\alpha^{(i')} = \begin{cases} \hat{\alpha}^X, & i' = 0 \\ \hat{\alpha}^{\mathbf{Y}_i}, & i' = 2i - 1, i \in [\text{sid}] \\ -\hat{\alpha}^{\mathbf{A}_i}, & i' = 2i, i \in [\text{sid}] \\ 0, & o.w. \end{cases}, \quad (5.16)$$

$$\beta^{(i')} = \begin{cases} -\hat{\beta}^X, & i' = 0 \\ -\hat{\beta}^{\mathbf{Y}_i}, & i' = 2i - 1, i \in [\text{sid}] \\ \hat{\beta}^{\mathbf{A}_i}, & i' = 2i, i \in [\text{sid}] \\ 0, & o.w. \end{cases}.$$

After receiving the output $(\delta_{\text{hid}}, \text{hid})$, \mathcal{B} sets $T(\text{str}) \leftarrow \delta_{\text{hid}}$ and $\text{Hid}(\text{str}) \leftarrow \text{hid}$.

- Otherwise, if $\xi^A \notin T(\text{Cur})$ or $\xi^Y \notin T(\text{Cur})$ (or if the decomposition of str is not possible), \mathcal{B} samples $T(\text{str})$ uniformly from \mathbb{Z}_p and sets $\text{Hid}(\text{str}) = \perp$.

Finally, \mathcal{B} returns $T(\text{str})$.

After \mathcal{A} outputs $\{(m_k^*, \sigma_k^*)\}_{k \in [Q_{S_1}+1]}$, \mathcal{B} aborts if the signatures are not valid, i.e., one of the following conditions is not satisfied:

$$\forall k_1, k_2 \in [Q_{S_1} + 1] \text{ and } k_1 \neq k_2 : (m_{k_1}^*, \sigma_{k_1}^*) \neq (m_{k_2}^*, \sigma_{k_2}^*), \quad (5.17)$$

$$\forall k \in [Q_{S_1} + 1] : y_k^* \neq 0 \wedge c_k^* = \tilde{\mathbf{H}}(\text{str}_k^*), \quad (5.18)$$

where $(c_k^*, s_k^*, y_k^*) = \sigma_k^*$ and $\text{str}_k^* = \tilde{\Phi}(s_k^* - c_k^* \cdot y_k^* \cdot \mathbf{X}) \parallel \tilde{\Phi}(y_k^* \cdot \mathbf{X}) \parallel m_k^*$. (Here, $\tilde{\mathbf{H}}$ and $\tilde{\Phi}$ are the oracles described previously.) Further, \mathcal{B} aborts if the following condition does not hold:

$$\forall k \in [Q_{S_1} + 1] : \text{Hid}(\text{str}_k^*) \neq \perp . \quad (5.19)$$

Otherwise, \mathcal{B} outputs $\mathcal{J} := \{\text{Hid}(\text{str}_k^*)\}_{k \in [Q_{S_1}+1]}$.

ANALYSIS OF \mathcal{B} . Note that \mathcal{B} queries to H at most once when it receives a query to \tilde{H} and makes $Q_{S_1} + 1$ more queries to \tilde{H} when checking the validity of the output. Therefore, \mathcal{B} makes at most $Q_H + Q_{S_1} + 1$ queries to H . Also, it is clear that \mathcal{B} simulates oracles S_1, S_2 in Game_4 perfectly. For the simulation of \tilde{H} , the only difference is that the distribution of δ_{hid} outputting from H in WFROS is uniformly over \mathbb{Z}_p^* , where in Game_4 it is always uniformly from \mathbb{Z}_p . However, the statistical distance between the two distributions is $1/p$. Since \mathcal{B} makes at most $Q_H + Q_{S_1} + 1$ queries to H , the statistical difference between the view of \mathcal{A} in Game_4 and that in the one simulated by \mathcal{B} is bounded by $(Q_H + Q_{S_1} + 1)/p$.

Denote the event E_1 such that when \mathcal{B} checks the output from \mathcal{A} , both (5.17) and (5.18) hold. As these are exactly the winning conditions of Game_4 , which is simulated statistically close to perfect, we have $\Pr[E_1] + \frac{Q_H + Q_{S_1} + 1}{p} \geq \Pr[\text{Game}_4^{\mathcal{A}} = 1]$. Also, let E_2 be the event for which the condition (5.19) holds immediately afterward. If E_2 does not happen, but E_1 does, then we know \mathcal{A} outputs a valid message-signature pair (m_k^*, σ_k^*) such that $\text{Hid}(\text{str}_k^*) = \perp$, which is unlikely to happen. The following formalizes this, and the proof is in Section 5.2.4.

Claim 5.2.7. $\Pr[E_1 \wedge (\neg E_2)] \leq \frac{2Q_\Phi(Q_H + Q_{S_1} + 1)}{p - Q_\Phi}$.

Then, we can conclude the proof with the following claim.

Claim 5.2.8. *If both E_1 and E_2 happen, then \mathcal{B} outputs a valid WFROS solution \mathcal{J} , which in turn implies that $\Pr[E_1 \wedge E_2] \leq \text{Adv}_{Q_{S_1}, p}^{\text{wfros}}(\mathcal{B})$.*

Before we proceed with a proof, we state a simple lemma for Game_4 that is used in the proof of Claim 5.2.8. The proof is immediate and follows from the uniqueness of values returned by the oracle.

Lemma 5.2.9. *At any step of Game_4 , for any two polynomials P and P' , suppose we make queries P and P' to Φ . If $\Phi(P) = \Phi(P')$, then $P =_L P'$. Equivalently, if $P \neq_L P'$, then we have $\Phi(P) \neq \Phi(P')$.*

Proof of Claim 5.2.8. Suppose both E_1 and E_2 happen. We first show that for any $k_1, k_2 \in [Q_{S_1} + 1]$ and $k_1 \neq k_2$, it holds that $\text{str}_{k_1}^* \neq \text{str}_{k_2}^*$, which implies $|\mathcal{J}| = Q_{S_1} + 1$. We then show that \mathcal{J} is valid for the WFROS game.

For $k_1, k_2 \in [Q_{S_1} + 1]$ and $k_1 \neq k_2$, suppose $\text{str}_{k_1}^* = \text{str}_{k_2}^*$. Then, we have

$$c_{k_1}^* = \tilde{H}(\text{str}_{k_1}^*) = \tilde{H}(\text{str}_{k_2}^*) = c_{k_2}^*, \quad m_{k_1}^* = m_{k_2}^*,$$

$$\tilde{\Phi}(s_{k_1}^* - c_{k_1}^* \cdot y_{k_1}^* \cdot \mathbf{X}) = \tilde{\Phi}(s_{k_2}^* - c_{k_2}^* \cdot y_{k_2}^* \cdot \mathbf{X}), \quad \tilde{\Phi}(y_{k_1}^* \cdot \mathbf{X}) = \tilde{\Phi}(y_{k_2}^* \cdot \mathbf{X}).$$

By Lemma 5.2.9, it holds that $(m_{k_1}^*, (c_{k_1}^*, s_{k_1}^*, y_{k_1}^*)) = (m_{k_2}^*, (c_{k_2}^*, s_{k_2}^*, y_{k_2}^*))$. However, since E_1 happens, by (5.17), we have $(m_{k_1}^*, \sigma_{k_1}^*) \neq (m_{k_2}^*, \sigma_{k_2}^*)$, which yields a contradiction. Therefore, we know $\text{str}_{k_1}^* \neq \text{str}_{k_2}^*$. From the simulation of $\tilde{\mathbf{H}}$, we have $\text{Hid}(\text{str}_{k_1}^*) \neq \text{Hid}(\text{str}_{k_2}^*)$, and thus we have $|\mathcal{J}| = \ell + 1$.

We now show that for each $j \in \mathcal{J}$, it holds that

$$\alpha_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}} y_i \left(\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)} \right) = \delta_j \left(\beta_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) \right), \quad (\text{C1})$$

$$\beta_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right) \neq 0, \quad (\text{C2})$$

which implies \mathcal{J} is valid for the WFROS game.

Let us fix a $j \in \mathcal{J}$. Since $j \in \mathcal{J}$, there exists $k \in [Q_{S_1} + 1]$ such that $\text{Hid}(\text{str}_k^*) = j$, and there exists $P_j^A, P_j^Y \in \text{Cur}$ and m_j such that $\text{str}_k^* = \Xi(P_j^A) \parallel \Xi(P_j^Y) \parallel m_j$. Let $\hat{\alpha}_j$ and $\hat{\beta}_j$ denote the coefficients of P_j^A and P_j^Y . Since E_1 happens, by (5.18) and Lemma 5.2.9, we have $P_j^A =_L s_k^* - \delta_j \cdot y_k^* \cdot \mathbf{X}$, $P_j^Y =_L y_k^* \cdot \mathbf{X}$, which implies there exists $\{r_{1,i}^{P_j^A}, r_{2,i}^{P_j^A}, r_{1,i}^{P_j^Y}, r_{2,i}^{P_j^Y}\}_{i \in \mathcal{I}_{\text{fin}}}$ such that

$$\begin{aligned} P_j^A &= s_k^* - \delta_j \cdot y_k^* \cdot \mathbf{X} + \sum_{i \in [Q_{S_1}]} r_{1,i}^{P_j^A} (A_i + c_i Y_i - s_i) + \sum_{i \in [Q_{S_1}]} r_{2,i}^{P_j^A} (Y_i - y_i \mathbf{X}), \\ P_j^Y &= y_k^* \cdot \mathbf{X} + \sum_{i \in [Q_{S_1}]} r_{1,i}^{P_j^Y} (A_i + c_i Y_i - s_i) + \sum_{i \in [Q_{S_1}]} r_{2,i}^{P_j^Y} (Y_i - y_i \mathbf{X}). \end{aligned} \quad (5.20)$$

By looking into the coefficients of \mathbf{X} , $\{A_i, Y_i\}_{i \in [Q_{S_1}]}$ on both sides of (5.20), we have $\hat{\alpha}_j^{A_i} = r_{1,i}^{P_j^A}$, $\hat{\alpha}_j^{Y_i} = r_{1,i}^{P_j^A} \cdot c_i + r_{2,i}^{P_j^A}$, $\hat{\beta}_j^{A_i} = r_{1,i}^{P_j^Y}$, $\hat{\beta}_j^{Y_i} = r_{1,i}^{P_j^Y}$ and $c_i + r_{2,i}^{P_j^Y}$ for each $i \in [Q_{S_1}]$, $\hat{\beta}_j^{\mathbf{X}} = y_k^* - \sum_{i \in [Q_{S_1}]} r_{2,i}^{P_j^Y} \cdot y_i$, and $\hat{\alpha}_j^{\mathbf{X}} = -\delta_j \cdot y_k^* - \sum_{i \in [Q_{S_1}]} r_{2,i}^{P_j^A} \cdot y_i$. By sorting out the equations, we have

$$y_k^* = \beta_j^{\mathbf{X}} + \sum_{i \in \mathcal{I}_{\text{fin}}} y_i \left(\hat{\beta}_j^{Y_i} - c_i \cdot \hat{\beta}_j^{A_i} \right),$$

$$\hat{\alpha}_j^{\mathbf{X}} + \sum_{i \in [Q_{S_1}]} y_i (\hat{\alpha}_j^{Y_i} - c_i \cdot \hat{\alpha}_j^{A_i}) = -\delta_j \cdot \left(\hat{\beta}_j^{\mathbf{X}} + \sum_{i \in [Q_{S_1}]} y_i \left(\hat{\beta}_j^{Y_i} - c_i \cdot \hat{\beta}_j^{A_i} \right) \right).$$

By the definition of α and β in (5.16), we know (C1) holds and

$$y_k^* = \beta_j^{(0)} + \sum_{i \in \mathcal{I}_{\text{fin}}} y_i \left(\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)} \right).$$

Since E_1 happens, by (5.18), we know $y_k^* \neq 0$, which implies (C2) happens. \square

5.2.4 Proof of Claim 5.2.7

Suppose $E_1 \wedge (\neg E_2)$ occurs. Denote str_j as the input of the j -th query to $\tilde{\text{H}}$. Denote the total number of queries to $\tilde{\text{H}}$ as $\text{num}_{\tilde{\text{H}}}^{\text{tot}}$. Denote the decomposition of str_j as $\text{str}_j = \xi_j^A \parallel \xi_j^Y \parallel m_j$. Denote Cur_j as the set Cur by the step when the j -th query to $\tilde{\text{H}}$ is made and denote Cur^{tot} as the set Cur after \mathcal{B} finishes the check of the condition (5.17) and (5.18). Since \mathcal{B} makes a query str_k^* to $\tilde{\text{H}}$ to check the condition (5.18), there exists $j \in [\text{num}_{\tilde{\text{H}}}^{\text{tot}}]$ such that $\text{str}_j = \text{str}_k^*$. Let j_{\min} be the smallest index such that $\text{str}_{j_{\min}} = \text{str}_k^*$. Since $\text{Hid}(\text{str}_k^*) = \perp$, from the simulation of $\tilde{\text{H}}$, we know $\xi_{j_{\min}}^A \notin \Xi(\text{Cur}_{j_{\min}})$ or $\xi_{j_{\min}}^Y \notin \Xi(\text{Cur}_{j_{\min}})$. However, since $\xi_{j_{\min}}^A = \tilde{\Phi}(s_k^* - c_k^* \cdot y_k^* \cdot \mathcal{X})$ and $\xi_{j_{\min}}^Y = \tilde{\Phi}(y_k^* \cdot \mathcal{X})$, we know $\xi_{j_{\min}}^A, \xi_{j_{\min}}^Y \in \text{Cur}^{\text{tot}}$. Therefore, denote the set of all ξ_j^Y and ξ_j^A that do not correspond to any encoding of polynomials when the j -th query to $\tilde{\text{H}}$ is made as

$$D^{\text{tot}} := \{ \xi_j^A \mid j \in [\text{num}_{\tilde{\text{H}}}^{\text{tot}}], \xi_j^A \notin \Xi(\text{Cur}_j) \} \cup \{ \xi_j^Y \mid j \in [\text{num}_{\tilde{\text{H}}}^{\text{tot}}], \xi_j^Y \notin \Xi(\text{Cur}_j) \},$$

and then we have at least one of $\xi_{j_{\min}}^A$ and $\xi_{j_{\min}}^Y$ is in $D \cap \Xi(\text{Cur}^{\text{tot}})$, which implies $D \cap \Xi(\text{Cur}^{\text{tot}}) \neq \emptyset$. Therefore, we have the event E occurs implies $D \cap \Xi(\text{Cur}^{\text{tot}}) \neq \emptyset$, which means

$$\Pr[E_1 \wedge (\neg E_2)] \leq \Pr[D^{\text{tot}} \cap \Xi(\text{Cur}^{\text{tot}}) \neq \emptyset]. \quad (5.21)$$

It is left to bound $\Pr[D \cap \Xi(\text{Cur}^{\text{tot}}) \neq \emptyset]$.

Denote

$$D_j := \{ \xi_{j'}^A \mid j' \in [j], \xi_{j'}^A \notin \Xi(\text{Cur}_{j'}) \} \cup \{ \xi_{j'}^Y \mid j' \in [j], \xi_{j'}^Y \notin \Xi(\text{Cur}_{j'}) \}.$$

Denote $\text{Cur}^{(i)}$ as the set Cur after the i -th query to $\tilde{\Phi}$ is finished and $\text{Cur}^{(0)} = \emptyset$. Consider the step when the i -th query to $\tilde{\Phi}$ is made. Denote the number of queries to $\tilde{\text{H}}$ before the i -th query to $\tilde{\Phi}$ is made as $\text{num}_{\tilde{\text{H}}}^{(i)}$. Denote the event E'_i as $D_{\text{num}_{\tilde{\text{H}}}^{(i)}} \cap \Xi(\text{Cur}^{(i-1)}) = \emptyset$ and $D_{\text{num}_{\tilde{\text{H}}}^{(i)}} \cap \Xi(\text{Cur}^{(i)}) \neq \emptyset$. We first show that if $D^{\text{tot}} \cap \Xi(\text{Cur}^{\text{tot}}) \neq \emptyset$, then there exists i such that E'_i occurs, and then bound $\Pr[E'_i]$ for each i .

Denote the total number of queries to $\tilde{\Phi}$ as $\text{num}_{\tilde{\Phi}}^{\text{tot}}$. Suppose none of $E'_1, \dots, E'_{\text{num}_{\tilde{\Phi}}^{\text{tot}}}$ occurs. We show that at any time step, supposing the number of queries to $\tilde{\Phi}$ made so far is i and the number of queries to \tilde{H} made so far is j , we have $D_j \cap T(\text{Cur}^{(i)}) = \emptyset$, which implies $D^{\text{tot}} \cap \Xi(\text{Cur}^{\text{tot}}) = \emptyset$. We show the statement by induction. At the beginning, we know $i = 0, j = 0, \text{Cur}^{(0)} = \emptyset$, and $D_0 = \emptyset$. Thus, the statement holds trivially. For any time step with $i > 0$ or $j > 0$, suppose the latest query is made to \tilde{H} and we have $D_{j-1} \cap T(\text{Cur}^{(i)}) = \emptyset$. Consider the step when the j -th query to \tilde{H} is made. If $T(\text{str}_j) \neq \perp$, we have $D_j = D_{j-1}$ and $D_j \cap T(\text{Cur}^{(i)}) = \emptyset$. Otherwise, if $T(\text{str}_j) = \perp$, we have $D_j = D_{j-1} \cup (\{\xi_j^A, \xi_j^Y\} \setminus T(\text{Cur}_j))$. Since $\text{Cur}_j = \text{Cur}^{(i)}$, we have $D_j \cap T(\text{Cur}^{(i)}) = D_{j-1} \cap T(\text{Cur}^{(i)}) = \emptyset$. Therefore, we have $D_j \cap T(\text{Cur}^{(i)}) = \emptyset$. Otherwise, suppose the latest query is made to $\tilde{\Phi}$ and we have $D_j \cap T(\text{Cur}^{(i-1)}) = \emptyset$. Since we have $j = \text{num}_{\tilde{H}}^{(i)}$ and E'_i does not occur, we have $D_j \cap T(\text{Cur}^{(i-1)}) = D_{\text{num}_{\tilde{H}}^{(i)}} \cap \Xi(\text{Cur}^{(i)}) = \emptyset$. Therefore, by induction, the statement holds. Then, considering the step when \mathcal{B} finishes the check of the condition (5.17) and (5.18), we have $D^{\text{tot}} \cap \Xi(\text{Cur}^{\text{tot}}) = D_{\text{num}_{\tilde{H}}^{\text{tot}}} \cap \Xi(\text{Cur}^{\text{tot}}) = \emptyset$. Therefore, if $D^{\text{tot}} \cap \Xi(\text{Cur}^{\text{tot}}) \neq \emptyset$, then at least one of $\{E'_i\}_{i \in [\text{num}_{\tilde{\Phi}}^{\text{tot}}]}$ occurs.

Finally, to bound $\Pr[E'_i]$, consider the i -th query to $\tilde{\Phi}$. Denote the input of the i -th query to $\tilde{\Phi}$ as P_i . Denote $j = \text{num}_{\tilde{H}}^{(i)}$ for simplicity. Suppose E'_i occurs. We know $\text{Cur}^{(i)} \neq \text{Cur}^{(i-1)}$, which implies $\text{Cur}^{(i)} = \text{Cur}^{(i-1)} \cup \{P_i\}$ and $\Xi(\text{Cur}^{(i)}) = \Xi(\text{Cur}^{(i-1)}) \cup \{\Xi(P_i)\}$. Since $D_j \cap \Xi(\text{Cur}^{(i-1)}) = \emptyset$ and $D_j \cap \Xi(\text{Cur}^{(i)}) \neq \emptyset$, we know $\Xi(P_i) \in D_j$. Therefore, we have

$$\Pr[E'_i] \leq \Pr[\text{Cur}^{(i)} = \text{Cur}^{(i-1)} \cup \{P_i\} \wedge \Xi(P_i) \in D_j].$$

Consider the step when $\Xi(P_i)$ is generated. We know D_j is already determined. Therefore, we know $\Xi(P_i)$ is sampled from $\{0, 1\}^{\log(p)} \setminus \Xi(\text{Cur}^{(i-1)})$ uniformly at random independent of D_j , which implies

$$\begin{aligned} \Pr[E'_i] &\leq \Pr[\text{Cur}^{(i)} = \text{Cur}^{(i-1)} \cup \{P_i\} \wedge \Xi(P_i) \in D_j] \\ &\leq \Pr[\Xi(P_i) \in D_j | \text{Cur}^{(i)} = \text{Cur}^{(i-1)} \cup \{P_i\}] \\ &\leq \frac{|D_j|}{p - |\text{Cur}^{(i-1)}|} \leq \frac{|D^{\text{tot}}|}{p - |\text{Cur}^{\text{tot}}|}. \end{aligned}$$

Therefore, we have

$$\Pr[D^{\text{tot}} \cap \Xi(\text{Cur}^{\text{tot}}) \neq \emptyset] \leq \Pr \left[\bigvee_{i \in [\text{num}_{\tilde{\Phi}}^{\text{tot}}]} E'_i \right] \leq \sum_{i \in [\text{num}_{\tilde{\Phi}}^{\text{tot}}]} \Pr[E'_i] \leq \frac{\text{num}_{\tilde{\Phi}}^{\text{tot}} \cdot |D^{\text{tot}}|}{p - |\text{Cur}^{\text{tot}}|}.$$

Since $|D^{\text{tot}}| \leq 2\text{num}_{\tilde{H}}^{\text{tot}} \leq 2(Q_H + Q_{S_1} + 1)$ and $|\text{Cur}^{\text{tot}}| \leq \text{num}_{\tilde{\Phi}}^{\text{tot}} \leq Q_\Phi$, by (5.21), we have

$$\Pr[E_1 \wedge (\neg E_2)] \leq \frac{2\text{num}_{\tilde{\Phi}}^{\text{tot}} \cdot \text{num}_{\tilde{H}}^{\text{tot}}}{p - \text{num}_{\tilde{\Phi}}^{\text{tot}}} = \frac{2Q_\Phi(Q_H + Q_{S_1} + 1)}{p - Q_\Phi}.$$

Chapter 6

EFFICIENT BLIND SIGNATURES IN THE AGM

We now present schemes that are secure in the *algebraic group model* (AGM) [64]. This model considers security for *algebraic adversaries* - these are adversaries that, when used within a reduction, provide a representation of any group element they output in terms of all prior group elements input to the adversary. (We dispense with a more formal definition since the use of the AGM is self-evident in our proofs.)

6.1 Construction and Security

In this section, we introduce a scheme, which we refer to as BS_2 , that relies on the hardness of the (plain) discrete logarithm (DL) problem, which is formalized in Figure 2.1. In contrast to BS_1 , our new scheme (described in Figure 6.1) requires an extra group element Z in the public key, and the commitment X^y in is replaced by $g^t Z^y$. (This will necessary result in an additional scalar in the signature.) However, one could generate Z as an output of a hash function (assuming the hash function is a random oracle, which we assume anyways), although, interestingly, our proof for BS_2 will show that blindness holds even when Z is chosen maliciously by the signer (who may consequently also know its discrete logarithm).

The additional group element Z will in fact allow us to develop a *partially blind* version of BS_2 , which we refer to as PBS, which we discuss in Section 6.3 below. We note that in fact *all* results about BS_2 can be obtained as a corollary of our analysis of PBS, because a blind signature scheme is of course a special case of a partially blind one. However, we are opting for a separate presentation, as the main ideas behind the reduction are much simpler to understand in (plain) BS_2 , and the proof of PBS adds some extra complexity (in particular, in order to obtain a tighter bound), which obfuscates the main ideas.

SECURITY ANALYSIS. The following theorem establishes the blindness of BS_2 . (Its proof is very similar to the blindness proof of BS_1 [GGen].)

<p><u>Algorithm Setup(1^κ) :</u> $(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$ Select $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p^*$ Return $par \leftarrow (p, \mathbb{G}, g, H)$</p> <p><u>Algorithm KeyGen(par) :</u> $(p, \mathbb{G}, g, H) \leftarrow par$ $x \leftarrow_{\\$} \mathbb{Z}_p ; X \leftarrow g^x ; Z \leftarrow_{\\$} \mathbb{G}$ $sk \leftarrow x ; pk \leftarrow (X, Z)$ Return (sk, pk)</p> <p><u>Algorithm $S_1(sk)$:</u> $x \leftarrow sk ; X \leftarrow g^x$ $a, t \leftarrow_{\\$} \mathbb{Z}_p ; y \leftarrow_{\\$} \mathbb{Z}_p^*$ $A \leftarrow g^a ; C \leftarrow g^t Z^y$ $st^s \leftarrow (a, y, t, x) ; msg_1 \leftarrow (A, C)$ Return (st^s, msg_1)</p> <p><u>Algorithm $S_2(st^s, c)$:</u> If $c = 0$ then return \perp $(a, y, t, x) \leftarrow st^s$ $s \leftarrow a + c \cdot y \cdot x$ Return $msg_2 \leftarrow (s, y, t)$</p>	<p><u>Algorithm $U_1(pk, msg_1, m)$:</u> $X \leftarrow pk ; (A, C) \leftarrow msg_1$ $r_1 \leftarrow_{\\$} \mathbb{Z}_p ; \gamma_1, \gamma_2 \leftarrow_{\\$} \mathbb{Z}_p^*$ $A' \leftarrow g^{r_1} \cdot A^{\gamma_1/\gamma_2} \cdot C^{\gamma_1}$ $c' \leftarrow H(A' \ m)$ $c \leftarrow c' \cdot \gamma_2$ $st^u \leftarrow (c, c', r_1, \gamma_1, \gamma_2, X, Z, A, C)$ Return (st^u, c)</p> <p><u>Algorithm $U_2(st^u, msg_2)$:</u> $(c, c', r_1, \gamma_1, \gamma_2, X, Z, A, C) \leftarrow st^u$ $(s, y, t) \leftarrow msg_2$ If $y = 0$ or $C \neq g^t Z^y$ or $g^s \neq A \cdot X^{c \cdot y}$ then return \perp $s' \leftarrow r_1 + (\gamma_1/\gamma_2) \cdot s + \gamma_1 \cdot t$ $y' \leftarrow \gamma_1 \cdot y$ Return $\sigma \leftarrow (c', s', y')$</p> <p><u>Algorithm Ver(pk, σ, m) :</u> $(c, s, y) \leftarrow \sigma$ If $y = 0$ then return 0 $A \leftarrow g^s \cdot X^{-c \cdot y} \cdot Z^y$ If $c \neq H(A \ m)$ then return 0 Return 1</p>
--	--

Figure 6.1: The blind signature scheme $BS_2 = BS_2[\text{GGen}]$.

Theorem 6.1.1. *For a group generation algorithm GGen, the blind signature scheme $BS_2[\text{GGen}]$ is perfectly blind.*

Proof. Let \mathcal{A} be an adversary playing the $\text{Blind}_{BS_2[\text{GGen}]}^{\mathcal{A}}$ game. Similar to the blindness proof of $BS_1[\text{GGen}]$, we can assume the randomness of \mathcal{A} is fixed and \mathcal{A} always finishes both signing sessions and receives valid signatures (σ_0, σ_1) without loss of generality.

Define the view of \mathcal{A} after its execution as $\pi = (X, Z, m_0, m_1, T_0, T_1, \sigma_0, \sigma_1)$, where $T_i := (A_i,$

C_i, c_i, s_i, y_i, t_i), denoting the transcripts learned from interactions with the i -th signing session and $\sigma_i = (c'_i, s'_i, y'_i, t'_i)$. Since the randomness of \mathcal{A} is fixed, the only randomness left is the randomness in U_1 and U_2 . Denote $\eta := (r_1^{(0)}, \gamma_1^{(0)}, \gamma_2^{(0)}, r_1^{(1)}, \gamma_1^{(1)}, \gamma_2^{(1)})$ as the total randomness. To prove the theorem, we need only show that the distribution of π is identical in both the cases $b = 0$ and $b = 1$. We prove this by showing that for any fixed view Δ such that $\Pr[\pi = \Delta | b = 1] > 0$, there exists a unique value of the randomness η that makes $\pi = \Delta$ for the cases $b = 0$ and $b = 1$.

For both the cases $b = 0$ and $b = 1$, we now show that $\pi = \Delta$ if and only if for each $i \in \{0, 1\}$, it holds that

$$\begin{aligned}\gamma_1^{(i)} &= y'_{b_i} \Delta / y_i \Delta, \\ \gamma_2^{(i)} &= c_i \Delta / c'_{b_i} \Delta, \\ r_1^{(i)} &= s'_{b_i} \Delta - s_i \Delta \cdot (\gamma_1^{(i)} / \gamma_2^{(i)}) - \gamma_1^{(i)} \cdot t_i \Delta,\end{aligned}\tag{6.1}$$

where the superscript $(\cdot)^\Delta$ represents the corresponding value in Δ . From the algorithms $BS_2.U_1$ and $BS_2.U_2$, it is clear that the “only if” part holds. For the “if” part, suppose (6.1) holds. Since the randomness of \mathcal{A} is fixed, the view of \mathcal{A} can differ only on the outputs c_0, c_1 from the oracle U_1 or the output (σ_0, σ_1) from the oracle U_2 . Since both signatures in Δ are valid, we have

$$A_i \Delta = g^{s_i \Delta} X^{\Delta - c_i \Delta \cdot y_i \Delta}, \quad C_i \Delta = g^{t_i \Delta} Z^{\Delta y_i \Delta}.\tag{6.2}$$

$$c'_{b_i} \Delta = H(g^{s'_{b_i} \Delta} X^{\Delta - y'_{b_i} \Delta \cdot c'_{b_i} \Delta} Z^{\Delta y'_{b_i} \Delta} \| m_{b_i} \Delta).\tag{6.3}$$

For c_i where $i \in \{0, 1\}$, suppose the values in the view of \mathcal{A} that have already determined when c_i is generated, which must include (X, m_i, A_i, C_i) , are consistent with Δ . By (6.1), we have

$$\begin{aligned}c_i &= \gamma_2^{(i)} \cdot H(g^{r_1^{(i)}} A_i^{\gamma_1^{(i)} / \gamma_2^{(i)}} C_i^{\gamma_1^{(i)}} \| m_{b_i}) \\ &= \gamma_2^{(i)} \cdot H(g^{r_1^{(i)}} A_i^{\Delta \gamma_1^{(i)} / \gamma_2^{(i)}} C_i^{\Delta \gamma_1^{(i)}} \| m_{b_i} \Delta) \\ &= \gamma_2^{(i)} \cdot H(g^{r_1^{(i)} + s_i \Delta \cdot (\gamma_1^{(i)} / \gamma_2^{(i)}) + \gamma_1^{(i)} \cdot t_i \Delta} X^{\Delta - y_i \Delta \cdot c_i \Delta \cdot (\gamma_1^{(i)} / \gamma_2^{(i)})} Z^{\Delta y_i \Delta \cdot \gamma_1^{(i)}} \| m_{b_i} \Delta) \\ &= \gamma_2^{(i)} \cdot H(g^{s'_{b_i} \Delta} X^{\Delta - y'_{b_i} \Delta \cdot c'_{b_i} \Delta} Z^{\Delta y'_{b_i} \Delta} \| m_{b_i} \Delta) \\ &= \gamma_2^{(i)} \cdot c'_{b_i} \Delta = c_i \Delta.\end{aligned}$$

where the third equality is due to (6.2), the fourth equality is due to (6.1), and the final equality is due to (6.3). Then, consider the step when (σ_0, σ_1) are output. Suppose the current view, which

contains T_i , is consistent with Δ . By (6.1), we have

$$\begin{aligned} y'_{b_i} &= \gamma_1^{(i)} \cdot y_i = \gamma_1^{(i)} \cdot y_i^\Delta = y'_{b_i}{}^\Delta, \\ s'_{b_i} &= r_1^{(i)} + s_i(\gamma_1^{(i)}/\gamma_2^{(i)}) + \gamma_1^{(i)} \cdot t_i = r_1^{(i)} + s_i^\Delta(\gamma_1^{(i)}/\gamma_2^{(i)}) + \gamma_1^{(i)} \cdot t_i^\Delta = s'_{b_i}{}^\Delta, \\ c'_{b_i} &= c_i/\gamma_2^{(i)} = c_i^\Delta/\gamma_2^{(i)} = c'_{b_i}{}^\Delta, \end{aligned}$$

which implies $(\sigma_0, \sigma_1) = (\sigma_0^\Delta, \sigma_1^\Delta)$. Therefore, by induction, if (6.1) holds, we know $\pi = \Delta$. \square

The core of the analysis is once again a proof that the scheme is one-more unforgeable in the AGM, i.e., we only prove security against algebraic adversaries. In particular, we model the selected hash function as a random oracle H, to which the adversary is given explicit access.

Theorem 6.1.2. *Let \mathbb{G} be a group generation algorithm. For any algebraic adversary \mathcal{A}_{alg} for the game $\text{OMUF}_{\text{BS}_2[\text{GGen}]}(\kappa)$ making at most Q_{S_1} queries to S_1 and Q_H queries to the random oracle H, there exists an adversary $\mathcal{B}_{\text{dlog}}$ for the DLog problem running in a similar running time as \mathcal{A}_{alg} such that*

$$\text{Adv}_{\text{BS}_2[\text{GGen}]}^{\text{omuf}}(\mathcal{A}_{\text{alg}}, \kappa) \leq 2\text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}, \kappa) + \frac{(Q_H + Q_{S_1} + 1)(Q_H + 3Q_{S_1} + 1)}{p - 1}.$$

6.2 Proof of OMUF (Theorem 6.1.2)

Let us fix an adversary \mathcal{A}_{alg} that makes at most Q_{S_1} queries to S_1 and Q_H queries to the random oracle H. Without loss of generality, assume \mathcal{A}_{alg} makes exactly Q_{S_1} queries to S_1 and exactly one query (i, c_i) to S_2 for each $i \in [Q_{S_1}]$. Then, after \mathcal{A}_{alg} returns, we know $\ell = Q_{S_1}$ and $\mathcal{I}_{\text{fin}} = [Q_{S_1}]$.

The $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game is formally defined in Figure 6.2. In addition to the original OMUF game (defined in Figure 3.1), for each query $(A \parallel m)$ to H, its corresponding hid is recorded in $\text{Hid}(A \parallel m)$, and the output of the query is recorded as δ_{hid} . Also, since \mathcal{A}_{alg} is algebraic, it also provides the representations of A , and the corresponding coefficient $\hat{\alpha}$ are recorded as $\hat{\alpha}_{\text{hid}}$.

Denote the event WIN as \mathcal{A}_{alg} wins the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game, i.e., all output message-signature pairs $\{m_k^*, \sigma_k^*\}_{k \in [Q_{S_1} + 1]}$ are distinct and valid. Furthermore, denote $\text{str}_k^* := g^{s_k^*} X^{-c_k^*} y_k^* Z^{y_k^*} \parallel m_k^*$. We let E be the event in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game for which, after the validity of the output is

<p>Game OMUF$_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$($\lambda$):</p> <p>$(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$; $x \leftarrow_{\\$} \mathbb{Z}_p$; $X \leftarrow g^x$; $Z \leftarrow \mathbb{G}$</p> <p>$\text{sid} \leftarrow 0$; $\ell \leftarrow 0$; $\mathcal{I}_{\text{fin}} \leftarrow \emptyset$; $T \leftarrow ()$; $\text{hid} \leftarrow 0$;</p> <p>$\text{Hid} \leftarrow ()$</p> <p>$\{(m_k^*, \sigma_k^*)\}_{k \in [\ell+1]} \leftarrow_{\\$} \mathcal{A}_{\text{alg}}^{\text{S}_1, \text{S}_2, \text{H}}(p, g, \mathbb{G}, X, Z)$</p> <p>If $\exists k_1 \neq k_2$ such that $(m_{k_1}^*, \sigma_{k_1}^*) = (m_{k_2}^*, \sigma_{k_2}^*)$ then</p> <p style="padding-left: 20px;">Return 0</p> <p>If $\exists k \in [\ell + 1]$ such that $y_k^* = 0$</p> <p style="padding-left: 20px;">or $c_k^* \neq \text{H}(g^{s_k^*} X^{-c_k^*} y_k^* Z^{y_k^*} \ m_k^*)$</p> <p style="padding-left: 20px;">where $(c_k^*, s_k^*, y_k^*) = \sigma_k^*$ then return 0</p> <p>Return 1</p> <p>Oracle H($A \ m$) :</p> <p>If $T(A \ m) = \perp$ then</p> <p style="padding-left: 20px;">$T(A \ m) \leftarrow_{\\$} \mathbb{Z}_p$</p> <p style="padding-left: 20px;">$\text{hid} \leftarrow \text{hid} + 1$; $\text{Hid}(A \ m) \leftarrow \text{hid}$</p> <p style="padding-left: 20px;">$\ A = g^{\hat{\alpha}^g} X^{\hat{\alpha}^x} Z^{\hat{\alpha}^z} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}^i} C_i^{\hat{\alpha}^i}$</p> <p style="padding-left: 20px;">$\delta_{\text{hid}} \leftarrow T(A \ m)$; $\hat{\alpha}_{\text{hid}} \leftarrow \hat{\alpha}$; $\hat{\beta}_{\text{hid}} \leftarrow \hat{\beta}$</p> <p>Return $T(A \ m)$</p>	<p>Oracle S₁ :</p> <p>$\text{sid} \leftarrow \text{sid} + 1$</p> <p>$a_{\text{sid}}, t_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p$</p> <p>$y_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p^*$</p> <p>$\text{st}_{\text{sid}}^s \leftarrow (a_{\text{sid}}, y_{\text{sid}}, t_{\text{sid}})$</p> <p>$A_{\text{sid}} \leftarrow g^{a_{\text{sid}}}$</p> <p>$C_{\text{sid}} \leftarrow g^{t_{\text{sid}}} Z^{y_{\text{sid}}}$</p> <p>$\text{msg}_1 \leftarrow (A_{\text{sid}}, C_{\text{sid}})$</p> <p>Return ($\text{sid}, \text{msg}_1$)</p> <p>Oracle S₂(i, c_i) :</p> <p>If $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$</p> <p style="padding-left: 20px;">or $c_i = 0$ then</p> <p style="padding-left: 20px;">Return \perp</p> <p style="padding-left: 20px;">$(a_i, y_i, t_i) \leftarrow \text{st}_i^s$</p> <p style="padding-left: 20px;">$s_i \leftarrow a_i + c_i \cdot y_i \cdot x$</p> <p style="padding-left: 20px;">$\text{msg}_2 \leftarrow (s_i, y_i, t_i)$</p> <p style="padding-left: 20px;">$\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$</p> <p style="padding-left: 20px;">$\ell \leftarrow \ell + 1$</p> <p>Return msg_2</p>
---	--

Figure 6.2: The OMUF security game for the blind signature scheme $\text{BS}_2[\text{GGen}]$.

checked, for each $k \in [Q_{\text{S}_1} + 1]$ and $j = \text{Hid}(\text{str}_k^*)$,¹ the following conditions hold:

$$\hat{\alpha}_j^x - \sum_{i \in [Q_{\text{S}_1}]} y_i \cdot c_i \cdot \hat{\alpha}_j^{A_i} = -\delta_j \cdot y_k^*, \quad (6.4)$$

$$\hat{\alpha}_j^z + \sum_{i \in [Q_{\text{S}_1}]} y_i \cdot \hat{\alpha}_j^{C_i} = y_k^*. \quad (6.5)$$

Since $\text{Adv}_{\text{BS}_2[\text{GGen}]}^{\text{omuf}}(\mathcal{A}_{\text{alg}}, \lambda) = \Pr[\text{WIN}] = \Pr[\text{WIN} \wedge E] + \Pr[\text{WIN} \wedge (\neg E)]$, the theorem follows by combining the following two lemmas with Theorem 4.1.1.

¹Here, $\text{Hid}(\text{str}_k^*)$ must be defined since a query str_k^* is made to H when checking the validity of the output (m_k^*, σ_k^*) .

Lemma 6.2.1. *There exists an adversary $\mathcal{B}_{\text{wfros}}$ for the $\text{WFROS}_{Q_{S_1}, p}$ problem making at most $Q_H + Q_{S_1} + 1$ queries to the random oracle H such that*

$$\text{Adv}_{Q_{S_1}, p}^{\text{wfros}}(\mathcal{B}_{\text{wfros}}) \geq \Pr[\text{WIN} \wedge E]. \quad (6.6)$$

Lemma 6.2.2. *There exists an adversary $\mathcal{B}_{\text{dlog}}$ for the DLog problem running in a similar running time as \mathcal{A}_{alg} such that*

$$\text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}, \lambda) \geq \frac{1}{2} \Pr[\text{WIN} \wedge (\neg E)]. \quad (6.7)$$

6.2.1 Proof of Lemma 6.2.1

We first give a detailed description of $\mathcal{B}_{\text{wfros}}$ playing the game $\text{WFROS}_{Q_{S_1}, p}$. To start with, $\mathcal{B}_{\text{wfros}}$ initializes sid , \mathcal{I}_{fin} , ℓ , T , hid , and Hid as described in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}$ game. In addition, $\mathcal{B}_{\text{wfros}}$ samples x, z uniformly from \mathbb{Z}_p , sets X to g^x and Z to g^z .

Then, $\mathcal{B}_{\text{wfros}}$ runs \mathcal{A}_{alg} on input (p, g, \mathbb{G}, X, Z) , and with access to the oracles $\tilde{\mathcal{S}}_1$, $\tilde{\mathcal{S}}_2$, and \tilde{H} . These oracles operate as follows:

Oracle $\tilde{\mathcal{S}}_1$: Same as the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}$ game except that instead of sampling $y_{\text{sid}}, t_{\text{sid}}$ randomly and setting $C_{\text{sid}} \leftarrow g^{t_{\text{sid}}} Z^{y_{\text{sid}}}$, $\mathcal{B}_{\text{wfros}}$ samples a new variable t'_{sid} uniformly from \mathbb{Z}_p and sets $C_{\text{sid}} = g^{t'_{\text{sid}}}$.

Oracle $\tilde{\mathcal{S}}_2$: After receiving a query (i, c_i) to $\tilde{\mathcal{S}}_2$ from \mathcal{A}_{alg} , if $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ or $c_i = 0$, $\mathcal{B}_{\text{wfros}}$ returns \perp . Otherwise, $\mathcal{B}_{\text{wfros}}$ makes a query (i, c_i) to S and uses its output as the value y_i . Also, $\mathcal{B}_{\text{wfros}}$ sets $t_i = t'_i - y_i \cdot z$. With the value (a_i, y_i, t_i) , the rest of $\tilde{\mathcal{S}}_2$ is the same as S_2 in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}$ game.

Oracle \tilde{H} : After receiving a query $(A \parallel m)$ to \tilde{H} from \mathcal{A}_{alg} , if $T(A \parallel m) \neq \perp$, the value $T(A \parallel m)$ is returned. Otherwise, since \mathcal{A}_{alg} is algebraic, $\mathcal{B}_{\text{wfros}}$ also knows the coefficient $\hat{\alpha}$ such that

$$A = g^{\hat{\alpha}^g} X^{\hat{\alpha}^x} Z^{\hat{\alpha}^z} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}^{A_i}} C_i^{\hat{\alpha}^{c_i}}.$$

Then, $\mathcal{B}_{\text{wfros}}$ issues the query (α, β) to H , where $\alpha, \beta \in \mathbb{Z}_p^{2Q_{S_1}+1}$ are such that

$$\alpha^{(i')} = \begin{cases} \hat{\alpha}^X, & i' = 0 \\ -\hat{\alpha}^{A_i}, & i' = 2i, i \in [\text{sid}] , \\ 0, & o.w. \end{cases} \quad (6.8)$$

$$\beta^{(i')} = \begin{cases} -\hat{\alpha}^Z, & i' = 0 \\ -\hat{\alpha}^{C_i}, & i' = 2i - 1, i \in [\text{sid}] . \\ 0, & o.w. \end{cases}$$

After receiving the output $(\delta_{\text{hid}}, \text{hid})$, $\mathcal{B}_{\text{wfros}}$ sets $T(A \parallel m) \leftarrow \delta_{\text{hid}}$ and $\text{Hid}(A \parallel m) \leftarrow \text{hid}$. Finally, $\mathcal{B}_{\text{wfros}}$ returns $T(A \parallel m)$.

After \mathcal{A}_{alg} outputs $\{(m_k^*, \sigma_k^*)\}_{k \in [Q_{S_1}+1]}$, $\mathcal{B}_{\text{wfros}}$ aborts if the conditions from the event $\text{WIN} \wedge E$ do not occur. Otherwise, $\mathcal{B}_{\text{wfros}}$ outputs $\mathcal{J} := \{\text{Hid}(\text{str}_k^*) \mid k \in [Q_{S_1} + 1]\}$. Since $\mathcal{B}_{\text{wfros}}$ simulates the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game perfectly, the probability that $\text{WIN} \wedge E$ occurs when running $\mathcal{B}_{\text{wfros}}$ is the same as in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game with \mathcal{A}_{alg} .

Following the similar analysis of \mathcal{B} in the GGM proof (Section 5.2.3), we know $\mathcal{B}_{\text{wfros}}$ makes at most $Q_H + Q_{S_1} + 1$ queries to H .

It is left to show that if $\text{WIN} \wedge E$ occurs within the simulation, then $\mathcal{B}_{\text{wfros}}$ wins the WFROS game. We first show that $|\mathcal{J}| = Q_{S_1} + 1$. Suppose $|\mathcal{J}| \leq Q_{S_1}$. Since \mathcal{A}_{alg} wins the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game, we know there exists $k_1, k_2 \in [Q_{S_1} + 1]$ such that $k_1 \neq k_2$ and $\text{Hid}(\text{str}_{k_1}^*) = \text{Hid}(\text{str}_{k_2}^*)$, which implies $\text{str}_{k_1}^* = \text{str}_{k_2}^*$. Therefore, we have

$$g^{s_{k_1}^*} X^{-c_{k_1}^*} \cdot y_{k_1}^* Z^{y_{k_1}^*} = g^{s_{k_2}^*} X^{-c_{k_2}^*} \cdot y_{k_2}^* Z^{y_{k_2}^*}, \quad m_{k_1}^* = m_{k_2}^*. \quad (6.9)$$

Also, let $j = \text{Hid}(\text{str}_{k_1}^*) = \text{Hid}(\text{str}_{k_2}^*)$. Since E occurs in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game simulated by $\mathcal{B}_{\text{wfros}}$, by (6.5), we have

$$y_{k_1}^* = \hat{\alpha}_j^X + \sum_{i \in [Q_{S_1}]} y_i \cdot \hat{\alpha}_j^{C_i} = y_{k_2}^* .$$

Since $y_{k_1}^* = y_{k_2}^*$ and $c_{k_1}^* = c_{k_2}^*$, by (6.9), we have $s_{k_1}^* = s_{k_2}^*$. However, since $(m_{k_1}^*, \sigma_{k_1}^*)$ and $(m_{k_2}^*, \sigma_{k_2}^*)$ are different message-signature pairs, we have

$$(m_{k_1}^*, c_{k_1}^*, s_{k_1}^*, y_{k_1}^*) \neq (m_{k_2}^*, c_{k_2}^*, s_{k_2}^*, y_{k_2}^*),$$

which yields a contradiction. Therefore, we have $|\mathcal{J}| = Q_{S_1} + 1$.

Then, since in particular E occurs, by (6.4) and (6.5), it holds that for any $j \in \mathcal{J}$

$$\hat{\alpha}_j^X - \sum_{i \in [Q_{S_1}]} y_i \cdot c_i \cdot \hat{\alpha}_j^{A_i} = -\delta_j \left(\hat{\alpha}_j^Z + \sum_{i \in [Q_{S_1}]} y_i \cdot \hat{\alpha}_j^{C_i} \right).$$

From the simulation of $\tilde{\mathbf{H}}$, by (6.8), we have for any $j \in \mathcal{J}$

$$\alpha_j^{(0)} + \sum_{i \in [Q_{S_1}]} y_i (\alpha_j^{(2i-1)} + c_i \cdot \alpha_j^{(2i)}) = \delta_j \left(\beta_j^{(0)} + \sum_{i \in [Q_{S_1}]} y_i (\beta_j^{(2i-1)} + c_i \cdot \beta_j^{(2i)}) \right).$$

Therefore, $\mathcal{B}_{\text{wfros}}$ wins the $\text{WFROS}_{Q_{S_1}, p}$ game, which concludes the proof.

6.2.2 Proof of Lemma 6.2.2

Proof. We first partition the event $\text{WIN} \wedge (\neg E)$ into two cases. Denote F_1 as the event in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game that there exists $k \in [Q_{S_1} + 1]$ such that (6.4) does not hold, and denote F_2 as the event that there exists $k \in [Q_{S_1} + 1]$ such that (6.5) does not hold. Then, if E does not occur, we know either F_1 or F_2 occurs. Therefore, we have $\text{WIN} \wedge (\neg E) = (\text{WIN} \wedge F_1) \vee (\text{WIN} \wedge F_2)$. We then prove the following two claims.

Claim 6.2.3. *There exists $\mathcal{B}_{\text{dlog}}^{(0)}$ for the DLog problem running in a similar running time as \mathcal{A}_{alg} such that $\Pr[\text{WIN} \wedge F_1] \leq \text{Adv}_{\mathbb{G}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(0)}, \lambda)$.*

Claim 6.2.4. *There exists $\mathcal{B}_{\text{dlog}}^{(1)}$ for the DLog problem running in a similar running time as \mathcal{A}_{alg} such that $\Pr[\text{WIN} \wedge F_2] \leq \text{Adv}_{\mathbb{G}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(1)}, \lambda)$.*

By the above two claims, we can construct an adversary $\mathcal{B}_{\text{dlog}}$ for the DLog problem that runs either $\mathcal{B}_{\text{dlog}}^{(0)}$ or $\mathcal{B}_{\text{dlog}}^{(1)}$ with 1/2 probability, and we can conclude the lemma since

$$\begin{aligned} \Pr[\text{WIN} \wedge (\neg E)] &\leq \Pr[\text{WIN} \wedge F_1] + \Pr[\text{WIN} \wedge F_2] \\ &\leq \text{Adv}_{\mathbb{G}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(0)}, \lambda) + \text{Adv}_{\mathbb{G}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(1)}, \lambda) = 2\text{Adv}_{\mathbb{G}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}, \lambda). \end{aligned}$$

□

Proof of Claim 6.2.3. We first give a detailed description of $\mathcal{B}_{\text{dlog}}^{(0)}$ playing the $\text{DLog}_{\mathbb{G}}$ game.

THE ADVERSARY $\mathcal{B}_{\text{dlog}}^{(0)}$. Initially, $\mathcal{B}_{\text{dlog}}^{(0)}$ initializes sid , \mathcal{I}_{fin} , ℓ , T , hid , and Hid as described in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\text{Aalg}}$ game. After $\mathcal{B}_{\text{dlog}}^{(0)}$ receives (p, g, \mathbb{G}, W) from the $\text{DLog}_{\mathbb{G}}$ game, $\mathcal{B}_{\text{dlog}}^{(0)}$ samples z uniformly from \mathbb{Z}_p and sets $X \leftarrow W, Z \leftarrow g^z$. Then, $\mathcal{B}_{\text{dlog}}^{(0)}$ runs \mathcal{A}_{alg} on input (p, g, \mathbb{G}, X) and with access to the oracles $\tilde{\mathcal{S}}_1, \tilde{\mathcal{S}}_2$, and $\tilde{\mathcal{H}}$. These oracles operate as follows:

Oracle $\tilde{\mathcal{S}}_1$: $\mathcal{B}_{\text{dlog}}^{(0)}$ samples $s_{\text{sid}}, t'_{\text{sid}}$ uniformly from \mathbb{Z}_p and y'_{sid} uniformly from \mathbb{Z}_p^* and sets $A_{\text{sid}} = g^{s_{\text{sid}}} X^{-y'_{\text{sid}}}$ and $C_{\text{sid}} = g^{t'_{\text{sid}}}$. Then, $\mathcal{B}_{\text{dlog}}^{(0)}$ returns $(\text{sid}, A_{\text{sid}}, C_{\text{sid}})$.

Oracle $\tilde{\mathcal{S}}_2$: Same as in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\text{Aalg}}$ game if $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ or $c_i = 0$. Otherwise, after receiving a query (i, c_i) to $\tilde{\mathcal{S}}_2$ from \mathcal{A}_{alg} , $\mathcal{B}_{\text{dlog}}^{(0)}$ sets $y_i \leftarrow y'_i/c_i$ and $t_i \leftarrow t'_i - y_i \cdot z$. Then, $\mathcal{B}_{\text{dlog}}^{(0)}$ returns (s_i, y_i, t_i) to \mathcal{A}_{alg} .

Oracle $\tilde{\mathcal{H}}$: Same as in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\text{Aalg}}$ game.

After receiving the output $\{(m_k^*, \sigma_k^*)\}_{k \in [Q_{S_1} + 1]}$, $\mathcal{B}_{\text{dlog}}^{(0)}$ aborts if the event $\text{WIN} \wedge F_1$ does not occur.

It is clear that $\mathcal{B}_{\text{dlog}}^{(0)}$ simulates the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\text{Aalg}}$ game perfectly. Therefore, it is left to show that if the event $\text{WIN} \wedge F_1$ occurs within the simulation, $\mathcal{B}_{\text{dlog}}^{(0)}$ can compute the discrete log of X , which equals to W .

Suppose $\text{WIN} \wedge F_1$ occurs. There exists $k \in [Q_{S_1} + 1]$ and $j = \text{Hid}(\text{str}_k^*)$ such that (6.4) does not hold. Since $\text{Hid}(\text{str}_k^*) = j$ and $\delta_j = c_k^*$, we have

$$g^{s_k^*} X^{-\delta_j \cdot y_k^*} Z^{y_k^*} = g^{\hat{\alpha}_j^g} X^{\hat{\alpha}_j^X} Z^{\hat{\alpha}_j^Z} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}_j^{A_i}} C_i^{\hat{\alpha}_j^{C_i}}. \quad (6.10)$$

Similar to the preceding case, since $\mathcal{B}_{\text{dlog}}^{(0)}$ knows the discrete log of Z as z , by substituting $A_i = g^{s_i} X^{-c_i \cdot y_i}$, $C_i = g^{t_i} Z^{y_i}$, and $Z = g^z$ into (6.10), we have

$$g^{s_k^* + y_k^* \cdot z} X^{-\delta_j \cdot y_k^*} = g^{\hat{\alpha}_j^g + \hat{\alpha}_j^Z \cdot z + \sum_{i \in [Q_{S_1}]} (\hat{\alpha}_j^{A_i} \cdot s_i + \hat{\alpha}_j^{C_i} \cdot (t_i + y_i \cdot z))} X^{\hat{\alpha}_j^X - \sum_{i \in [Q_{S_1}]} y_i \cdot c_i \cdot \hat{\alpha}_j^{A_i}}.$$

Since (6.4) does not hold, $\mathcal{B}_{\text{dlog}}^{(0)}$ can compute the discrete log of X as

$$x := \frac{s_k^* + y_k^* \cdot z - \hat{\alpha}_j^g - \hat{\alpha}_j^Z \cdot z - \sum_{i \in [Q_{S_1}]} (\hat{\alpha}_j^{A_i} \cdot s_i + \hat{\alpha}_j^{C_i} \cdot (t_i + y_i \cdot z))}{\hat{\alpha}_j^X - \sum_{i \in [Q_{S_1}]} y_i \cdot c_i \cdot \hat{\alpha}_j^{A_i} + \delta_j \cdot y_k^*}.$$

□

Proof of Claim 6.2.4. We first give a detailed description of $\mathcal{B}_{\text{dlog}}^{(1)}$ playing the $\text{DLog}_{\mathbb{G}}$ game.

THE ADVERSARY $\mathcal{B}_{\text{dlog}}^{(1)}$. Initially, $\mathcal{B}_{\text{dlog}}^{(1)}$ initializes sid , \mathcal{I}_{fin} , ℓ , T , hid , and Hid as described in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\text{Aalg}}$ game. After $\mathcal{B}_{\text{dlog}}^{(1)}$ receives (p, g, \mathbb{G}, W) from the $\text{DLog}_{\mathbb{G}}$ game, $\mathcal{B}_{\text{dlog}}^{(1)}$ samples x uniformly from \mathbb{Z}_p and sets $X \leftarrow g^x, Z \leftarrow W$. Then, $\mathcal{B}_{\text{dlog}}^{(1)}$ runs \mathcal{A}_{alg} on input (p, g, \mathbb{G}, X) and with access to the oracles $\tilde{\mathcal{S}}_1, \tilde{\mathcal{S}}_2$, and $\tilde{\mathcal{H}}$. Since $\mathcal{B}_{\text{dlog}}^{(1)}$ knows $X = g^x$, $\mathcal{B}_{\text{dlog}}^{(1)}$ can simulate all the oracles $\tilde{\mathcal{S}}_1, \tilde{\mathcal{S}}_2$, and $\tilde{\mathcal{H}}$ the same as in the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\text{Aalg}}$ game. After receiving the output $\{(m_k^*, \sigma_k^*)\}_{k \in [Q_{S_1} + 1]}$, $\mathcal{B}_{\text{dlog}}^{(1)}$ aborts if the event $\text{WIN} \wedge F_2$ does not occur.

It is clear that $\mathcal{B}_{\text{dlog}}^{(1)}$ simulates the $\text{OMUF}_{\text{BS}_2[\text{GGen}]}^{\text{Aalg}}$ game perfectly. Therefore, it is left to show that if the event $\text{WIN} \wedge F_2$ occurs within the simulation, $\mathcal{B}_{\text{dlog}}^{(1)}$ can compute the discrete log of Z , which equals to W .

Suppose $\text{WIN} \wedge F_2$ occurs. There exists $k \in [Q_{S_1} + 1]$ and $j = \text{Hid}(\text{str}_k^*)$ such that (6.5) does not hold. Since $\text{Hid}(\text{str}_k^*) = j$, we have

$$g^{s_k^*} X^{-c_k^* \cdot y_k^*} Z^{y_k^*} = g^{\hat{\alpha}_j^g} X^{\hat{\alpha}_j^x} Z^{\hat{\alpha}_j^z} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}_j^{A_i}} C_i^{\hat{\alpha}_j^{C_i}}. \quad (6.11)$$

From the simulation of $\tilde{\mathcal{S}}_2$, for each $i \in [Q_{S_1}]$, we have

$$g^{s_i} = A_i X^{c_i \cdot y_i}, \quad g^{t_i} = C_i Z^{-y_i}.$$

Also, $\mathcal{B}_{\text{dlog}}^{(1)}$ knows the discrete log of X as x . By substituting $A_i = g^{s_i} X^{-c_i \cdot y_i}$, $C_i = g^{t_i} Z^{y_i}$, and $X = g^x$ into (6.11), we have

$$g^{s_k^* - x \cdot c_k^* \cdot y_k^*} Z^{y_k^*} = g^{\hat{\beta}_j^g + \hat{\beta}_j^x \cdot x + \sum_{i \in [Q_{S_1}]} (\hat{\beta}_j^{A_i} \cdot (s_i - c_i \cdot y_i \cdot x) + \hat{\beta}_j^{C_i} \cdot t_i)} Z^{\hat{\beta}_j^z + \sum_{i \in [Q_{S_1}]} y_i \cdot \hat{\beta}_j^{C_i}}.$$

Since (6.5) does not hold, $\mathcal{B}_{\text{dlog}}^{(1)}$ can compute the discrete log of Z as

$$z := \frac{s_k^* - x \cdot c_k^* \cdot y_k^* - \hat{\beta}_j^g - \hat{\beta}_j^x \cdot x - \sum_{i \in [Q_{S_1}]} (\hat{\beta}_j^{A_i} \cdot (s_i - c_i \cdot y_i \cdot x) + \hat{\beta}_j^{C_i} \cdot t_i)}{\hat{\beta}_j^z + \sum_{i \in [Q_{S_1}]} y_i \cdot \hat{\beta}_j^{C_i} - y_k^*}.$$

□

6.3 Partially Blind Signatures

This section presents our partially blind signature scheme, PBS, which is detailed in Figure 6.3. The scheme builds on top of the BS_2 scheme by replacing the extra generator Z contained in the public key with the output of a hash function F (also modeled as a random oracle in the OMUF

<p><u>Algorithm Setup(1^λ) :</u> $(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\lambda)$ Select $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p^*$ Select $F : \{0, 1\}^* \rightarrow \mathbb{G}$ Return $par \leftarrow (p, \mathbb{G}, g, H, F)$</p> <p><u>Algorithm KeyGen(par) :</u> $(p, \mathbb{G}, g, H, F) \leftarrow par$ $x \leftarrow_{\\$} \mathbb{Z}_p$; $X \leftarrow g^x$ $sk \leftarrow x$; $pk \leftarrow X$ Return (sk, pk)</p> <p><u>Algorithm $S_1(sk, \text{info})$:</u> $x \leftarrow sk$; $X \leftarrow g^x$; $Z \leftarrow F(\text{info})$ $a, t \leftarrow_{\\$} \mathbb{Z}_p$; $y \leftarrow_{\\$} \mathbb{Z}_p^*$ $A \leftarrow g^a$; $C \leftarrow g^t Z^y$ $st^s \leftarrow (a, y, t, x)$; $msg_1 \leftarrow (A, C)$ Return (st^s, msg_1)</p> <p><u>Algorithm $S_2(st^s, c)$:</u> If $c = 0$ then return \perp $(a, y, t, x) \leftarrow st^s$ $s \leftarrow a + c \cdot y \cdot x$ Return $msg_2 \leftarrow (s, y, t)$</p>	<p><u>Algorithm $U_1(pk, msg_1, \text{info}, m)$:</u> $X \leftarrow pk$; $(A, C) \leftarrow msg_1$; $Z \leftarrow F(\text{info})$ $r_1 \leftarrow_{\\$} \mathbb{Z}_p$; $\gamma_1, \gamma_2 \leftarrow_{\\$} \mathbb{Z}_p^*$ $A' \leftarrow g^{r_1} \cdot A^{\gamma_1/\gamma_2} C^{\gamma_1}$ $c' \leftarrow H(\text{info} \parallel A' \parallel m)$ $c \leftarrow c' \cdot \gamma_2$ $st^u \leftarrow (c, c', r_1, \gamma_1, \gamma_2, X, Z, A, C)$ Return (st^u, c)</p> <p><u>Algorithm $U_2(st^u, msg_2)$:</u> $(c, c', r_1, \gamma_1, \gamma_2, X, Z, A, C) \leftarrow st^u$ $(s, y, t) \leftarrow msg_2$ If $y = 0$ or $C \neq g^t Z^y$ or $g^s \neq A \cdot X^{c \cdot y}$ then return \perp $s' \leftarrow r_1 + (\gamma_1/\gamma_2) \cdot s + \gamma_1 \cdot t$ $y' \leftarrow \gamma_1 \cdot y$ Return $\sigma \leftarrow (c', s', y')$</p> <p><u>Algorithm Ver($pk, \text{info}, \sigma, m$) :</u> $X \leftarrow pk$; $Z \leftarrow F(\text{info})$; $(c, s, y) \leftarrow \sigma$ If $y = 0$ then return 0 $A \leftarrow g^s \cdot X^{-c \cdot y} \cdot Z^y$ If $c \neq H(\text{info} \parallel A \parallel m)$ then return 0 Return 1</p>
--	---

Figure 6.3: The partially blind signature scheme $\text{PBS} = \text{PBS}[\text{GGen}]$.

proof) applied to the public input info . We do not formally redefine the syntax of partially blind signatures, but we note that it simply extends that of blind signatures by adding the extra input $\text{info} \in \{0, 1\}^*$ to the signer, the user, and the verification algorithm.

BLINDNESS. We first study the blindness of PBS. The $\text{PBlind}_{\text{PBS}}^A$ game is defined in Figure 6.4. The only difference between PBlind and Blind is that initially, the adversary \mathcal{A} also picks a public information info and interacts with PBS.U_1 and PBS.U_2 for signing (info, m_0) and (info, m_1) .

<p><u>Game PBlind_{PBS}^A(λ) :</u> $par \leftarrow \text{SetupBS}(1^\lambda)$ $b \leftarrow_s \{0, 1\}; b_0 \leftarrow b; b_1 \leftarrow 1-b$ $b' \leftarrow_s \mathcal{A}^{\text{INIT}, U_1, U_2}(par)$ If $b' = b$ then return 1 Return 0</p> <p><u>Oracle INIT($\tilde{pk}, \tilde{\text{info}}, \tilde{m}_0, \tilde{m}_1$) :</u> $\text{sess}_0 \leftarrow \text{init}$ $\text{sess}_1 \leftarrow \text{init}$ $pk \leftarrow \tilde{pk}$ $\text{info} \leftarrow \tilde{\text{info}}$ $m_0 \leftarrow \tilde{m}_0; m_1 \leftarrow \tilde{m}_1$</p>	<p><u>Oracle U₁($i, \text{msg}_1^{(i)}$) :</u> If $i \notin \{0, 1\}$ or $\text{sess}_i \neq \text{init}$ then return \perp $\text{sess}_i \leftarrow \text{open}$ $(\text{st}_i^u, \text{chl}^{(i)}) \leftarrow U_1(pk, \text{msg}_1^{(i)}, \text{info}, m_{b_i})$ Return $\text{chl}^{(i)}$</p> <p><u>Oracle U₂($i, \text{msg}_2^{(i)}$) :</u> If $i \notin \{0, 1\}$ or $\text{sess}_i \neq \text{open}$ then return \perp $\text{sess}_i \leftarrow \text{closed}$ $\sigma_{b_i} \leftarrow U_2(\text{st}_i^u, \text{msg}_2^{(i)})$ If $\text{sess}_0 = \text{sess}_1 = \text{closed}$ then If $\sigma_0 = \perp$ or $\sigma_1 = \perp$ then return (\perp, \perp) Return (σ_0, σ_1) Return (i, closed)</p>
---	--

Figure 6.4: The PBlind security game for a partially blind signature scheme PBS.

Denote the advantage of the adversary \mathcal{A} as

$$\text{Adv}_{\text{PBS}}^{\text{pblind}}(\mathcal{A}, \lambda) := \left| \Pr[\text{PBlind}_{\text{PBS}}^{\mathcal{A}}(\lambda) = 1] - \frac{1}{2} \right|.$$

We say a partially blind signature scheme PBS is perfectly blind if and only if $\text{Adv}_{\text{PBS}}^{\text{pblind}}(\mathcal{A}) = 0$ for any \mathcal{A} .

Theorem 6.3.1. *For a group generation algorithm \mathbb{G} , the partially blind signature scheme $\text{PBS}[\mathbb{G}\text{Gen}]$ is perfectly blind.*

Since the algorithm PBS.U_1 and PBS.U_2 are almost the same as $U_1\text{BS}_2$ and $U_2\text{BS}_2$, we can use a proof similar to the one for BS_2 (Section 6.1) to show $\text{PBS}[\mathbb{G}\text{Gen}]$ is perfectly blind. The only difference is that in BS_2 , Z is given in the public key, while in $\text{PBS}[\mathbb{G}\text{Gen}]$, Z is given by $F(\text{info})$.

OMUF SECURITY. We next study the OMUF security of PBS. Note that the definition must also be adjusted: The main difference is that the adversary wins as long as it can produce $\ell + 1$ valid message-signature pairs for some info for which it has run only ℓ signing sessions, regardless of

how many signing sessions are run with $\text{info}' \neq \text{info}$ (i.e., their number could be higher than ℓ). The corresponding game is defined in Figure 6.5, for the specific case of the scheme PBS. We prove the following theorem.

<p>Game OMUF$_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$($\lambda$): $(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$; $x \leftarrow_{\\$} \mathbb{Z}_p$; $X \leftarrow g^x$ $\text{sid} \leftarrow 0$; $\mathcal{I}_{\text{fin}} \leftarrow \emptyset$; $T_1 \leftarrow ()$; $T_2 \leftarrow ()$ $\ell \leftarrow$ a table where all entry are initially set to 0 $\text{fid} \leftarrow 0$; $\text{Fid} \leftarrow ()$; $\text{Hid} \leftarrow ()$ $(\text{info}^*, \{(m_k^*, \sigma_k^*)\}_{k \in [\ell(\text{info}^*)+1]})$ $\leftarrow_{\\$} \mathcal{A}_{\text{alg}}^{\text{S}_1, \text{S}_2, \text{H}, \text{F}}(p, g, \mathbb{G}, X)$ If $\exists k_1 \neq k_2 : (m_{k_1}^*, \sigma_{k_1}^*) = (m_{k_2}^*, \sigma_{k_2}^*)$ then Return 0 If $\exists k \in [\ell(\text{info}^*) + 1]$ such that $y_k^* = 0$ or $c_k^* \neq \text{H}(\text{info}^* \ g^{s_k^*} X^{-c_k^*} y_k^* Z^{y_k^*} \ m_k^*)$ where $(c_k^*, s_k^*, y_k^*) = \sigma_k^*$ and $Z = \text{F}(\text{info}^*)$ then return 0 Return 1</p> <p>Oracle H(info A m) : If $T_1(\text{info} \ A \ m) = \perp$ then $T_1(\text{info} \ A \ m) \leftarrow_{\\$} \mathbb{Z}_p$ $\text{hid} \leftarrow \text{hid} + 1$ $\text{Hid}(\text{info} \ A \ m) \leftarrow \text{hid}$ $\ A = g^{\hat{\alpha}^g} X^{\hat{\alpha}^x} \prod_{i \in [\text{fid}]} Z_i^{\hat{\alpha}^{Z_i}} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}^{A_i}} C_i^{\hat{\alpha}^{C_i}}$ $\delta_{\text{hid}} \leftarrow T_1(\text{info} \ A \ m)$ $\hat{\alpha}_{\text{hid}} \leftarrow \hat{\alpha}$; $\hat{\beta}_{\text{hid}} \leftarrow \hat{\beta}$ Return $T_1(\text{info} \ A \ m)$</p>	<p>Oracle S₁(info) : $Z \leftarrow \text{F}(\text{info})$ $\text{sid} \leftarrow \text{sid} + 1$; $\text{info}_{\text{sid}} \leftarrow \text{info}$ $a_{\text{sid}}, t_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p$; $y_{\text{sid}} \leftarrow_{\\$} \mathbb{Z}_p^*$ $\text{st}_{\text{sid}}^s \leftarrow (a_{\text{sid}}, y_{\text{sid}}, t_{\text{sid}})$ $A_{\text{sid}} \leftarrow g^{a_{\text{sid}}}$; $C_{\text{sid}} \leftarrow g^{t_{\text{sid}}} Z^{y_{\text{sid}}}$ $\text{msg}_1 \leftarrow (A_{\text{sid}}, C_{\text{sid}})$ Return $(\text{sid}, \text{msg}_1)$</p> <p>Oracle S₂(i, c_i) : If $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ or $c_i = 0$ then Return \perp $(a_i, y_i, t_i) \leftarrow \text{st}_i^s$ $s_i \leftarrow a_i + c_i \cdot y_i \cdot x$ $\text{msg}_2 \leftarrow (s_i, y_i, t_i)$ $\mathcal{I}_{\text{fin}} \leftarrow \mathcal{I}_{\text{fin}} \cup \{i\}$ $\mathcal{I}_{\text{fin}}^{(\text{info}_i)} \leftarrow \mathcal{I}_{\text{fin}}^{(\text{info}_i)} \cup \{i\}$ $\ell(\text{info}) \leftarrow \ell(\text{info}) + 1$ Return msg_2</p> <p>Oracle F(info) : If $T_2(\text{info}) = \perp$ then $T_2(\text{info}) \leftarrow_{\\$} \mathbb{G}$ $\text{fid} \leftarrow \text{fid} + 1$; $\text{Fid}(\text{info}) \leftarrow \text{fid}$ $Z_{\text{fid}} = T_2(\text{info})$ $\mathcal{I}_{\text{fin}}^{(\text{info})} \leftarrow \emptyset$ Return $T_2(\text{info})$</p>
--	---

Figure 6.5: The OMUF security game for the partially blind signature scheme PBS[GGen].

Theorem 6.3.2. *Let GGen be a group generation algorithm. Let \mathcal{A}_{alg} be an algebraic adversary for the game $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{PBS}[\text{GGen}]}(\lambda)$ such that for each public information info , makes at most Q_{S_1} queries to S_1 and Q_H queries to the random oracle H that start with info . Also, let the total number of distinct public information info 's queried by \mathcal{A}_{alg} to S_1 be bounded by Q_{info} . Then, there exists an adversary $\mathcal{B}_{\text{dlog}}$ for the DLog problem running in similar running time as \mathcal{A}_{alg} such that*

$$\text{Adv}_{\text{PBS}[\text{GGen}]}^{\text{omuf}}(\mathcal{A}_{\text{alg}}, \lambda) \leq 2\text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}, \lambda) + \frac{Q_{\text{info}}(Q_H + 3Q_{S_1} + 1)^2 + 2}{2^\kappa - 1}.$$

The proof is very similar to that for BS_2 except we need to additionally perform a hybrid argument over queries to F , guessing which info will be the one leading to a one-more forgery. However, we need to work harder here to ensure the discrete logarithm advantage does not scale with Q_{info} .

We also note that we have no argument supporting the fact that the information-theoretic term in Theorem 6.3.2 is tight and the inclusion of info in H is necessary. However, a tighter analysis appears to require studying a more general version of WFROS. We leave this to future work.

Proof of Theorem 6.3.2. Let \mathcal{A}_{alg} be an algebraic adversary described in the theorem. The game $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ is formally defined in Figure 6.5. Without loss of generality, we assume that if \mathcal{A}_{alg} outputs the public information info^* , then \mathcal{A}_{alg} makes exactly Q_{S_1} queries to S_1 and Q_{S_1} queries to S_2 that do not return \perp for info^* . Then, when \mathcal{A}_{alg} returns, we know $\ell(\text{info}^*) = Q_{S_1}$.

In the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game, the corresponding hid for each query $(\text{info} \| A \| m)$ to H is recorded in $\text{Hid}(\text{info} \| A \| m)$, and the output of the query is recorded as δ_{hid} . Also, since \mathcal{A}_{alg} is algebraic, \mathcal{A}_{alg} also provides the representation of A , and the corresponding coefficients α are recorded as α_{hid} . The corresponding fid for each new query info to S_1 is recorded in $\text{Fid}(\text{info})$. Also, $\mathcal{I}_{\text{fin}}^{(\text{info})}$ records the subset of \mathcal{I}_{fin} corresponding to signing sessions with public information info .

Denote the event WIN as \mathcal{A}_{alg} wins the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game, i.e., all the output message-signature pairs $\{m_k^*, \sigma_k^*\}_{k \in [Q_{S_1} + 1]}$ are distinct and valid for info^* . Furthermore, we denote $\text{str}_k^* := \text{info}^* \| g^{s_k^*} X^{-c_k^*} y_k^* Z_{\text{Fid}(\text{info}^*)}^{y_k^*} \| m_k^*$. We let E be the event in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game that after the validity of the output is checked, for each $k \in [Q_{S_1} + 1]$, $j = \text{Hid}(\text{str}_k^*)$, and $i^* = \text{Fid}(\text{info}^*)$ ², the

²Here, $\text{Hid}(\text{str}_k^*)$ must be defined since a query str_k^* is made to H when checking the validity of the output (m_k^*, σ_k^*) .

following conditions hold:

$$\hat{\alpha}^{Z_{i^*}} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(\text{info}^*)}} y_i \cdot \hat{\alpha}_j^{C_i} = y_k^*, \quad (6.12)$$

$$\hat{\alpha}_j^X - \sum_{i \in \mathcal{I}_{\text{fin}}} y_i \cdot c_i \cdot \hat{\alpha}_j^{A_i} = -\delta_j \cdot y_k^*, \quad (6.13)$$

$$\forall i \in [\text{sid}] \setminus \mathcal{I}_{\text{fin}} : \hat{\alpha}_j^{A_i} = 0. \quad (6.14)$$

Since $\text{Adv}_{\text{PBS}[\text{GGen}]}^{\text{omuf}}(\mathcal{A}_{\text{alg}}, \lambda) = \Pr[\text{WIN}] = \Pr[\text{WIN} \wedge E] + \Pr[\text{WIN} \wedge (\neg E)]$, the theorem follows by combining the following two lemmas with Theorem 4.1.1.

Lemma 6.3.3. *There exists an adversary $\mathcal{B}_{\text{wfros}}$ for the $\text{WFROS}_{Q_{S_1}, p}$ problem making at most $Q_{\text{H}} + Q_{S_1} + 1$ queries to the random oracle H such that*

$$\text{Adv}_{Q_{S_1}, p}^{\text{wfros}}(\mathcal{B}_{\text{wfros}}) \geq \frac{1}{Q_{\text{info}}} \Pr[\text{WIN} \wedge E]. \quad (6.15)$$

Lemma 6.3.4. *There exists an adversary $\mathcal{B}_{\text{dlog}}$ for the DLog problem running in a similar running time as \mathcal{A}_{alg} such that*

$$\Pr[\text{WIN} \wedge (\neg E)] \leq 2\text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}, \lambda) + \frac{2}{p-1}. \quad (6.16)$$

□

6.3.1 Proof of Lemma 6.3.3

We first give a detailed description of $\mathcal{B}_{\text{wfros}}$ playing the WFROS game.

THE ADVERSARY $\mathcal{B}_{\text{wfros}}$. To start with, $\mathcal{B}_{\text{wfros}}$ first samples a label \hat{i}^* uniformly from $[Q_{\text{info}}]$. Also, $\mathcal{B}_{\text{wfros}}$ samples x uniformly from \mathbb{Z}_p , sets X to g^x , and initializes sid , hid , fid , Hid , Fid , \mathcal{I}_{fin} , T_1 , and T_2 as described in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game. In addition, $\mathcal{B}_{\text{wfros}}$ initializes tfid to 0 and tFid to an empty table, which are used to record the labels of info queries to S_1 , and initializes tsid to 0 and tSid to an empty table, which are used to record the labels of session IDs for info such that $\text{tFid}(\text{info}) = \hat{i}^*$.

Then, $\mathcal{B}_{\text{wfros}}$ runs \mathcal{A}_{alg} on input (p, g, \mathbb{G}, X) and with access to the oracles $\tilde{\text{F}}$, \tilde{S}_1 , \tilde{S}_2 , and $\tilde{\text{H}}$. These oracles, operate as follows:

Oracles $\tilde{\text{F}}$: Same as in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game except instead of sampling $T_2(\text{info})$ uniformly from \mathbb{G} , if $T_2(\text{info}) = \perp$, $\mathcal{B}_{\text{wfros}}$ samples z_{fid} uniformly from \mathbb{Z}_p and sets $T_2(\text{info}) \leftarrow g^{z_{\text{fid}}}$.

Oracles \tilde{S}_1 : After receiving a query info to \tilde{S}_1 from \mathcal{A}_{alg} , if $\text{tFid}(\text{info}) = \perp$, $\mathcal{B}_{\text{wfros}}$ increases tfid by 1 and sets $\text{tFid}(\text{info}) = \text{tfid}$. Then, there are two cases:

- If $\text{tFid}(\text{info}) \neq \hat{i}^*$, $\mathcal{B}_{\text{wfros}}$ samples $s_{\text{sid}}, t'_{\text{sid}}$ uniformly from \mathbb{Z}_p and samples y'_{sid} uniformly from \mathbb{Z}_p^* . Then, $\mathcal{B}_{\text{wfros}}$ sets $A_{\text{sid}} = g^{s_{\text{sid}}} X^{-y'_{\text{sid}}}$ and $C_{\text{sid}} = g^{t'_{\text{sid}}}$.
- If $\text{tFid}(\text{info}) = \hat{i}^*$, $\mathcal{B}_{\text{wfros}}$ samples $a_{\text{sid}}, t'_{\text{sid}}$ uniformly from \mathbb{Z}_p and sets $A_{\text{sid}} = g^{a_{\text{sid}}}$ and $C_{\text{sid}} = g^{t'_{\text{sid}}}$. Also, $\mathcal{B}_{\text{wfros}}$ increases tsid by 1 and sets $\text{tSid}(\text{tsid}) \leftarrow \text{sid}$.

Finally, $\mathcal{B}_{\text{wfros}}$ returns $(\text{sid}, A_{\text{sid}}, C_{\text{sid}})$.

Oracles \tilde{S}_2 : After receiving a query (i, c_i) to \tilde{S}_2 from \mathcal{A}_{alg} , if $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ or $c_i = 0$, $\mathcal{B}_{\text{wfros}}$ returns \perp . Otherwise, there are two cases:

- If $\text{tFid}(\text{info}_i) \neq \hat{i}^*$, $\mathcal{B}_{\text{wfros}}$ computes $y_i \leftarrow y'_i/c_i$ and $t_i \leftarrow t'_i - y_i \cdot z_{\text{Fid}(\text{info}_i)}$.
- If $\text{tFid}(\text{info}_i) = \hat{i}^*$, let i' be the index in $[\text{sid}]$ such that $\text{tSid}(i') = i$ and $\mathcal{B}_{\text{wfros}}$ sets $\tilde{c}_{i'} \leftarrow c_i$. Then, $\mathcal{B}_{\text{wfros}}$ makes a query $(i', \tilde{c}_{i'})$ to S . After $\mathcal{B}_{\text{wfros}}$ receives $\tilde{y}_{i'}$ from S , \mathcal{B} sets $y_i \leftarrow \tilde{y}_{i'}$ and $t_i \leftarrow t'_i - y_i \cdot z_{\text{Fid}(\text{info}_i)}$.

Finally, $\mathcal{B}_{\text{wfros}}$ returns (s_i, y_i, t_i) .

Oracles \tilde{H} : After receiving a query $(\text{info} \| A \| m)$ to \tilde{H} from \mathcal{A}_{alg} , if $T_1(\text{info} \| A \| m) \neq \perp$ or $\text{tFid}(\text{info}) \neq \hat{i}^*$, then \tilde{H} is the same as H in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game. Otherwise, since \mathcal{A}_{alg} is algebraic, $\mathcal{B}_{\text{wfros}}$ also knows $\hat{\alpha}$ such that

$$A = g^{\hat{\alpha}^g} X^{\hat{\alpha}^x} \prod_{i \in [\text{fid}]} Z_i^{\hat{\alpha}^{z_i}} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}^{a_i}} C_i^{\hat{\alpha}^{c_i}} .$$

Then, $\mathcal{B}_{\text{wfros}}$ issues the query (α, β) to H , where $\alpha, \beta \in \mathbb{Z}_p^{2Q_{s_1} + 1}$ such that

$$\alpha^{(i')} = \begin{cases} \hat{\alpha}^x - \sum_{i \in [\text{sid}], \text{tFid}(\text{info}_i) \neq \hat{i}^*} \hat{\alpha}^{a_i} \cdot y'_i, & i' = 0 \\ -\hat{\alpha}^{A_{\text{tSid}(i)}}, & i' = 2i, i \in [\text{tsid}] , \\ 0, & o.w. \end{cases} \quad (6.17)$$

$$\beta^{(i')} = \begin{cases} -\hat{\alpha}^{Z_{i^*}}, & i' = 0 \\ -\hat{\alpha}^{C_{\text{tSid}(i)}}, & i' = 2i - 1, i \in [\text{tsid}] . \\ 0, & o.w. \end{cases} \quad (6.18)$$

After receiving the output $(\delta_{\text{hid}}, \text{hid})$, $\mathcal{B}_{\text{wfros}}$ sets $T_1(\text{info} \parallel A \parallel m) \leftarrow \delta_{\text{hid}}$ and $\text{Hid}(\text{info} \parallel A \parallel m) \leftarrow \text{hid}$. Finally, $\mathcal{B}_{\text{wfros}}$ returns $T_1(\text{info} \parallel A \parallel m)$.

After receiving the output $\{\text{info}^*, (m_k^*, \sigma_k^*)\}_{k \in [Q_{S_1} + 1]}$ from \mathcal{A}_{alg} , $\mathcal{B}_{\text{wfros}}$ aborts if the conditions from the event $\text{WIN} \wedge E$ do not occur. Otherwise, $\mathcal{B}_{\text{wfros}}$ outputs $\mathcal{J} := \{\text{Hid}(\text{str}_k^*) \mid k \in [Q_{S_1} + 1]\}$.

ANALYSIS OF $\mathcal{B}_{\text{wfros}}$. Note that $\mathcal{B}_{\text{wfros}}$ makes a query to H at most once when it receives a query to $\tilde{\text{H}}$ for $\text{info}_{\text{tfid}}$ and at most $Q_{S_1} + 1$ more queries to $\tilde{\text{H}}$ when checking the validity of the output. Therefore, \mathcal{B} makes at most $Q_{\text{H}} + Q_{S_1} + 1$ queries to H . Also, it is clear that \mathcal{B} simulates oracles F , S_1 , S_2 , H in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game perfectly no matter what label is assigned to tfid . Therefore, the probability that $\text{tFid}(\text{info}^*) = \hat{i}^*$ and $\text{WIN} \wedge E$ occurs when running $\mathcal{B}_{\text{wfros}}$ is equal to $\frac{1}{Q_{\text{info}}} \Pr[\text{WIN} \wedge E]$.

It is left to show that if $\text{tFid}(\text{info}^*) = \hat{i}^*$ and $\text{WIN} \wedge E$ occurs within the simulation, then $\mathcal{B}_{\text{wfros}}$ wins the WFROS game. Suppose $\text{WIN} \wedge E$ occurs and $\text{tFid}(\text{info}^*) = \hat{i}^*$. Following the similar analysis of $\mathcal{B}_{\text{wfros}}$ in the proof of Lemma 6.2.1, we have $|\mathcal{J}| = Q_{S_1} + 1$.

Denote $\mathcal{I}_{\text{fin}}^{\text{tot}}$ and sid^{tot} as the values of \mathcal{I}_{fin} and sid when \mathcal{A}_{alg} returns. Then, since E occurs, by (6.12) and (6.13), for any $j \in \mathcal{J}$ it holds that

$$\hat{\alpha}_j^{\text{X}} - \sum_{i \in \mathcal{I}_{\text{fin}}^{\text{tot}}} y_i \cdot c_i \cdot \hat{\alpha}_j^{\text{A}_i} = -\delta_j \left(\hat{\alpha}_j^{\text{Z}_{i^*}} + \sum_{i \in \mathcal{I}_{\text{fin}}^{\text{(info}^*)}} y_i \cdot \hat{\alpha}_j^{\text{C}_i} \right). \quad (6.19)$$

Then, by (6.14), we have

$$\begin{aligned} -\delta_j \left(\hat{\alpha}_j^{\text{Z}_{i^*}} + \sum_{i \in \mathcal{I}_{\text{fin}}^{\text{(info}^*)}} y_i \cdot \hat{\alpha}_j^{\text{C}_i} \right) x &= \hat{\alpha}_j^{\text{X}} - \sum_{i \in \mathcal{I}_{\text{fin}}^{\text{tot}}} y_i \cdot c_i \cdot \hat{\alpha}_j^{\text{A}_i} \\ &= \hat{\alpha}_j^{\text{X}} - \sum_{i \in \mathcal{I}_{\text{fin}}^{\text{tot}}} y_i \cdot c_i \cdot \hat{\alpha}_j^{\text{A}_i} - \sum_{i \in [\text{sid}^{\text{tot}}] \setminus \mathcal{I}_{\text{fin}}^{\text{tot}}} y'_i \cdot \hat{\alpha}_j^{\text{A}_i} \\ &= \hat{\alpha}_j^{\text{X}} - \sum_{i \in [\text{sid}^{\text{tot}}], \text{tFid}(\text{info}_i) \neq \hat{i}^*} y'_i \cdot \hat{\alpha}_j^{\text{A}_i} - \sum_{i \in \mathcal{I}_{\text{fin}}^{\text{(info}^*)}} y_i \cdot c_i \cdot \hat{\alpha}_j^{\text{A}_i}. \end{aligned}$$

Then, from the simulation, by (6.17), we have for any $j \in \mathcal{J}$

$$\alpha_j^{(0)} + \sum_{i \in [Q_{S_1}]} \tilde{y}_i (\alpha_j^{(2i-1)} + \tilde{c}_i \cdot \alpha_j^{(2i)}) = \delta_j \left(\beta_j^{(0)} + \sum_{i \in [Q_{S_1}]} \tilde{y}_i (\beta_j^{(2i-1)} + \tilde{c}_i \cdot \beta_j^{(2i)}) \right).$$

Therefore, $\mathcal{B}_{\text{wfros}}$ wins the $\text{WFROS}_{Q_{S_1}, p}$ game.

<p style="margin: 0;">Game rel-DLog$_{\text{GGen},n}^{\mathcal{A}}$($\lambda$) :</p> <p style="margin: 0;">$(\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\kappa)$</p> <p style="margin: 0;">$\{X_i\}_{i \in [n]} \leftarrow_{\text{s}} \mathbb{G}$</p> <p style="margin: 0;">$y_0, y_1, \dots, y_n \leftarrow \mathcal{A}(p, g, \mathbb{G}, \{X_i\}_{i \in [n]})$</p> <p style="margin: 0;">If $\forall i \in \{1, \dots, n\} : y_i = 0$ then return 0</p> <p style="margin: 0;">If $g^{y_0} \prod_{i \in [n]} X_i^{y_i} = 1_{\mathbb{G}}$ then return 1</p> <p style="margin: 0;">Return 0</p>
--

Figure 6.6: The rel-DLog game.

6.3.2 Proof of Lemma 6.3.4

We first partition the event $\text{WIN} \wedge (\neg E)$ into two cases. Denote F_1 as the event in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game that there exists $k \in [Q_{S_1} + 1]$ such that either (6.13) or (6.14) does not hold, and denote F_2 as the event that there exists $k \in [Q_{S_1} + 1]$ such that (6.12) does not hold. Then, if E does not occur, we know either F_1 or F_2 occurs. Therefore, we have $\text{WIN} \wedge (\neg E) = (\text{WIN} \wedge F_1) \vee (\text{WIN} \wedge F_2)$. For the case that $\text{WIN} \wedge F_1$ occurs, we show the following claim.

Claim 6.3.5. *There exists $\mathcal{B}_{\text{dlog}}^{(0)}$ for the DLog problem running in a similar running time as \mathcal{A}_{alg} such that*

$$\Pr[\text{WIN} \wedge F_1] \leq \text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(0)}, \lambda) + \frac{1}{p-1}. \quad (6.20)$$

For the case that $\text{WIN} \wedge F_2$ occurs, we construct an adversary $\mathcal{B}_{\text{rel-dlog}}$ for the $\text{rel-DLog}_{\mathbb{G}, Q_{\text{F}}}$ game (defined in 6.6) with advantage equals to the probability that $\text{WIN} \wedge F_2$ occurs, where Q_{F} denotes the maximum number of queries to F issued in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game, and we summarize it into the following claim.

Claim 6.3.6. *There exists $\mathcal{B}_{\text{rel-dlog}}$ for the rel-DLog problem running in a similar running time as \mathcal{A}_{alg} such that*

$$\Pr[\text{WIN} \wedge F_2] \leq \text{Adv}_{\text{GGen}, Q_{\text{F}}}^{\text{rel-dlog}}(\mathcal{B}_{\text{rel-dlog}}, \lambda). \quad (6.21)$$

The rel-DLog problem is equivalent to the DLog problem, as shown in the following lemma from [84].

Lemma 6.3.7 (Lemma 3 in [84]³). *Let GGen be a group generation algorithm. For any $n > 0$ and any adversary $\mathcal{B}_{\text{rel-dlog}}$ for the $\text{rel-DLog}_{\text{GGen},n}$ game, there exists an adversary $\mathcal{B}_{\text{dlog}}$ for the $\text{DLog}_{\text{GGen}}$ game such that*

$$\text{Adv}_{\text{GGen},n}^{\text{rel-dlog}}(\mathcal{B}_{\text{rel-dlog}}, \lambda) \leq \text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}, \lambda) + 1/p.$$

By the lemma and Claim 6.3.6, there exists an adversary $\mathcal{B}_{\text{dlog}}^{(1)}$ for the $\text{DLog}_{\text{GGen}}$ problem such that $\Pr[\text{WIN} \wedge F_1] \leq \text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(1)}, \lambda) + \frac{1}{p}$. Therefore, together with Claim 6.3.5, we can construct an adversary $\mathcal{B}_{\text{dlog}}$ for the $\text{DLog}_{\text{GGen}}$ problem that runs either $\mathcal{B}_{\text{dlog}}^{(0)}$ or $\mathcal{B}_{\text{dlog}}^{(1)}$ with $1/2$ probability, and we can conclude the lemma since

$$\begin{aligned} \Pr[\text{WIN} \wedge (\neg E)] &\leq \Pr[\text{WIN} \wedge F_1] + \Pr[\text{WIN} \wedge F_2] \\ &\leq \text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(0)}, \lambda) + \text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}^{(1)}, \lambda) + \frac{2}{p-1} = 2\text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}_{\text{dlog}}, \lambda) + \frac{2}{p-1}. \end{aligned}$$

Proof of Claim 6.3.5. We first give a detailed description of $\mathcal{B}_{\text{dlog}}^{(0)}$ playing the $\text{DLog}_{\text{GGen}}$ game.

THE ADVERSARY $\mathcal{B}_{\text{dlog}}^{(0)}$. To start with, $\mathcal{B}_{\text{dlog}}^{(0)}$ initializes sid , hid , fid , Hid , Fid , \mathcal{I}_{fin} , T_1 , and T_2 as described in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game. After $\mathcal{B}_{\text{dlog}}^{(0)}$ receives (p, g, \mathbb{G}, W) from the $\text{DLog}_{\text{GGen}}$ game, sets $X \leftarrow W$. Then, $\mathcal{B}_{\text{dlog}}^{(0)}$ runs \mathcal{A}_{alg} on input (p, g, \mathbb{G}, X) , and with access to the oracles $\tilde{\text{F}}$, $\tilde{\text{S}}_1$, $\tilde{\text{S}}_2$, and $\tilde{\text{H}}$. These oracles operate as follows:

Oracle $\tilde{\text{F}}$: Same as in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game except instead of sampling $T_2(\text{info})$ uniformly from \mathbb{G} , if $T_2(\text{info}) = \perp$, $\mathcal{B}_{\text{dlog}}^{(0)}$ samples z uniformly from \mathbb{Z}_p and sets $T_2(\text{info}) \leftarrow g^{z\text{fid}}$.

Oracle $\tilde{\text{S}}_1$: After receiving a query info from \mathcal{A}_{alg} , $\mathcal{B}_{\text{wfros}}$ samples $s_{\text{sid}}, t'_{\text{sid}}$ uniformly from \mathbb{Z}_p and samples y'_{sid} uniformly from \mathbb{Z}_p^* . Then, $\mathcal{B}_{\text{wfros}}$ sets $A_{\text{sid}} \leftarrow g^{s_{\text{sid}}} X^{-y'_{\text{sid}}}$ and $C_{\text{sid}} \leftarrow g^{t'_{\text{sid}}}$ and returns $(\text{sid}, A_{\text{sid}}, C_{\text{sid}})$.

Oracle $\tilde{\text{S}}_2$: After receiving a query (i, c_i) to $\tilde{\text{S}}_2$ from \mathcal{A}_{alg} , if $i \notin [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$ or $c_i = 0$, $\mathcal{B}_{\text{dlog}}^{(0)}$ returns \perp . Otherwise, let $\hat{i} := \text{Fid}(\text{info}_i)$, and $\mathcal{B}_{\text{dlog}}^{(0)}$ computes $y_i \leftarrow y'_i/c_i$ and $t_i \leftarrow t'_i - y_i \cdot z_{\hat{i}}$. Then, $\mathcal{B}_{\text{dlog}}^{(0)}$ returns (s_i, y_i, t_i) .

³The DLog and rel-DLog games defined in [84] differ slightly from our descriptions, but the lemma follows by a similar proof.

Oracle $\tilde{\mathbf{H}}$: Same as in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game.

After receiving the output $(\text{info}^*, \{(m_k^*, \sigma_k^*)\}_{k \in [Q_{S_1}+1]})$, $\mathcal{B}_{\text{dlog}}^{(0)}$ aborts if the event $\text{WIN} \wedge F_1$ does not occur.

It is clear that $\mathcal{B}_{\text{dlog}}^{(0)}$ simulates the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game perfectly, and thus it is left to show that if $\text{WIN} \wedge F_1$ occurs, $\mathcal{B}_{\text{dlog}}^{(0)}$ can compute the discrete log of W except for probability $1/p$.

Suppose $\text{WIN} \wedge F_1$ occurs in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\mathcal{A}_{\text{alg}}}$ game simulated by $\mathcal{B}_{\text{dlog}}^{(0)}$. There exists $k \in [Q_{S_1}+1]$ and $j = \text{Hid}(\text{str}_k^*)$ such that either (6.13) or (6.14) does not hold. Since $j = \text{Hid}(\text{str}_k^*)$ and $\delta_j = c_k^*$, we have

$$g^{s_k^*} X^{-\delta_j \cdot y_k^*} Z_{i^*}^{y_k^*} = g^{\hat{\alpha}_j^g} X^{\hat{\alpha}_j^x} \prod_{i \in [\text{fid}]} Z_i^{\hat{\alpha}_j^z} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}_j^a} C_i^{\hat{\alpha}_j^c}. \quad (6.22)$$

From the simulation of \tilde{S}_1 , for each $i \in [\text{sid}]$, we have

$$A_i = g^{s_i} X^{-y'_i}, \quad C_i = g^{t'_i}.$$

Also, $\mathcal{B}_{\text{dlog}}^{(0)}$ knows the discrete log of Z_i as z_i for each $i \in [\text{fid}]$. By substituting $A_i = g^{s_i} X^{-y'_i}$, $C_i = g^{t'_i}$, and $Z_i = g^{z_i}$ into (6.22), we have

$$g^{s_k^* + y_k^* \cdot z_{i^*}} X^{-\delta_j \cdot y_k^*} = g^{\eta_j^g} X^{\eta_j^z},$$

where

$$\begin{aligned} \eta_j^g &:= \hat{\alpha}_j^g + \sum_{i \in [\text{fid}]} \hat{\alpha}_j^z \cdot z_i + \sum_{i \in [\text{sid}]} (\hat{\alpha}_j^a \cdot s_i + \hat{\alpha}_j^c \cdot t'_i), \\ \eta_j^x &:= \hat{\alpha}_j^x - \sum_{i \in [\text{sid}]} y'_i \cdot \hat{\alpha}_j^a. \end{aligned}$$

If $\eta_j^x \neq -\delta_j \cdot y_k^*$, $\mathcal{B}_{\text{dlog}}^{(0)}$ can compute the discrete log of X , which is also W , as

$$x := \frac{s_k^* + y_k^* \cdot z_{i^*} - \eta_j^g}{\eta_j^x + \delta_j \cdot y_k^*}.$$

Therefore, it is left to bound the probability that $\eta_j^x = -\delta_j \cdot y_k^*$, and there are the following two cases.

(6.14) does not hold for k, j . Consider the transcript π^{tot} that the adversary sees before it returns. Given the transcript π^{tot} , since for each $i \in [\text{sid}] \setminus \mathcal{I}_{\text{fin}}$, the adversary sees only A_i but does not know

either s_i or y'_i , the value y'_i is uniformly distributed over \mathbb{Z}_p^* independent of all other $y'_{i'}$ for $i' \neq i$. Therefore, the probability that $\eta_j^X = -\delta_j \cdot y_k^*$ is $\frac{1}{p-1}$.

(6.14) holds but (6.13) does not hold for k, j . Since (6.14) holds and for each $i \in \mathcal{I}_{\text{fin}}$ it holds that $y'_i = y_i \cdot c_i$, we have

$$\eta_j^X = \hat{\alpha}_j^X - \sum_{i \in \mathcal{I}_{\text{fin}}} y_i \cdot c_i \cdot \hat{\alpha}_j^{A_i}.$$

Then, since (6.13) does not hold, we have

$$\eta_j^X \neq -\delta_j \cdot y_k^*,$$

which means the probability that $\eta_j^X = -\delta_j \cdot y_k^*$ is 0. Therefore, for both cases, the probability that $\eta_j^X = -\delta_j \cdot y_k^*$ is bounded by $\frac{1}{p-1}$. \square

Proof of Claim 6.3.6. We first give a detailed description of $\mathcal{B}_{\text{rel-dlog}}$ playing the $\text{rel-DLog}_{\mathbb{G}, Q_F}$ game.

THE ADVERSARY $\mathcal{B}_{\text{rel-dlog}}$. To start with, $\mathcal{B}_{\text{rel-dlog}}$ initializes sid , hid , fid , Hid , Fid , \mathcal{I}_{fin} , T_1 , and T_2 as described in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game. Also, $\mathcal{B}_{\text{rel-dlog}}$ samples x uniformly from \mathbb{Z}_p and sets $X \leftarrow g^x$. After $\mathcal{B}_{\text{rel-dlog}}$ receives $(p, g, \mathbb{G}, Z_1, \dots, Z_{Q_F})$ from the $\text{rel-DLog}_{\mathbb{G}, Q_F}$ game, $\mathcal{B}_{\text{rel-dlog}}$ runs \mathcal{A}_{alg} on input (p, g, \mathbb{G}, X) and with access to the oracles $\tilde{\mathbf{F}}$, $\tilde{\mathbf{S}}_1$, $\tilde{\mathbf{S}}_2$, and $\tilde{\mathbf{H}}$. These oracles operate as follows:

Oracle $\tilde{\mathbf{F}}$: Same as in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game except instead of sampling $T_2(\text{info})$ uniformly from \mathbb{G} , if $T_2(\text{info}) = \perp$, $\mathcal{B}_{\text{rel-dlog}}$ sets $T_2(\text{info}) \leftarrow Z_{\text{fid}}$.

Oracle $\tilde{\mathbf{S}}_1, \tilde{\mathbf{S}}_2, \tilde{\mathbf{H}}$: The same as in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game.

After receiving the output $(\text{info}^*, \{(m_k^*, \sigma_k^*)\}_{k \in [Q_{S_1} + 1]})$, $\mathcal{B}_{\text{rel-dlog}}$ aborts if $\text{WIN} \wedge F_2$ does not occur.

It is clear that $\mathcal{B}_{\text{rel-dlog}}$ simulates the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game perfectly, and thus it is left to show that if $\text{WIN} \wedge F_2$ occurs, $\mathcal{B}_{\text{rel-dlog}}$ can win the $\text{rel-DLog}_{\mathbb{G}, Q_F}$ game.

Suppose $\text{WIN} \wedge F_2$ occurs in the $\text{OMUF}_{\text{PBS}[\text{GGen}]}^{\text{Aalg}}$ game simulated by $\mathcal{B}_{\text{rel-dlog}}$. There exists $k \in [Q_{S_1} + 1]$ and $j = \text{Hid}(\text{str}_k^*)$ such that (6.12) does not hold. Since $j = \text{Hid}(\text{str}_k^*)$, we have

$$g^{s_k^*} X^{-\delta_j \cdot y_k^*} Z_{i^*}^{y_k^*} = g^{\hat{\alpha}_j^g} X^{\hat{\alpha}_j^X} \prod_{i \in [\text{fid}]} Z_i^{\hat{\alpha}_j^{Z_i}} \prod_{i \in [\text{sid}]} A_i^{\hat{\alpha}_j^{A_i}} C_i^{\hat{\alpha}_j^{C_i}}. \quad (6.23)$$

From the simulation of \widetilde{S}_1 , for each $i \in [\text{sid}]$, we have

$$A_i = g^{a_i}, \quad g^{t_i} = C_i Z_i^{-y_i}.$$

Also, $\mathcal{B}_{\text{rel-dlog}}$ knows the discrete log of X as x . By substituting $A_i = g^{a_i}$, $C_i = g^{t_i} Z_i^{y_i}$, and $X = g^x$ into (6.23), we have

$$g^{s_k^*} X^{-\delta_j \cdot y_k^*} Z_{i^*}^{y_k^*} = g^{\hat{\alpha}_j^g + \hat{\alpha}_j^x \cdot x + \sum_{i \in [\text{sid}]} (\hat{\alpha}_j^{A_i} \cdot a_i + \hat{\alpha}_j^{C_i} \cdot t_i)} \prod_{i \in [\text{fid}]} Z_i^{\hat{\alpha}_j^{Z_i} + \sum_{i' \in [\text{sid}], \text{tSid}(i')=i} y_{i'} \cdot \hat{\alpha}_j^{C_{i'}}}.$$

Therefore, $\mathcal{B}_{\text{rel-dlog}}$ can compute (w_0, \dots, w_{Q_F}) such that $g^{w_0} \prod_{i \in [Q_F]} W_i^{w_i} = g^{w_0} \prod_{i \in [\text{fid}]} Z_i^{w_i} = 1_{\mathbb{G}_\lambda}$ as

$$w_i := \begin{cases} \hat{\alpha}_j^g + \hat{\alpha}_j^x \cdot x + \sum_{i \in [\text{sid}]} (\hat{\alpha}_j^{A_i} \cdot a_i + \hat{\alpha}_j^{C_i} \cdot t_i) - s_k^* + x \cdot \delta_j \cdot y_k^*, & i = 0 \\ \hat{\alpha}_j^{Z_i} + \sum_{i' \in [\text{sid}], \text{tSid}(i')=i} y_{i'} \cdot \hat{\alpha}_j^{C_{i'}}, & i \in [\text{fid}], i \neq i^* \\ -y_k^* + \hat{\alpha}_j^{Z_i} + \sum_{i' \in [\text{sid}], \text{tSid}(i')=i} y_{i'} \cdot \hat{\alpha}_j^{C_{i'}}, & i = i^* \\ 0, & o.w. \end{cases}$$

Since (6.12) does not hold, we have

$$w_{i^*} = -y_k^* + \hat{\alpha}_j^{Z_{i^*}} + \sum_{i \in \mathcal{I}_{\text{fin}}^{(\text{info}^*)}} y_{i'} \cdot \hat{\alpha}_j^{C_{i'}} \neq 0.$$

Therefore, $\mathcal{B}_{\text{rel-dlog}}$ wins the $\text{rel-DLog}_{\mathbb{G}, Q_F}$ game by outputting (w_0, \dots, w_{Q_F}) defined above. \square

Part II

THRESHOLD SIGNATURES IN PAIRING-FREE GROUPS

Chapter 7

THRESHOLD SIGNATURES: INTRODUCTION

Threshold signatures, which originated in the late 1980s [58, 59], are seeing renewed attention, driven in particular by an interest in using them to secure digital wallets in the cryptocurrencies ecosystem [70]. Parallel IETF [92] and NIST [102] standardization efforts are evidence as to the speed at which the area is moving into practice.

Whether securing a user’s digital wallet, or being used by a CA to create a certificate, forgery of a digital signature is costly. The rising tide of system breaches and phishing attacks makes exposure of a signing key too likely to ignore. The idea of a threshold signature scheme is to distribute the secret signing key across multiple parties who then interact to produce a signature, the intent being to retain security even in the face of compromise of up to a threshold number of these parties. Over the years, threshold versions of many schemes have been presented, including RSA [53, 74, 119], DSA/ECDSA [71, 73, 70, 28, 69, 96, 38], Schnorr signatures [120, 72, 91] and BLS signatures [27].

In this thesis, we focus on FROST [91], the state-of-the-art construction of threshold Schnorr signatures, and its variants, referred to as FROST2 [50] and FROST3 [113]. (We also refer to the original version of FROST as FROST1.) FROST is a pairing-free scheme that is *partially* non-interactive, meaning signing additionally involves a message-independent pre-processing round.

Our first main contribution is to show that FROST and its variants satisfy stronger security definitions than those previously considered in the literature or proven for these schemes. These refined definitions capture natural and practically relevant strengths that may be desirable in real-world applications. Moreover, our fine-grained perspective reveals meaningful differences between existing schemes. In particular, we demonstrate that FROST2 is strictly less secure than FROST1, and FROST3 is strictly less secure than FROST2. We elaborate on this in Section 7.1.

Our second main contribution is a new approach that can transform FROST into threshold signature schemes whose security relies on weaker assumptions. Specifically, while all three FROST variants rely on the Algebraic One-More Discrete Logarithm (AOMDL) assumption [105] in the Random Oracle Model (ROM), our new schemes are provably secure under the plain DL assump-

tion in the ROM. The resulting schemes produce Okamoto signatures [107], which are just one scalar larger than Schnorr signatures. Further details are provided in Section 7.2.

7.1 Stronger Security for Non-Interactive Threshold Signatures

OUR PATH. The classical development paradigm in theoretical cryptography is to ask what security we would like, define it, and then seek schemes that meet it. Yet if we look back, there has been another path alongside; canonical, reference schemes guided a choice of definitions that model them, and, once made, these definitions went on to be influential targets for future schemes. (The formal definition of trapdoor permutations [78], for example, was crafted to model RSA.) We are inspired by the latter path. Within the space of Schnorr threshold schemes, FROST appears to be the simplest candidate. Examining it, we see strengths not captured by current definitions or results. We step back to create corresponding abstractions, including a syntax and a hierarchy of definitions of security for non-interactive threshold signature schemes. We then return to ask where, in this hierarchy, we can fit the starting schemes, giving proofs that FROST and its variants as high as possible. In terms of the proofs this needs, and that we give, this turns out to be challenging, so that we offer also some content of technical interest.

Although inspired by specific schemes, our definitional development, once started, unfolds in a logical way, and yields definitions that even go beyond what FROST achieves. These make intriguing new targets. We show how to achieve them, with minimal modifications to the existing schemes.

NON-INTERACTIVE THRESHOLD SCHEMES. We consider schemes where the signing operations involve a leader and a set of n nodes, which we refer to as servers, with server i holding a secret share sk_i of the secret signing key sk . Signing is done via an interactive protocol that begins with a leader request to some set of at least t number of servers and culminates with the leader holding the signature, where $t \leq n$, the threshold, is a protocol parameter.

In a *fully non-interactive* threshold signature scheme, this protocol is a simple, one-round one. The leader sends a leader request lr , which specifies a message M and possibly other things, to any server i , and obtains in response a *partial signature*, $psig_i$, that i computes as a function of sk_i and M . The leader can request partial signatures asynchronously, at any time, and independently for each server, and there is no server-to-server communication. Once it has enough partial signatures, the leader aggregates them into a signature sig of M under the verification key pk corresponding to

sk. The canonical example is the threshold BLS scheme [27, 32], where $sk, sk_1, \dots, sk_n \in \mathbb{Z}_p$ for a public prime p , and $psig_i \leftarrow H(\mu)^{sk_i}$ where $H : \{0, 1\}^* \rightarrow \mathbb{G}$ is a public hash function with range a group \mathbb{G} of order p . Aggregation produces sig as a weighted product of the partial signatures.

A *partially non-interactive* threshold signature scheme adds to the above a message-independent pre-processing round in which, pinged by the leader at any point, a server i returns a pre-processing token pp_i . The leader’s request for a partial signatures will now depend on tokens it has received. The canonical example is FROST [91].

This understanding of a non-interactive scheme encompasses what FROST calls flexibility; obtaining $psig_i$ from *any* $\geq t$ servers will allow us to reconstruct a signature.

WHICH FORGERIES ARE NON-TRIVIAL? For a regular (non-threshold) signature scheme, the first and most basic notion of security is un-forgeability (UF) [78]. The adversary (given access to a signing oracle) outputs a forgery consisting of a message μ and a valid signature for it. To win, the forgery must be *non-trivial*, meaning not legitimately obtained. This is naturally captured, in this context, as meaning that μ was not a signing query.

Turning to define un-forgeability for a non-interactive threshold signature scheme, we assume the adversary has corrupted the leader, and up to $t - 1$ servers, where $1 \leq t \leq n$ is the threshold. Furthermore it has access to the honest servers. Again, it outputs a forgery consisting of a message μ and valid signature for it, and, to win, the forgery must be *non-trivial*, meaning not legitimately obtained. Deciding what “non-trivial” means, however, is now a good deal more delicate, and interesting, than it was for regular signatures.

In this regard, we suggest that many of the prior works have set a low bar, being more generous than necessary in declaring a forgery trivial, leading to definitions that are weaker than one can desire, and weaker even than what their own schemes seem to meet. The definitions we formulate rectify this by giving a hierarchy of five non-triviality conditions of increasing stringency, yielding a corresponding hierarchy $TS\text{-}UF\text{-}0 \leftarrow TS\text{-}UF\text{-}1 \leftarrow TS\text{-}UF\text{-}2 \leftarrow TS\text{-}UF\text{-}3 \leftarrow TS\text{-}UF\text{-}4$ of five notions of un-forgeability of increasing strength. (Here an arrow $B \leftarrow A$ means A implies B : any scheme that is A -secure is also B -secure.) $TS\text{-}UF\text{-}0$, the lowest in the hierarchy, is the notion used in [74, 71, 73, 27]. $TS\text{-}UF\text{-}1$ was introduced by Shoup [119].

Returning to regular (non-threshold) signature schemes, strong un-forgeability (SUF) has the same template as UF, but makes the non-triviality condition more stringent, asking that there have been no signing query μ that returned sig . We ask if SUF has any analogue in the threshold setting. For non-interactive schemes, we suggest it does and give a hierarchy of three definitions

of strong unforgeability $\text{TS-SUF-2} \leftarrow \text{TS-SUF-3} \leftarrow \text{TS-SUF-4}$. The numbering reflects that $\text{TS-UF-}i \leftarrow \text{TS-SUF-}i$ for $i = 2, 3, 4$.

TS-UF-0 AND TS-UF-1. For the basic security notion TS-UF-0, the non-triviality condition is that *no* server was asked to issue a partial signature on the forgery message μ . However, allowing asynchronous requests is a feature of non-interactive schemes. A corrupted leader could ask one honest server i for a partial signature. No other server would even be aware of this request, but the adversary would now have $psig_i$. Under TS-UF-0, the forgery is now trivial, and the adversary does not win. Yet (assuming a threshold $t \geq 2$), there is no reason possession of just $psig_i$ should allow creation of a signature, and indeed for FROST there is no attack that seems able to create such a signature, indicating the scheme is achieving more than TS-UF-0. This leads to the next level of the hierarchy, TS-UF-1, where, following [119], the non-triviality condition is that a partial signature of μ was requested from at most $t - 1 - c$ honest servers, where c is the number of corrupted servers.

The distinction between TS-UF-1 and TS-UF-0 is not just academic. Implicit in applications of threshold signing in wallets is the fact that servers also perform well-formedness checks of what is being signed (typically, as part of a transaction). TS-UF-1 guarantees that every issued signature has been inspected by sufficiently many servers, but TS-UF-0 does not.

BEYOND TS-UF-1. Yet the hierarchy needs to go higher, and this becomes apparent when looking at partially non-interactive schemes like FROST [91]. Here, the discussion becomes more subtle, and interesting.

In more detail, a FROST pre-processing token takes the form of a pair $pp_i = (g^{r_i}, g^{s_i})$ of group elements for one-time use. (A server will ensure that the pre-processing token in its name in the leader request is one it has previously sent, and will never use it again.) An honest request lr includes, along with the message μ to be signed, a sufficiently large server set $lr.SS \subseteq [1..n]$, and, for each i in this set, a pre-processing token pp_i that i previously sent. Each server $i \in lr.SS$ will then generate a signature share $psig_i = (R, z_i)$, where R is a value which can be computed (publicly) from the tokens included in lr , whereas z_i depends on the discrete logarithms of the server's token and its own key share sk_i . The z_i 's can then be aggregated into a value z such that (R, z) is a valid Schnorr signature for μ . Three variants of FROST are known, and differ in the way in which R is computed from lr . The original version, which we refer to as FROST1 [91], requires $|lr.SS|$ exponentiations, whereas a more recent optimization, which we call FROST2 [50], only requires a single exponentiation. A further optimized version, FROST3 [113], improves the communica-

tion complexity by only sending an aggregated nonce. The current security analyses [50, 113] do not surface any security difference between the three variants, but only because they consider a non-triviality notion as in our TS-UF-0 notion.

In terms of our framework, we show that FROST1 achieves TS-SUF-3 security. The proof is under the AOMDL assumption of [105] in the ROM of [18]. This considers a signature trivial even if some of the honest servers in $lr.SS$ do not respond to a (malicious) leader request, as long as the tokens associated with these servers are not honestly generated. In particular, the honest servers may not respond because they recognize these tokens as invalid, or because the malicious leader did not submit the request to them. We show that, while FROST2 fails to achieve TS-SUF-3, it achieves the next step down in our hierarchy, TS-SUF-2. (Again the proof is under AOMDL in the ROM.) This is still stronger than the notions lower in the hierarchy. Finally, we show FROST3 achieves TS-UF-1 under AOMDL in ROM but fails to achieve TS-UF-2.

STRONGER GOALS. A stronger security goal (TS-UF-4 in our hierarchy) is to expect that the *only* way to obtain a signature for a message μ is to follow the above blueprint, i.e., to issue the same honest leader request lr to all servers in $lr.SS$. In fact, we may even ask for more, in terms of *strong* unforgeability — the value R is uniquely defined by lr , and, along with the message μ , it defines a *unique* signature (although not efficiently computable given the verification key alone). An ideal goal, which corresponds to our strongest security goal, is to ensure that the *only* way to generate the signature associated with lr is to obtain a signature share for lr from every honest server whose tokens are included in lr . This is a notion we refer to as TS-SUF-4.

We will however show that neither FROST1 nor FROST2 meet TS-SUF-4. To overcome this, we will show a general transformation which can boost the security of a TS-SUF-3-secure scheme like FROST1 to achieve TS-SUF-4. Our framework allows schemes more general than the FROST ones, and also leaves the question open of better and more efficient designs achieving the stronger notions. Moreover, we provide simple reference schemes for all of our notions, which, while inefficient, guide us in understanding the subtle differences among notions and baseline requirements. In particular, these schemes will enable us to separate the proposed notions.

A SUMMARY FOR OUR NOTIONS. In summary, our unforgeability notions declare a signature for a message μ trivial in the following cases:

- TS-UF-0: A partial signature for the message μ was generated by at least one honest server.
- TS-UF-1: A partial signature for the message μ was generated by at least $t - c$ honest servers,

where c is the number of corrupted servers.

- TS-UF-2: There exists a leader request lr for the message μ which was answered by at least $t - c$ honest servers.
- TS-UF-3: There exists a leader request lr for the message μ such that every honest server $i \in lr.SS$ either answered lr or the token pp_i associated with i in lr is maliciously generated.
- TS-UF-4: There exists a leader request lr for the message μ such that every honest server $i \in lr.SS$ answered lr .

Analogous notions of strong unforgeability are obtained by further associating a request lr to a (unique) signature, in addition to a message μ .

We stress that it is not clear which scenarios demand which notions in our hierarchy. This is especially true because we are still lacking formal analyses of full-fledged systems using threshold signatures, but it is not hard to envision a potential mismatch between natural expectations from such schemes and what they actually achieve. In both FROST variants, for example, it is natural to expect that a signature can only be generated by a sufficient number of honest servers answering the *same* request, a property which we show is actually achieved. Further, one may also expect that all honest servers that generated these honest tokens need to be involved in the generation of a valid signature, but this stronger property is actually not achieved by either of the FROST variants.

WHAT WE DO NOT DO. Some schemes like FROST come with a concrete *distributed key-generation* (DKG) protocol. Security proofs frequently (but not always) consider, monolithically, the composition of DKG and threshold signing. This lack of modular treatment is due to the fact that efficient DKG protocols like Pedersen's [109] are not secure [73] in the strongest possible sense *in isolation*, but it may still be possible to show security when they are used with a particular threshold signature scheme. Here, instead, we idealize DKG protocols, as the points we are trying to express are orthogonal to the concrete choice of a DKG. Our result would still guarantee security of the schemes when used with truly secure DKGs (such as the DKG from [73]), but further investigation is needed to extend our proofs to consider more efficient DKGs.

Our framework does not handle adaptive corruptions, i.e., we demand instead that the adversary declares its corruption set initially. We could extend our definitions to adaptive corruptions rather

easily, but our concrete bounds would be impacted. In particular, we would resort to a generic reduction guessing the corrupted set beforehand, with a multiplicative loss of 2^n , which is acceptable for the smaller values of the number n of parties that we consider common in practice.

Our framework cannot cover recent protocols, like that of Canetti et al. [38], which combine a *multi-round* message-independent pre-processing phase with a final, message-dependent, round. (Conversely, their UC security analysis does not give definitions which help our fine-grained framework.)

Finally, many prior works [73, 113, 23] also consider *robustness*, i.e., the guarantee that a signature is always produced. Here, we follow the same viewpoint as in FROST, and do not focus on robustness explicitly. This allows us to prevent imposing a small t (relative to n) just for the sake of ensuring it. However, our schemes all implicitly give verification keys pk_i for each server, and it is not hard to verify individual partial signatures $psig_i$. Any t valid partial signatures will always aggregate into a valid signature.

7.2 Partially Non-Interactive Threshold Signatures from Weaker Assumptions

As our second main contribution, we develop new two-round protocols that are secure under the discrete logarithm assumption in the ROM. The signatures produced by both schemes resemble those proposed by Okamoto [107]. Furthermore, our signing protocols are partially non-interactive, i.e., the first round messages do not depend on the message being signed, which is a desirable property in practice.

SIGNIFICANCE. Our DL-based schemes are the first partially non-interactive 2-round schemes based solely on the hardness of the DL assumption. For threshold signatures, in particular, no two-round scheme is known from only the DL assumption.

OUR APPROACH. Our schemes are the outcome of the same paradigm applied to FROST [91, 14] and its variants. It is not known how to prove the security of either scheme under the plain discrete logarithm assumption, and they are instead proved secure under the (stronger) *AOMDL* [105] assumption, which is a variant of the *OMDL* assumption [16], an assumption that has been the subject of criticism [90, 89]. As we explain next, our paradigm can be seen as a general recipe to remove the *AOMDL* assumption from these schemes.

The main ingredient of our approach are *linear hash functions*, which have also been used in recent works [10, 80, 81] to abstract identification schemes from which signature variants are de-

rived. Here, we observe that FROST can naturally be generalized by replacing the exponentiation map $x \mapsto g^x$ with a linear hash function $F : \mathcal{D} \rightarrow \mathcal{R}$, where \mathcal{D}, \mathcal{R} are \mathcal{S} -modules for a field \mathcal{S} . We generically refer to these instantiations as FROST-H. (In fact, we present two variants for FROST-H but make no distinction in the introduction.) In particular, we require that:

- F is an epimorphism of \mathcal{S} -modules from \mathcal{D} to \mathcal{R} , i.e., F is a surjection from \mathcal{D} to \mathcal{R} such that for any $r \in \mathcal{S}$ and $x, y \in \mathcal{D}$, $F(x + r \cdot y) = F(x) + r \cdot F(y)$.
- F is *not a monomorphism*, which is equivalent to postulating that there exists $z^* \in \mathcal{D}$ such that $z^* \neq 0$ and $F(z^*) = 0$.

We then define a natural analogue of the AOMDL assumption, which we refer to as the *Algebraic One-More Preimage Resistance* (AOMPR). Roughly speaking, the corresponding security game allows the attacker to obtain multiple *challenges* $X_i = F(x_i)$ for a random element $x_i \leftarrow \mathcal{D}$, and the attacker also gets access to an *inversion oracle* which, on input $X \in \mathcal{R}$, returns a element in the preimage set of X under F . The restriction here, and hence the term *algebraic*, is that X must be an affine combination of previously obtained X_i 's, and this affine combination is given to the inversion oracle, along with X . (This makes the assumption falsifiable since the oracle can efficiently answer such inversion queries.) To win the game, the attacker is then asked to invert $q + 1$ challenges after querying the inversion oracle at most q times.

Our results then follow from the combination of the following two theorems, which we state here informally:

Theorem (informal). The security of FROST-H follows from the AOMPR assumption on the underlying linear hash function.

Theorem (informal). If F is collision-resistant, then the AOMPR assumption holds with respect to F .

The proof of the first theorem is, on its own, not particularly surprising and mostly generalizes the prior proofs in the literature, in particular those of [105] and [14]. Our main contribution here is to notice that these proofs, and the resulting schemes, can be abstracted in terms of linear hash functions. In particular, for threshold signatures, as in [14], we consider an abstract setting with an ideal distributed key generation, and we target the security notions of TS-SUF-2 and TS-SUF-3, which

were shown to be achieved by two variants of FROST, both of which we model here abstractly. Since we are targeting feasibility, we are less concerned with the concrete round complexity of distributed key generation and could use any secure multi-party computation protocol for this task.

In contrast, the rough intuition behind a proof of the latter theorem is that for any execution of a (wlog deterministic) adversary \mathcal{A} playing the AOMPR game with challenges $\mathbf{X} = F(\mathbf{x})$, since F is not a monomorphism, there exists another execution with challenges $\mathbf{X} = F(\mathbf{x}')$ such that $\mathbf{x} \neq \mathbf{x}'$, but the views of \mathcal{A} are identical in the two executions. Then, if \mathcal{A} wins the game given \mathbf{x} by outputting \mathbf{y} such that $F(\mathbf{y}) = \mathbf{X}$, \mathcal{A} also wins the game given \mathbf{x}' by outputting \mathbf{y} . Therefore, we have $F(\mathbf{x}) = F(\mathbf{y}) = F(\mathbf{x}') \wedge (\mathbf{x} \neq \mathbf{y} \vee \mathbf{x}' \neq \mathbf{y})$, which implies that we can find a collision in at least one of the executions. Indeed, special cases of this technique already underlie several works, including Okamoto's [107], but our main challenges are to prove the concrete mapping of \mathbf{x}' from \mathbf{x} and to package this in terms of the AOMPR abstraction.

7.2.1 DL-based Instantiations

To obtain an instantiation of FROST-H based on the hardness of the discrete logarithm (DL) problem, we can use the Pedersen linear hash function [109]

$$F(x_1, x_2) = g^{x_1} Z^{x_2} ,$$

which is well known to be collision-resistant under the hardness of DL whenever g, Z are generators of a group with prime size p . While FROST produces valid Schnorr signatures [114], the signatures produced by our DL-based instantiations of FROST-H are slightly less efficient, and effectively compatible with Okamoto's signatures [107]. Here, as in Schnorr signatures, the secret signing key is $x \in \mathbb{Z}_p$, and the public verification key $\text{pk} = g^x$, and a signature for a message $m \in \{0, 1\}^*$ has format

$$\sigma = (R = g^a Z^b, a + H(\text{pk}, m, R) \cdot x, b) ,$$

where H is a hash function that is modeled as a random oracle in our proofs. To verify a signature (R, a, b) , we check that $g^a Z^b = R \cdot \text{pk}^{H(\text{pk}, M, R)}$. The only difference from Okamoto's scheme [107] is that the latter uses a secret key $(x_1, x_2) \in \mathbb{Z}_p^2$, and a signature has form $(R = g^a Z^b, a + c \cdot x_1, b + c \cdot x_2)$, where $c = H(\text{pk}, m, R)$, i.e., here, we restrict the scheme to the case where $(x_1, x_2) = (x, 0)$. This optimization is generic and could have been applied to Okamoto's scheme directly; however, it is particularly advantageous for threshold signatures since it lets us leverage any distributed

key generation protocol for Schnorr signatures. Here, we need a trusted setup to generate Z as a random group element independent of g , but we note that this is a minimal setup since it can be made transparent, e.g., g, Z can be generated as outputs of a hash function.

RELATED WORK. Our DL-based threshold signatures are the first two-round scheme with security proved based solely on the discrete logarithm assumption in the ROM. The most efficient protocol is FROST [91, 14], which is slightly more efficient than our scheme since it generates plain Schnorr signatures; however, FROST relies on the stronger AOMDL assumption. Though schemes based solely on the discrete logarithm assumption exist [120, 73, 95], they use more rounds. We stress that not all schemes achieve the same security goals, and here we target our security notions, whereas Lindell [95] targets UC security.

Chapter 8

A FRAMEWORK FOR NON-INTERACTIVE THRESHOLD SIGNATURES

We present our hierarchy of definitions of security for non-interactive threshold schemes, formalizing both unforgeability (UF) and strong unforgeability (SUF) in several ways. We provide relations between all notions considered.

8.1 Syntax and Correctness

MAINTAINING STATE. Parties as implemented in protocols would maintain state. When activated with some inputs (which include messages from other parties), they would apply some algorithm Alg to these and their current state to get outputs (including outgoing messages) and an updated state. To model this, we do not change our definition of algorithms, but make the state an explicit input and output that will, in definitions, be maintained by the overlying game. Thus, we would write something like $(\dots, st) \leftarrow \$ \text{Alg}(\dots, st)$.

SYNTAX. A non-interactive threshold signature scheme TS specifies a number $n \geq 1$ of signers, a reconstruction threshold t , a set HF of functions from which the random oracle is drawn, a key-generation algorithm KeyGen, a server pre-processing algorithm SPP, a leader pre-processing algorithm LPP, a leader signing-request algorithm LR, a server partial-signature algorithm PS, a leader partial-signature aggregation algorithm Agg and a verification algorithm Vf. If disambiguation is needed, we write $\text{TS}.n$, $\text{TS}.t$, $\text{TS}.HF$, $\text{TS}.Setup$, $\text{TS}.KeyGen$, $\text{TS}.SPP$, $\text{TS}.LPP$, $\text{TS}.LR$, $\text{TS}.PS$, $\text{TS}.Agg$, $\text{TS}.Vf$, respectively. We now explain the operation and use of these components, the understanding of which may be aided by already looking at the correctness game $\text{TS-COR}_{\text{TS}}$ of Figure 8.1.

Parties involved are a leader (numbered 0, implicit in some prior works, but made explicit here) and signers numbered $1, \dots, n$, for a total of $n + 1$ parties. The setup algorithm $Setup(1^\kappa)$ initializes the state st_i for each signer $i \in [n]$, as well as the state st_0 for the leader, and additionally a public parameter par , which is given as input to all other algorithms. Algorithms have oracle

<p><u>Game TS-COR_{TS}(κ, μ, SS) :</u></p> <p>$par \leftarrow \text{Setup}(1^\kappa) ; H \leftarrow_{\\$} \text{HF}$</p> <p>$(pk, aux, \{sk_i\}_{i \in [n]}) \leftarrow \text{KeyGen}()$</p> <p>For $i \in [n]$ do $st_i.sk \leftarrow sk_i ; st_i.pk \leftarrow pk$</p> <p>For $i \in SS$ do</p> <p style="padding-left: 2em;">$(pp_i, st_i) \leftarrow \text{SPP}(st_i) ; st_0 \leftarrow \text{LPP}(i, pp_i, st_0)$</p> <p>$(lr, st_0) \leftarrow \text{LR}(\mu, SS, st_0)$</p> <p>For $i \in SS$ do $(psig_i, st_i) \leftarrow \text{PS}(lr, i, st_i)$</p> <p>$sig \leftarrow \text{Agg}(lr, \{psig_i\}_{i \in SS}, st_0)$</p> <p>Return $\text{Vf}(pk, \mu, sig) = 0$</p>

Figure 8.1: The TS-COR game for a threshold signature scheme TS for n signers and threshold t .

access to a function H that is drawn at random from HF in games and plays the role of the random oracle. Specifying HF as part of the scheme allows the domain and range of the random oracle to be scheme-dependent. We omit par and H from the input of all algorithms for simplicity.

The key-generation algorithm KeyGen , run once at the beginning, creates a public signature-verification key pk , associated public auxiliary information aux , and an individual secret signing key sk_i for each server $i \in [1..n]$. (Usually, sk_1, \dots, sk_n will be shares of a global secret key sk , but the definitions do not need to make sk explicit. The leader does not hold any secrets associated to pk .) While key-generation may in practice be performed by a distributed key-generation protocol, our syntax assumes it done by a trusted algorithm to allow a modular treatment. Keys are held by parties in their state, encoded into dedicated fields of the latter as shown in Figure 8.1. For specific schemes, we will typically use aux to model additional information that can be leaked by key generation step without violating security (e.g., the values g^{sk_i} in most cases).

The signing protocol can be seen as having two rounds, which we think as a pre-processing and online stage. In a pre-processing round, any server i can run $(pp, st_i) \leftarrow_{\$} \text{SPP}(st_i)$ to get a *pre-processing token* pp which it sends to the leader. (Here st_i is the state of i .) Via $st_0 \leftarrow \text{LPP}(pp, st_0)$, the leader updates its state st_0 to incorporate token pp .

In a signing round, the leader begins with a message and a choice of a signer set $SS \subseteq [1..n]$

of size at least t . Via $(lr, st_0) \leftarrow^s \text{LR}(M, SS, st_0)$ it generates a leader request lr that, through st_0 , implicitly depends on a choice of pre-processing tokens. The leader request is sent to each $i \in SS$, who, via $(psig_i, st_i) \leftarrow^s \text{PS}(lr, i, st_i)$, computes a partial signature $psig_i$ and returns it to the leader. Via $(sig, st_0) \leftarrow^s \text{Agg}(lr, \{psig_i\}_{i \in SS}, st_0)$, the leader aggregates the partial signatures into a signature sig of M , the desired output of the protocol. Also, we require that for $i \notin SS$, $\text{PS}(lr, i, st_i)$ always outputs (\perp, st_i) .

The verification algorithm, like in a standard signature scheme, takes pk , a message μ and a candidate signature, and returns a boolean validity decision.

ECHO SCHEMES. We define a sub-class of non-interactive threshold schemes that we call *echo schemes*. Recall that a leader request lr is mandated to specify a message $lr.\text{msg}$ and a set $lr.SS \subseteq [1..n]$ of signers from whom partial signatures are being requested. In an echo scheme, lr additionally specifies a function $lr.PP : lr.SS \rightarrow \{0, 1\}^*$. If the leader is honest, $lr.PP(i)$ is a token pp that i had previously sent to the leader. That is, the leader is echoing tokens back to the signers, whence the name. In considering security, of course, $lr.PP(i)$ is picked by the adversary and may not be a prior token. For an echo scheme, we impose the additional requirement that the state st_i of party i includes a set $st_i.PP$ in which SPP stores prior tokens: more formally, the execution $(pp, st_i) \leftarrow \text{SPP}(st_i)$ is required to include the update $st_i.PP \leftarrow st_i.PP \cup \{pp\}$. In this way, st_i is recording the tokens that party i has previously sent to the leader. Finally, we require that $\text{PS}(lr, i, st_i)$ always outputs (\perp, st_i) if $lr.PP(i) \notin st_i.PP$. This is a natural requirement since an honest signer should not answer a leader request that is clearly malicious. As we will discuss in Section 9.1, FROST1 and FROST2 are examples of an echo scheme.

CORRECTNESS OF A TS SCHEME. We say that TS is *correct* with correctness error ε_{cor} if for $\mu \in \{0, 1\}^*$ and $SS \subseteq [n]$ with $|SS| \geq t$, we have $\Pr[\text{TS-COR}_{\text{TS}}(\kappa, \mu, SS) = 1] \leq \varepsilon_{\text{cor}}$. We say that TS is *perfectly correct* if $\varepsilon_{\text{cor}} = 0$. By default, when we say that TS is correct, we mean it is perfectly correct.

The way in which we are supposed to interpret the correctness definition is that a request lr is associated with a set SS and a message μ , and if such a request is issued successfully by the leader (i.e., $lr \neq \perp$), then the signers in SS would all accept lr producing partial signatures which aggregate into a valid signature for μ .

<p><u>Game TS-UF-$i_{\text{TS}}^A(\kappa)$ ($i = 0, 1, 2, 3, 4$) and TS-SUF-$i_{\text{TS}}^A(\kappa)$ ($i = 2, 3, 4$) :</u></p> <p>$par \leftarrow \text{Setup}(1^\kappa)$; $H \leftarrow_s \text{TS.HF}$; $\text{TF}_M \leftarrow \emptyset$; $\text{TF}_L \leftarrow \emptyset$; $\text{curSS}_M, \text{curSS}_L \leftarrow ()$</p> <p>$(\mu, sig) \leftarrow \mathcal{A}^{\text{INIT}, \text{PPO}, \text{PSIGNO}, \text{RO}}(par)$</p> <p>If $\forall f(\text{pk}, \mu, sig) = 0$ then return 0</p> <p>For each lr do</p> <p style="padding-left: 20px;">If $\text{tf}_i(lr.\text{msg})$ then $\text{TF}_M \leftarrow \text{TF}_M \cup \{lr.\text{msg}\}$ // For TS-UF-i_{TS}^A ($i = 0, 1$)</p> <p style="padding-left: 20px;">If $\text{tf}_i(lr)$ then $\text{TF}_M \leftarrow \text{TF}_M \cup \{lr.\text{msg}\}$ // For TS-UF-i_{TS}^A ($i = 2, 3, 4$)</p> <p style="padding-left: 20px;">If $\text{tf}_i(lr)$ then $\text{TF}_L \leftarrow \text{TF}_L \cup \{lr\}$ // For TS-SUF-i_{TS}^A</p> <p>Return $\mu \notin \text{TF}_M$ // For TS-UF-i_{TS}^A</p> <p>Return $(\nexists lr \in \text{TF}_L : lr.\text{msg} = \mu \wedge \text{SVf}(\text{pk}, lr, sig) = 1)$ // For TS-SUF-i_{TS}^A</p>	
<p><u>Oracle INIT(CS) :</u></p> <p>Require: $CS \subseteq [n]$ and $CS < t$</p> <p>$HS \leftarrow [n] \setminus CS$; $(\text{pk}, \text{aux}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{KeyGen}()$</p> <p>For $i \in HS$ do $\text{st}_i.\text{sk} \leftarrow \text{sk}_i$; $\text{st}_i.\text{pk} \leftarrow \text{pk}$</p> <p>Return $(\text{pk}, \text{aux}, \{\text{sk}_i\}_{i \in CS})$</p>	<p><u>Oracle PPO(i) :</u></p> <p>Require: $i \in HS$</p> <p>$(pp, \text{st}_i) \leftarrow_s \text{SPP}(\text{st}_i)$</p> <p>$\text{PP}_i \leftarrow \text{PP}_i \cup \{pp\}$</p> <p>Return pp</p>
<p><u>Oracle PSIGNO(i, lr) :</u></p> <p>Require: $lr.\text{SS} \subseteq [n]$ and $i \in HS \cap lr.\text{SS}$</p> <p>Require: $lr.\text{PP}(i) \in \text{PP}_i$ // Only for the case where TS is an echo scheme</p> <p>$(psig, \text{st}'_i) \leftarrow_s \text{PS}(lr, i, \text{st}_i)$; If $psig = \perp$ then return \perp</p> <p>$\text{curSS}_M(lr.\text{msg}) \leftarrow \text{curSS}_M(lr.\text{msg}) \cup \{i\}$; $\text{curSS}_L(lr) \leftarrow \text{curSS}_M(lr) \cup \{i\}$</p> <p>Return $psig$</p>	
<p><u>Oracle RO(x) :</u></p> <p>Return $H(x)$</p>	

Figure 8.2: The games TS-UF- i and TS-SUF- i of a threshold signature scheme TS. In the game TS-UF- i (resp. TS-SUF- i), the set TF_M (resp. TF_L) records the set of messages (resp. leader requests) that are considered trivial under the trivial forgery predicate tf_i (Figure 8.3). The oracle INIT must be invoked exactly once, and only before any other oracle is queried. In particular, the predicate tf_3 and, thus, TS-UF-3 and TS-SUF-3 are defined only if TS is an echo scheme.

8.2 Unforgeability and Strong Unforgeability

UNFORGEABILITY. Unforgeability as usual asks that the adversary be unable to produce a valid signature sig on some message M of its choice except in a trivial way. The question is what “trivial” means. For regular signatures, it means that the adversary did not obtain a signature of M from the signing oracle [78]. For threshold signatures, it is more subtle. We will give several definitions.

Figure 8.2 simultaneously describes several games, $\text{TS-UF-}i_{\text{TS}}$ for $i = 0, 1, 2, 3, 4$, where $\text{TS-UF-}3_{\text{TS}}$ is only defined if TS is an echo scheme. (We will get to the second set of games later.) They are almost the same, differing only in the winning conditions. The corresponding advantages of an adversary \mathcal{A} are $\text{Adv}_{\text{TS}}^{\text{ts-uf-}i}(\mathcal{A}, \kappa) = \Pr[\text{TS-UF-}i_{\text{TS}}^{\mathcal{A}}(\kappa)]$. The adversary calls INIT with a choice of a set of signers to corrupt. The oracle INIT must be invoked exactly once, and only before any other oracle is queried. It is also viewed as having corrupted the leader. Playing the leader role, it can request pre-processing tokens via oracle PPO . It can provide a server with a leader-request lr of its choice to obtain a partial signature PSIGNO . At the end, it outputs its forgery message M and signature sig . The adversary does not win if the signature is not valid. Now, to win, the signature must be non-trivial. It is in how this is defined that the games differ. Associated to i is a *trivial forgery* predicate tf_i that is invoked in the PSIGNO oracle to decide whether a message is added to the trivial forgery set TF_M . The choices for these predicates are shown in the table in Figure 8.3. When $i = 0$ we have the usual notion from the literature, used in particular in [27, 71, 73]. As i increases, we get more stringent (less generous) in declaring a forgery trivial, and the notion gets stronger.

Concretely, $\text{TS-UF-}0$ considers a signature for a message M trivial if a request lr with $lr.\text{msg}$ was answered by server with a partial signature. Moving on, $\text{TS-UF-}1$ strengthens this by declaring a signature trivial only if at least $t - |CS|$ signers have responded to some request for message M , where these requests could have been different. In turn, $\text{TS-UF-}2$ strengthens this even further by requiring that there was a single prior request lr for M which was answered by $t - |CS|$ signers.

The notions $\text{TS-UF-}3$ only deals with echo schemes. Recall that for these schemes, a request lr contains a map $lr.\text{PP} : lr.\text{SS} \rightarrow \{0, 1\}^*$, where $lr.\text{PP}(i)$ is meant to be a token issued by server i . Here, we consider a signature for message M trivial if there exists a request lr for M which is answered by all honest signers i for which $lr.\text{PP}(i)$ is a valid token previously output by i , and this set consists of at least $t - |CS|$ signers. Finally, our strongest notion, $\text{TS-UF-}4$ simply considers

$\mathbf{tf}_0(\mu)$: $\text{curSS}_M(\mu) \neq \emptyset$
$\mathbf{tf}_1(\mu)$: $ \text{curSS}_M(\mu) \geq t - CS $
$\mathbf{tf}_2(lr)$: $ \text{curSS}_L(lr) \geq t - CS $
$\mathbf{tf}_3(lr)$: $\mathbf{tf}_2(lr)$ and $\text{curSS}_L(lr) = \{i \in HS \cap lr.SS : lr.PP(i) \in PP_i\}$
$\mathbf{tf}_4(lr)$: $\mathbf{tf}_2(lr)$ and $\text{curSS}_L(lr) = HS \cap lr.SS$




Figure 8.3: **Top:** Trivial-forgery conditions $\mathbf{tf}_i(lr)$ ($i = 0, 1, 2, 3, 4$) used to define games TS-UF- i and TS-SUF- i (Figure 8.2). **Bottom:** Relations between notions of security.

a signature trivial if there exists a request lr for M which is answered by all honest signers in $i \in lr.SS$.

It is natural to expect TS-UF-3 and TS-UF-4 to be similar, but as we will see below, they are actually not equivalent. (Although we will give a transformation that boosts an TS-UF-3-secure scheme into an TS-UF-4-secure one.)

STRONG UNFORGEABILITY. For standard signatures, strong unforgeability asks, in addition to unforgeability, that the adversary be unable to produce a new signature on any message, where new means different from any obtained legitimately for that message. We ask, does this have any counterpart in threshold signatures? In fact, FROST seems to have such a property. We now provide formalisms to capture such properties.

It turns out that giving a general definition of strong unforgeability is rather complex, and we will restrict ourselves to a natural subclass of schemes (which includes FROST). Concretely, we ask that there is an algorithm SVf , called a *strong verification algorithm*, that takes a public key pk , a leader request lr , and a signature sig as inputs and outputs 1 or 0. We require that for any

pk, lr there exists at most one signature sig such that $\text{SVf}(\text{pk}, lr, sig) = 1$. Also, if there exists sig such that $\text{SVf}(\text{pk}, lr, sig) = 1$, it must hold that $\text{Vf}(\text{pk}, lr.\text{msg}, sig) = 1$, and we say that lr is *intended to produce* the signature sig for the message $lr.\text{msg}$. Moreover, TS is asked to satisfy a strong correctness property which is defined using the same game as $\text{TS-COR}_{\text{TS}}$ except the winning condition $\text{Vf}(\text{pk}, M, sig) = 0$ is replaced with $\text{SVf}(\text{pk}, lr, sig) = 0$. As we will discuss in Section 9.1, FROST schemes exactly have this property and we can define SVf easily.

For a scheme TS with a strong verification algorithm, we consider the $\text{TS-SUF-}i_{\text{TS}}$ ($i = 2, 3, 4$) games in Figure 8.2, where $\text{TS-SUF-}3_{\text{TS}}$ is only defined if TS additionally is an echo scheme. The differences (across the different values of i) are only in the trivial forgery predicates tf_i (defined in Figure 8.3). The difference between $\text{TS-SUF-}i$ and $\text{TS-SUF-}i$ is that we record the set of trivial leader requests TF_L instead of the set of trivial messages TF_M . More precisely, in $\text{TS-SUF-}i$, we record all leader requests lr that satisfy $\text{tf}_i(lr)$ in the set TF_L . A forgery (μ, sig) is considered non-trivial only if no leader request in TF_L is intended to produce the signature sig for μ . The corresponding advantage of an adversary \mathcal{A} is $\text{Adv}_{\text{TS}}^{\text{ts-suf-}i}(\mathcal{A}, \kappa) = \Pr[\text{TS-SUF-}i_{\text{TS}}(\mathcal{A})]$. The ensuing notion is called $\text{TS-SUF-}i$.

8.3 Relations and Transformations

RELATIONS BETWEEN NOTIONS. Figure 8.3 shows relations between the notions of unforgeability and strong unforgeability that we have defined. A (blue, non-dotted) arrow $A \rightarrow B$ is an implication, saying that A implies B: any scheme that is A-secure is also B-secure. Now see the nodes as forming a graph with edges the blue, non-dotted arrows. The thin arrow from $\text{TS-UF-}0$ to $\text{TS-UF-}1$ indicates us that the implication only holds under a quantitatively loose reduction. (We prove this in Theorem 8.3.2.) We claim that in this graph, if there is no path from a notion B to a notion A, they are separate or distinct: there exists a scheme that is B-secure but not A-secure. The dotted arrows are separations that we explicitly prove. These, together with the full arrows, prove the claim just made. The thick dotted arrows indicate the existence of a generic transformation lifting security of a scheme to achieve a stronger notion. (We establish this below as part of Theorem 8.3.3.)

8.3.1 Reference schemes and proofs of relations.

In this section, we give a set of (fully) non-interactive threshold schemes that we call reference schemes. They represent simple, canonical ways to achieve the different notions. They may not

be of practical interest, because they have key and signature sizes proportional to n , but the point is to embody notions in a representative way. A few things emanate from these schemes. One is that we use them, in the same Appendix, to establish the separations given by the dotted lines in Figure 8.3, thereby showing that any notions between which there is no path, in the graph given by the full arrows, are indeed separate. Second, we get a scheme that achieves our strongest notion, TS-SUF-4, which neither FROST schemes achieve. (Although we can get such a scheme by applying our transformation from Theorem 8.3.3 to FROST1.) Finally, reference schemes, as canonical examples, are ways to understand the notions.

PROOFS OF IMPLICATION RELATIONS. To show an unforgeability notion A implies B , since the unforgeability games defined in Figure 8.2 only differ on the winning conditions, we only need to show if the winning condition of the game of A does not hold when the adversary outputs (M, sig) , the winning condition of the game of B does not hold either.

- TS-UF-1 \Rightarrow TS-UF-0: Since $|CS| < t$, $|\text{curSS}_M(M)| \geq t - |CS|$ implies $\text{curSS}_M(M) \neq \emptyset$.
- TS-UF-2 \Rightarrow TS-UF-1: If there exists lr such that $lr.\text{msg} = M$ and $|\text{curSS}_L(lr)| \geq t - |CS|$, from the execution of the games, $|\text{curSS}_M(M)| \geq |\text{curSS}_L(lr)| \geq t - |CS|$.
- TS-(S)UF-3 \Rightarrow TS-(S)UF-2: It follows from the trivial fact that $\text{tf}_3(lr)$ is stronger than $\text{tf}_2(lr)$.
- TS-(S)UF-4 \Rightarrow TS-(S)UF-3: Due to the requirement that $lr.\text{PP}(i) \in \text{PP}_i$ in the oracle PSIGNO , $\text{curSS}_L(lr) \subseteq \{i \in HS \cap lr.\text{SS} : lr.\text{PP}(i) \in \text{PP}_i\} \subseteq HS \cap lr.\text{SS}$ for any lr . Therefore, $\text{curSS}_L(lr) = HS \cap lr.\text{SS}$ implies $\text{curSS}_L(lr) = \{i \in HS \cap lr.\text{SS} : lr.\text{PP}(i) \in \text{PP}_i\}$, which implies $\text{tf}_4(lr)$ is stronger than $\text{tf}_3(lr)$.
- TS-SUF- i \Rightarrow TS-UF- i , for $i \in \{2, 3, 4\}$: If the adversary outputs a valid forgery (μ, sig) but does not win TS-SUF- i , there exists lr such that $\text{tf}_i(lr) = 1$, $lr.\text{msg} = \mu$ and $\text{SVf}(\text{pk}, lr, sig) = 1$. Therefore, $lr.\text{msg}$ is added to TF_M , which means the adversary does not win TS-UF- i either.

REFERENCE SCHEMES. Fix t, n such that $2 \leq t < n$. The reference schemes are shown in Figure 8.4.

```

Algorithm  $\text{RTS}_\ell[\text{DS}].\text{KeyGen}$  :
For  $i = 1, \dots, n$  do  $(\text{pk}_i, \text{sk}_i) \leftarrow_{\$} \text{DS}.\text{KeyGen}$ 
 $\text{pk} \leftarrow (\text{pk}_1, \dots, \text{pk}_n)$  ; Return  $(\text{pk}, \perp, \text{sk}_1, \dots, \text{sk}_n)$ 
Algorithm  $\text{RTS}_\ell[\text{DS}].\text{SPP}(\text{st})$  :
Return  $(\perp, \text{st})$ 
Algorithm  $\text{RTS}_\ell[\text{DS}].\text{LPP}(pp, \text{st}_0)$  :
Return  $\text{st}_0$ 
Algorithm  $\text{RTS}_\ell[\text{DS}].\text{LR}(M, SS, \text{st}_0)$  :
 $lr.\text{msg} \leftarrow M$  ;  $lr.SS \leftarrow SS$ 
For  $i \in lr.SS$  do  $lr.PP(i) \leftarrow \perp$  //  $\ell = 2, 3, 4$ 
Return  $(lr, \text{st}_0)$ 
Algorithm  $\text{RTS}_\ell[\text{DS}].\text{PS}(lr, \text{st})$  :
If  $(\text{st}.\text{me} \notin lr.SS$  or  $|lr.SS| < t$ ) then return  $(\perp, \text{st})$ 
If  $(lr.PP(\text{st}.\text{me}) \neq \perp)$  then return  $(\perp, \text{st})$  //  $\ell = 3, 4$ 
 $psig \leftarrow_{\$} \text{DS}.\text{Sign}(\text{st}.\text{sk}, lr.\text{msg})$  //  $\ell = 1$ 
 $psig \leftarrow_{\$} \text{DS}.\text{Sign}(\text{st}.\text{sk}, lr)$  //  $\ell = 2, 3, 4$ 
Return  $(psig, \text{st})$ 
Algorithm  $\text{RTS}_\ell[\text{DS}].\text{Agg}(lr, \{psig_i\}_{i \in lr.SS}, \text{st}_0)$  :
 $sig \leftarrow (lr, lr.SS, \{psig_i\}_{i \in lr.SS})$  ; Return  $(sig, \text{st}_0)$ 

```

Figure 8.4: Reference threshold signature schemes $\text{RTS}_\ell[\text{DS}]$ associated to signature scheme DS for $\ell = 1, 2, 3, 4$. The verification and stronger verification algorithms are provided in Figure 8.5.

Proposition 8.3.1. *Suppose DS is a signature scheme and $2 \leq t < n$. Let $\text{RTS}_\ell[\text{DS}]$ ($\ell = 1, 2, 3, 4$) be the reference threshold schemes defined in Figure 8.4. Then: (1) For all $\ell \in \{1, 2, 3, 4\}$, if DS is UF-CMA-secure then $\text{RTS}_\ell[\text{DS}]$ is TS-UF- ℓ -secure, and (2) For all $\ell \in \{2, 3, 4\}$, if DS is unique and SUF-CMA-secure then $\text{RTS}_\ell[\text{DS}]$ is TS-SUF- ℓ -secure.*

The proof of the above proposition is rather straightforward. We now use these schemes to prove the separations claimed by the dotted arrows in Figure 8.3.

TS-UF-1 $\not\Rightarrow$ TS-UF-2. We need to exhibit a scheme TS that is TS-UF-1-secure but not TS-UF-2-secure. Let $\text{TS} = \text{RTS}_1[\text{DS}]$ (Figure 8.4) where DS is a UF-CMA-secure standard signature

```

Algorithm  $\text{RTS}_\ell[\text{DS}].\text{Vf}(\text{pk}, M, \text{sig}) :$ 
 $(\text{pk}_1, \dots, \text{pk}_n) \leftarrow \text{pk} ; (lr, F, \{\text{psig}_i\}_{i \in F}) \leftarrow \text{sig}$ 
If  $(lr.\text{msg} \neq M$  or  $F \not\subseteq lr.\text{SS})$  then return 0
 $T \leftarrow \{i \in F : \text{DS.Vf}(\text{pk}_i, M, \text{psig}_i)\}$  //  $\ell = 1$ 
 $T \leftarrow \{i \in F : \text{DS.Vf}(\text{pk}_i, lr, \text{psig}_i)\}$  //  $\ell = 2, 3, 4$ 
Return  $(T = F$  and  $|T| \geq t)$  //  $\ell = 1, 2$ 
 $E \leftarrow \{i \in lr.\text{SS} : lr.\text{PP}(i) = \perp\}$  ; Return  $(T = E = F$  and  $|T| \geq t)$  //  $\ell = 3$ 
Return  $(T = lr.\text{SS} = F$  and  $|T| \geq t)$  //  $\ell = 4$ 

Algorithm  $\text{RTS}_\ell[\text{DS}].\text{SVf}(\text{pk}, lr, \text{sig}) :$ 
// Only for  $\ell = 2, 3, 4$ 
 $(lr', F, \{\text{psig}_i\}_{i \in F}) \leftarrow \text{sig}$ 
Return  $(\text{RTS}_\ell[\text{DS}].\text{Vf}(\text{pk}, lr.\text{msg}, \text{sig})$  and  $lr = lr')$ 

```

Figure 8.5: The verification and strong verification algorithms of the reference threshold signature schemes $\text{RTS}_\ell[\text{DS}]$ associated to signature scheme DS for $\ell = 1, 2, 3, 4$.

scheme. Proposition 8.3.1 says this achieves TS-UF-1. We now give an attack showing it fails TS-UF-2. The idea is that the adversary can make partial-signing requests to t servers, all with the same message but with $lr.\text{SS}$, and thus lr itself, varying across the requests, so that $\text{curSS}_L(lr)$ stays small for any particular lr , and non-triviality under tf_2 is maintained. Proceeding to the details of the attack, fix a message M , and consider the following adversary \mathcal{A} for game TS-UF-2_{TS}:

Adversary \mathcal{A}

1. $CS \leftarrow \emptyset ; (\text{pk}, \emptyset, \emptyset) \leftarrow \text{INIT}(CS)$
2. $lr_1.\text{msg} \leftarrow \mu ; lr_1.\text{SS} \leftarrow [1..t] ; lr_2.\text{msg} \leftarrow M ; lr_2.\text{SS} \leftarrow [2..t + 1]$
3. For $i \in [1..t - 1]$ do $\text{psig}_i \leftarrow \text{PSIGNO}(i, lr_1)$
4. $\text{psig}_t \leftarrow \text{PSIGNO}(t, lr_2)$
5. $\text{sig} \leftarrow (lr_1, [1..t], \{\text{psig}_i\}_{i \in [1..t]})$; Return (μ, sig)

We claim that $\text{Adv}_{\text{TS}}^{\text{ts-uf-2}}(\mathcal{A}, \kappa) = 1$. The adversary has gathered t valid signatures, so the condition at line 18 of Figure 8.4 is met, and $\text{TS.Vf}(\text{pk}, \mu, \text{sig}) = 1$. Now we need to show that the forgery is non-trivial, meaning $\text{tf}_2(lr_1) = \text{tf}_2(lr_2) = 0$. Indeed, we have $|\text{curSS}_L(lr_1)| = t - 1 < t - |CS| = t$ and $|\text{curSS}_L(lr_2)| = 1 < t - |CS| = t$.

TS-SUF-2 $\not\Rightarrow$ TS-UF-3. We need to exhibit a scheme TS that is TS-SUF-2-secure but not TS-UF-3-secure. Let $TS = RTS_2[DS]$ where DS is a unique SUF-CMA-secure standard signature scheme. Proposition 8.3.1 says this achieves TS-SUF-2. We now give an attack showing it fails TS-UF-3. The idea is to allow the adversary to corrupt $t - 1$ signers, while setting $lr.SS$ to include all corrupted signers along with at least two honest signers. In this case, the adversary can produce a valid signature for lr by querying only a single honest signer in $lr.SS$. However, $curSS_L(lr)$ does not match the set $\{i \in HS \cap lr.SS : lr.PP(i) \in PP_i\}$, which contains at least two honest signers. Therefore, the adversary wins TS-UF-3. Proceeding to the details of the attack, fix a message μ , and consider the following adversary \mathcal{A} for game $TS-UF-3_{TS}$:

Adversary \mathcal{A}

1. $CS \leftarrow [t - 1]$; $(pk, aux, \{sk_i\}_{i \in [t-1]}) \leftarrow^s \text{INIT}(CS)$
2. $lr.msg \leftarrow \mu$; $lr.SS \leftarrow [t + 1]$; For $i \in [t + 1]$ do $lr.PP(i) \leftarrow \perp$
3. For $i \in [t - 1]$ do $psig_i \leftarrow^s \text{DS.Sign}(sk_i, lr)$
4. $psig_t \leftarrow^s \text{PSIGNO}(t, lr)$
5. $sig \leftarrow (lr, [t], \{psig_i\}_{i \in [t]})$; Return (μ, sig)

We claim that $\text{Adv}_{TS,t}^{\text{ts-uf-3}} \mathcal{A} = 1$. The adversary has gathered t valid signatures, so the verification condition in $RTS_2[DS].Vf$ (Figure 8.4) is met. Now we need to show that the forgery is non-trivial, meaning $\text{tf}_3(lr) = 0$. Indeed, we have $curSS_L(lr) = \{t\}$ but $\{i \in HS \cap lr.SS : lr.PP(i) \in PP_i\} = \{t, t + 1\}$.

TS-SUF-3 $\not\Rightarrow$ TS-UF-4. We need to exhibit a scheme TS that is TS-SUF-3-secure but not TS-UF-4-secure. Let $TS = RTS_3[DS]$ where DS is a unique SUF-CMA-secure standard signature scheme. Proposition 8.3.1 says this achieves TS-SUF-3. We now give an attack showing it fails TS-UF-4. Letting E be the set of all $i \in lr.SS$ such that $lr.PP(i)$ is correct, meaning equals \perp , the idea is that the adversary can let E be a strict subset of $lr.SS$ and then pass verification using only signatures from E . Proceeding to the details of the attack, fix a message μ , and consider the following adversary \mathcal{A} for game $TS-UF-4_{TS}$:

Adversary \mathcal{A}

1. $CS \leftarrow \emptyset$; $(pk, aux, \emptyset) \leftarrow^s \text{INIT}(CS)$
2. $lr.msg \leftarrow \mu$; $lr.SS \leftarrow [1..t + 1]$
3. $lr.PP(t + 1) \leftarrow 0$; For $i \in [1..t]$ do $lr.PP(i) \leftarrow \perp$

4. For $i \in [1..t]$ do $psig_i \leftarrow_s \text{PSIGNO}(i, lr)$
5. $sig \leftarrow (lr, [1..t], \{psig_i\}_{i \in [1..t]})$; Return (μ, sig)

We claim that $\text{Adv}_{\text{TS}, t}^{\text{ts-uf-4}}(\mathcal{A}, \kappa) = 1$. From Figure 8.5, we have $E = [1..t]$ and thus $\text{TS.Vf}(\text{pk}, \mu, sig) = 1$. Now we need to show that the forgery is non-trivial, meaning $\text{tf}_4(lr) = 0$. Indeed, we have $\text{curSSL}(lr) = [1..t]$ but $HS \cap lr.SS = [1..t + 1]$.

TS-UF-4 $\not\Rightarrow$ TS-SUF-2. We need to exhibit a scheme TS that is TS-UF-4-secure but not TS-SUF-2-secure. First, if $x \in \{0, 1\}^*$, it is convenient to let $\text{pre}(x)$ be the first bit of x and $\text{suff}(x)$ the rest. Now, let DS^* be a unique SUF-CMA-secure standard signature scheme, and modify it to a scheme DS as follows: Let $\text{DS.Sign}(\cdot, \cdot) \leftarrow 0 \parallel \text{DS}^*.\text{Sign}(\cdot, \cdot)$ and let $\text{DS.Vf}(\cdot, \cdot, psig) \leftarrow \text{DS}^*.\text{Vf}(\cdot, \cdot, \text{suff}(psig))$. Then DS is UF-CMA-secure, but, since flipping the first bit of $psig$ does not affect its validity, not SUF-CMA-secure. Now, consider $\text{RTS}_4[\text{DS}]$, and, for it, an alternative verification algorithm $\text{RTS}_4[\text{DS}].\text{Vf}'$ that is as in Figure 8.4 except that line 17 is changed to $T \leftarrow \{i \in F : \text{DS.Vf}(\text{pk}_i, lr, psig_i) \text{ and } \text{pre}(psig_i) = 0\}$. Let TS be the same as $\text{RTS}_4[\text{DS}]$ except that we change SVf as follows: at line 22 of Figure 8.4, invoke $\text{RTS}_4[\text{DS}].\text{Vf}'$ rather than $\text{RTS}_4[\text{DS}].\text{Vf}$. (Note that TS.Vf stays as in $\text{RTS}_4[\text{DS}]$ as shown in Figure 8.4. The modified verification algorithm is only used by TS.SVf . Also note the latter meets the required condition of accepting at most one signature per key and message, due to the uniqueness of DS^* .) Proposition 8.3.1 says $\text{RTS}_4[\text{DS}]$ is TS-UF-4-secure, and hence so is TS, because the only difference between these two is in SVf and TS-UF-4 does not depend on this. We now give an attack showing TS fails TS-SUF-2. The idea is to exploit lack of SUF-CMA-security of the base scheme DS. Proceeding to the details of the attack, fix a message μ , and consider the following adversary \mathcal{A} for game $\text{TS-SUF-2}_{\text{TS}}$, where $\text{flip1}(x)$ returns string x with its first bit flipped:

Adversary \mathcal{A}

1. $CS \leftarrow \emptyset$; $(\text{pk}, \text{aux}, \emptyset) \leftarrow_s \text{INIT}(CS)$
2. $lr.\text{msg} \leftarrow \mu$; $lr.SS \leftarrow [1..t]$; For $i \in [1..t]$ do $lr.PP(i) \leftarrow \perp$
3. For $i \in [1..t]$ do $psig_i^* \leftarrow_s \text{PSIGNO}(i, lr)$; $psig_i \leftarrow \text{flip1}(psig_i^*)$
4. $sig \leftarrow (lr, [1..t], \{psig_i\}_{i \in [1..t]})$; Return (M, sig)

It is not hard to see that $\text{Adv}_{\text{TS}, t}^{\text{ts-suf-2}}(\mathcal{A}, \kappa) = 1$ from Figure 8.5. Now we need to show that the forgery is non-trivial, meaning $\text{tsf}_2(lr, \text{pk}, sig) = 0$. Indeed, $\text{TS.SVf}(\text{pk}, lr, sig) = 0$ because the first bit of $psig_i$ is 1 for all $i \in [1..t]$.

8.3.2 From TS-UF-0 to TS-UF-1, loosely

The following theorem shows TS-UF-1 security is implied by TS-UF-0 security, although with an exponential loss in t , which is acceptable in settings where t is expected to be constant.

Theorem 8.3.2. *Let TS be a threshold signature scheme. For any TS-UF-1 adversary \mathcal{A} there exists a TS-UF-0 adversary \mathcal{B} such that $\text{Adv}_{\text{TS}}^{\text{ts-uf-1}}(\mathcal{A}, \kappa) \leq \binom{n}{t-1} \cdot \text{Adv}_{\text{TS}}^{\text{ts-uf-0}}(\mathcal{B}, \kappa)$. Moreover, \mathcal{B} runs in time roughly equal that of \mathcal{A} , and the number of \mathcal{B} 's queries to each oracle is at most that of \mathcal{A} .*

If the adversary always corrupts $t - 1$ parties, it is clear that TS-UF-0 and TS-UF-1 are equivalent. Otherwise, in general, for an adversary that breaks TS-UF-1 security and corrupts a subset CS of signers with size less than $t - 1$, if the adversary wins the game $\text{TS-UF-1}_{\text{TS}}$ by outputting (μ^*, sig^*) , we know $|\text{curSS}_M(\mu^*)| < t - |CS|$. Therefore, we can modify the adversary to initially guess a subset $\overline{CS} \subseteq [1..n] \setminus CS$ with size $t - |CS| - 1$ and corrupt all parties in \overline{CS} . If \overline{CS} happens to contain $\text{curSS}_M(\mu^*)$, the adversary actually wins. It is not hard to see that the probability that this is true is $1/\binom{n-|CS|}{t-|CS|-1} \geq 1/\binom{n}{t-1}$.

Proof of Theorem 8.3.2. We describe a construction of the adversary \mathcal{B} as follows. \mathcal{B} runs \mathcal{A} with access to the oracles $\widetilde{\text{INIT}}$, $\widetilde{\text{PPO}}$, $\widetilde{\text{PSIGNO}}$, $\widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$ query: \mathcal{B} randomly samples a set $\overline{CS} \in [1..n] \setminus CS$ of size $t - |CS| - 1$ and makes an oracle query $\text{INIT}(CS \cup \overline{CS})$ in the game $\text{TS-UF-0}_{\text{TS}}$. After receiving $\text{pk}, \text{aux}, \{\text{sk}_i\}_{i \in CS \cup \overline{CS}}$, \mathcal{B} sets $\text{st}_i.\text{sk} \leftarrow \text{sk}_i$, $\text{st}_i.\text{vk} \leftarrow \text{pk}$, and $\text{st}_i.\text{aux} \leftarrow \text{aux}$ for all $i \in \overline{CS}$. Finally, \mathcal{B} returns $\text{pk}, \text{aux}, \{\text{sk}_i\}_{i \in CS}$.

$\widetilde{\text{PPO}}(i)$ query: Same as in the game $\text{TS-UF-1}_{\text{TS}}$, except when $i \in [1..n] \setminus (CS \cup \overline{CS})$, \mathcal{B} directly relays the query to oracle PPO in the game $\text{TS-UF-0}_{\text{TS}}$.

$\widetilde{\text{PSIGNO}}(i, lr)$ query: Same as in the game $\text{TS-UF-1}_{\text{TS}}$, except when $i \in [1..n] \setminus (CS \cup \overline{CS})$ directly relays the query to oracle PSIGNO in the game $\text{TS-UF-0}_{\text{TS}}$. In addition, denote curSS'_M as curSS_M defined in the game $\text{TS-UF-1}_{\text{TS}}$. \mathcal{B} also updates the set $\text{curSS}'_M(lr.\text{msg})$ the same as in the game $\text{TS-UF-1}_{\text{TS}}$.

$\widetilde{\text{RO}}(x)$ query: \mathcal{B} directly relays the query to oracle RO in the game $\text{TS-UF-0}_{\text{TS}}$.

After receiving the output (M^*, sig^*) from \mathcal{A} , denote the event GoodECS as $curSS'_M(M^*) \subseteq \overline{CS}$. If \mathcal{A} wins the game TS-UF-1_{TS} and GoodECS occurs, \mathcal{B} returns (M^*, sig^*) . Otherwise, \mathcal{B} aborts.

Denote the event WIN as \mathcal{A} wins the game TS-UF-1_{TS} simulated by \mathcal{B} . We first show \mathcal{B} wins the game TS-UF-0_{TS} if WIN \wedge GoodECS occurs. From the simulation, we know $curSS_M(M^*) = curSS'_M(M^*) \setminus \overline{CS}$, where $curSS_M(M^*)$ is defined in the game TS-UF-0_{TS}. Since WIN \wedge GoodECS occurs, we know (M^*, sig^*) is valid for the public key pk and $curSS'_M(M^*) = \emptyset$, which implies \mathcal{B} wins the game TS-UF-0_{TS}.

Therefore, it is left to show that $\Pr[\text{WIN} \wedge \text{GoodECS}] \geq \frac{1}{\binom{n}{t-1}} \text{Adv}_{\text{TS}}^{\text{ts-uf-1}}(\mathcal{A}, \kappa)$. We first fix a set $S \in [1..n]$ with size less than t and consider the case when $\overline{CS} = S$. If WIN occurs, we know $|curSS'_M(M^*)| < t - |CS|$. Since \mathcal{B} perfectly simulates the game TS-UF-0_{TS} no matter which \overline{CS} is picked, we know the set $curSS'_M(M^*)$ is independent of the choice of \overline{CS} , which implies

$$\begin{aligned} \Pr[\text{GoodECS} \mid \text{WIN} \wedge CS = S] &= \Pr[\overline{CS} \in curSS'_M(M^*) \mid \text{WIN} \wedge CS = S] \\ &\geq \frac{1}{\binom{n-|S|}{t-1-|S|}} \geq \frac{1}{\binom{n}{t-1}}. \end{aligned}$$

Therefore,

$$\begin{aligned} \Pr[\text{GoodECS}] &= \sum_{\substack{S \subseteq [1..n], \\ |S| < t}} \Pr[\text{GoodECS} \mid \text{WIN} \wedge CS = S] \cdot \Pr[\text{WIN} \mid CS = S] \\ &\geq \sum_{\substack{S \subseteq [1..n], \\ |S| < t}} \frac{1}{\binom{n}{t-1}} \Pr[\text{WIN} \mid CS = S] \\ &= \frac{1}{\binom{n}{t-1}} \Pr[\text{WIN}] = \frac{1}{\binom{n}{t-1}} \text{Adv}_{\text{TS}}^{\text{ts-uf-1}}(\mathcal{A}, \kappa). \end{aligned}$$

□

8.3.3 From TS-(S)UF-3 to TS-(S)UF-4.

Figure 8.6 gives a general transformation from TS-(S)UF-3 security to TS-(S)UF-4 security. Concretely, we give a construction ATS from any TS-(S)UF-3-secure echo scheme TS and a digital signature scheme DS. The size of signatures produced by ATS and the verification algorithm Vf are exactly the same as TS. The main idea is to use signatures to authenticate each token con-

<p><u>Algorithm Setup(1^κ) :</u></p> <p>$par_1 \leftarrow \text{TS.Setup}(1^\kappa)$ $par_2 \leftarrow \text{DS.Setup}(1^\kappa)$ Return (par_1, par_2)</p> <p><u>Algorithm KeyGen :</u></p> <p>$pk, \tau_{aux}, \{\text{tsk}_i\}_{i \in [1..n]} \leftarrow \text{TS.KeyGen}$ For $i \in [1..n]$ do $(\text{svk}_i, \text{ssk}_i) \leftarrow \text{DS.KeyGen}$ $sk_i \leftarrow (\text{tsk}_i, \text{ssk}_i)$ $aux \leftarrow (\tau_{aux}, \text{svk}_1, \dots, \text{svk}_n)$ Return $pk, aux, \{sk_i\}_{i \in [1..n]}$</p> <p><u>Algorithm SPP(st_i) :</u></p> <p>$(tpp, st_i) \leftarrow \text{SPP}(st_i)$ $(\text{tsk}_i, \text{ssk}_i) \leftarrow st_i.sk$ $tsig \leftarrow \text{DS.Sign}(\text{ssk}_i, tpp)$ Return $((tpp, tsig), st_i)$</p> <p><u>Algorithm LPP(i, pp, st_0) :</u></p> <p>$(tpp, tsig) \leftarrow pp$ $st_0.\text{SigMap}(i, tpp) \leftarrow tsig$ Return $\text{TS.LPP}(i, tpp, st_0)$</p> <p><u>Algorithm Vf($pk, \mu, sig$) :</u></p> <p>Return $\text{TS.Vf}(pk, \mu, sig)$</p>	<p><u>Algorithm OriginLR(lr) :</u></p> <p>For $i \in lr.SS$ do $(tpp, tsig) \leftarrow lr.PP(i)$ $lr.PP(i) \leftarrow tpp$ Return lr</p> <p><u>Algorithm LR(μ, SS, st_0) :</u></p> <p>$(lr, st_0) \leftarrow \text{TS.LR}(\mu, SS, st_0)$ For $i \in SS$ do $tpp_i \leftarrow lr.PP(i)$ $lr.PP(i) \leftarrow (tpp_i, st_0.\text{SigMap}(i, tpp_i))$ Return (lr, st_0)</p> <p><u>Algorithm PS(lr, i, st_i) :</u></p> <p>$(\tau_{aux}, \text{svk}_1, \dots, \text{svk}_n) \leftarrow st_i.aux$ For $i \in lr.SS$ do $(tpp_i, tsig_i) \leftarrow lr.PP(i)$ If $\text{DS.Vf}(\text{svk}_i, tpp_i, tsig_i) = 0$ then Return \perp Return $\text{TS.PS}(\text{OriginLR}(lr), i, st_i)$</p> <p><u>Algorithm Agg(PS, st_0) :</u></p> <p>Return $\text{TS.Agg}(\text{PS}, st_0)$</p> <p><u>Algorithm SVf(pk, lr, sig) :</u></p> <p>Return $\text{TS.SVf}(pk, \text{OriginLR}(lr), sig)$</p>
--	--

Figure 8.6: The threshold signature $\text{ATS}[\text{TS}, \text{DS}]$ constructed from an echo scheme TS and a digital signature scheme DS such that $\text{ATS}.n = \text{TS}.n$ and $\text{ATS}.t = \text{TS}.t$. The algorithm OriginLR transforms a well-formed leader request lr for ATS to a well-formed leader request in TS . $st_0.\text{SigMap}$ is a table that stores the signature corresponding to each token generated by honest signers, which is initially set to empty. PS denotes a set of partial signatures.

Game $\text{SUF-CMA}_{\text{DS}}^{\mathcal{A}}(\kappa)$:

$Q \leftarrow \emptyset$

$(\text{sk}, \text{pk}) \leftarrow_{\$} \text{DS.KeyGen}(1^\kappa)$

$(\mu^*, \text{sig}^*) \leftarrow \mathcal{A}^{\text{SIGN}}(\text{pk})$

Return $\text{DS.Vf}(\text{pk}, \mu^*, \text{sig}^*) \wedge (\mu^*, \text{sig}^*) \notin Q$

Oracle $\text{SIGN}(\mu)$:

$\text{sig} \leftarrow_{\$} \text{DS.Sign}(\text{sk}, \mu)$

$Q \leftarrow Q \cup \{(\mu, \text{sig})\}$

Return sig

Figure 8.7: The game $\text{SUF-CMA}_{\text{DS}}$, where DS is a digital signature scheme.

tained in a leader request lr from TS, so that an honest server only answers the request if all the authentications are valid. The rest of the protocol remains the same.

In the game $\text{TS-UF-4}_{\text{ATS}}$, we can show that as long as the adversary does not break the strong unforgeability of DS, for any leader request lr such that $\text{curSS}_L(lr) > 0$, it holds that $\{i \in HS \cap lr.SS : lr.PP(i) \in PP_i\} = HS \cap lr.SS$, which implies the conditions tf_3 and tf_4 are equivalent. Therefore, we can reduce TS-(S)UF-4 security of ATS to TS-(S)UF-3 security of TS and SUF-CMA security of DS. (The latter notion is formally defined via the game in Figure 8.7.) This is captured by the the following theorem.

Theorem 8.3.3. *Let $XX \in \{\text{SUF}, \text{UF}\}$. Let TS be an echo scheme and DS be a digital signature scheme. For any TS-XX-4 adversary \mathcal{A} there exists a TS-XX-3 adversary \mathcal{B} and a SUF-CMA adversary \mathcal{C} such that*

$$\text{Adv}_{\text{ATS}[\text{TS}, \text{DS}]}^{\text{ts-xx-4}}(\mathcal{A}, \kappa) \leq \text{Adv}_{\text{TS}}^{\text{ts-xx-3}}(\mathcal{B}, \kappa) + n \cdot \text{Adv}_{\text{DS}}^{\text{suf-cma}}(\mathcal{C}, \kappa).$$

Moreover, \mathcal{B} and \mathcal{C} run in time roughly equal that of \mathcal{A} . The number of \mathcal{B} 's queries to each oracle is at most that of \mathcal{A} . The number of \mathcal{C} 's SIGN queries is at most the number of PPO queries made by \mathcal{A} .

Proof. This proof only deals with TS-SUF-4 security, but a similar proof also works for TS-UF-4 security.

Let \mathcal{A} be the adversary described in the theorem. After \mathcal{A} returns, denote the event BadLR as there exists lr such that $\text{curSS}_L(lr) > 0$ and $\{i \in HS \cap lr.SS : lr.PP(i) \in PP_i\} \neq HS \cap lr.SS$. Denote the event WIN as \mathcal{A} wins the game TS-SUF-4_{ATS}. Then, we have

$$\text{Adv}_{\text{ATS}}^{\text{ts-suf-4}}(\mathcal{A}, \kappa) \leq \Pr[\text{WIN} \wedge (\neg \text{BadLR})] + \Pr[\text{BadLR}] .$$

Thus, we can conclude the theorem with the following two lemmas.

Lemma 8.3.4. *There exists a TS-SUF-3 adversary \mathcal{B} making at most q_{s1} queries to PPO, at most q_{s2} queries to PSIGNO, and at most q_h queries to RO such that*

$$\Pr[\text{WIN} \wedge (\neg \text{BadLR})] \leq \text{Adv}_{\text{TS}}^{\text{ts-suf-3}}(\mathcal{B}, \kappa) .$$

Moreover, \mathcal{B} runs in time roughly equal that of \mathcal{A}

Lemma 8.3.5. *There exists a SUF-CMA adversary \mathcal{C} making at most q_{s1} queries to SIGN such that*

$$\Pr[\text{BadLR}] \leq n \cdot \text{Adv}_{\text{DS}}^{\text{suf-cma}}(\mathcal{C}, \kappa) ,$$

Moreover, \mathcal{C} runs in time roughly equal that of \mathcal{A}

□

Proof of Lemma 8.3.4. We give a construction of the adversary \mathcal{B} in Figure 8.8, where \mathcal{B} runs \mathcal{A} by simulating the game TS-SUF-4_{ATS}. The simulation is done simply by relaying all queries from \mathcal{A} to the oracles in the game TS-SUF-3_{TS} and doing the extra authentication parts by \mathcal{B} itself. It is clear that \mathcal{B} simulates the game TS-SUF-4_{ATS} perfectly, which implies the probability that \mathcal{B} does not abort is equal to $\Pr[\text{WIN} \wedge (\neg \text{BadLR})]$.

Therefore, it is left to show that if \mathcal{B} does not abort, then \mathcal{B} wins the game TS-SUF-3_{TS}. Suppose $\text{WIN} \wedge (\neg \text{BadLR})$ occurs in the game TS-SUF-4_{ATS} simulated by \mathcal{B} . We use PP' and curSS'_L to denote the variables PP and curSS_L in the game TS-SUF-3_{TS}. For any PSIGNO query lr , denote $\tilde{lr} = \text{OriginLR}(lr)$, and we have $lr.\text{msg} = \tilde{lr}.\text{msg}$ and $\text{TS.SVf}(\text{pk}, \tilde{lr}, \text{sig}) = 1$ if and only if $\text{ATS.SVf}(\text{pk}, lr, \text{sig}) = 1$. Since the output (μ, sig) must be valid for the public key pk due to WIN occurs, to show \mathcal{B} wins the game TS-SUF-3_{TS}, we just need to show for any lr such that $lr.\text{msg} = \mu$ and $\text{ATS.SVf}(\text{pk}, lr, \text{sig}) = 1$, it holds that

$$(|\text{curSS}'_L(\tilde{lr})| < t - |CS|) \vee (\text{curSS}'_L(\tilde{lr}) \neq \{i \in HS \cap \tilde{lr}.SS : \tilde{lr}.PP(i) \in PP'_i\}) . \quad (8.1)$$

<p><u>Adversary $\mathcal{B}^{\text{INIT}, \text{PPO}, \text{PSIGNO}, \text{RO}}()$:</u></p> <p>For $i \in [1..n]$ do</p> <p style="padding-left: 20px;">$(\text{svk}_i, \text{ssk}_i) \leftarrow \text{DS.KeyGen}()$</p> <p style="padding-left: 20px;">$(\mu, \text{sig}) \leftarrow \mathcal{A}^{\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}, \widetilde{\text{RO}}}()$</p> <p>If $\text{WIN} \wedge (\neg \text{BadLR})$ occurs then</p> <p style="padding-left: 20px;">Return (μ, sig)</p> <p>Else abort</p> <p><u>Oracle $\widetilde{\text{INIT}}(CS)$:</u></p> <p>$\text{pk}, \text{taux}, \{\text{tsk}_i\}_{i \in CS} \leftarrow \text{INIT}(CS)$</p> <p>For $i \in CS$ do</p> <p style="padding-left: 20px;">$\text{sk}_i \leftarrow (\text{tsk}_i, \text{ssk}_i)$</p> <p>$\text{aux} \leftarrow (\text{taux}, \text{svk}_1, \dots, \text{svk}_n)$</p> <p>Return $\text{pk}, \text{aux}, \{\text{sk}_i\}_{i \in CS}$</p>	<p><u>Oracle $\widetilde{\text{PPO}}(i)$:</u></p> <p>$tpp \leftarrow \text{PPO}(i)$</p> <p>$\text{tsig} \leftarrow \text{DS.Sign}(\text{ssk}_i, tpp)$</p> <p>Return (tpp, tsig)</p> <p><u>Oracle $\widetilde{\text{PSIGNO}}(i, lr)$:</u></p> <p>For $i \in lr.\text{SS}$ do</p> <p style="padding-left: 20px;">$(pp_i, \text{tsig}_i) \leftarrow lr.\text{PP}(i)$</p> <p style="padding-left: 20px;">If $\text{DS.Vf}(\text{svk}_i, pp_i, \text{tsig}_i) = 0$ then</p> <p style="padding-left: 40px;">Return \perp</p> <p>Return $\text{PSIGNO}(i, \text{OriginLR}(lr))$</p> <p><u>Oracle $\widetilde{\text{RO}}(x)$:</u></p> <p>Return $\text{RO}(x)$</p>
---	--

Figure 8.8: Adversary \mathcal{B} for the proof of Lemma 8.3.4. \mathcal{B} also compute the sets L , PP , and $\text{curSS}_L(lr)$ for each $lr \in L$ following the same logic as in the game $\text{TS-SUF-4}_{\text{ATS}}$ and thus can check whether the event $\text{WIN} \wedge (\neg \text{BadLR})$ occurs.

Since WIN occurs, we know either $|\text{curSS}_L(lr)| < t - |CS|$ or $\text{curSS}_L(lr) \neq HS \cap lr.\text{SS}$. If $|\text{curSS}_L(lr)| < t - |CS|$, from the simulation, we know $|\text{curSS}'_L(\tilde{lr})| = |\text{curSS}_L(lr)| < t - |CS|$, which implies (8.1) holds. Otherwise, we have $|\text{curSS}_L(lr)| \geq t - |CS| > 0$ and $\text{curSS}_L(lr) \neq HS \cap lr.\text{SS}$. Since BadLR does not occur, we have $\{i \in HS \cap lr.\text{SS} : lr.\text{PP}(i) \in PP_i\} = HS \cap lr.\text{SS}$. Therefore, for any $i \in HS \cap lr.\text{SS}$, it holds that $lr.\text{PP}(i) \in PP_i$, which implies $\tilde{lr}.\text{PP}(i) \in PP'_i$. Since $\tilde{lr}.\text{SS} = lr.\text{SS}$, we have $\{i \in HS \cap \tilde{lr}.\text{SS} : \tilde{lr}.\text{PP}(i) \in PP'_i\} = HS \cap lr.\text{SS}$. Therefore, we have $\text{curSS}'_L(\tilde{lr}) = \text{curSS}_L(lr) \neq HS \cap lr.\text{SS} = \{i \in HS \cap \tilde{lr}.\text{SS} : \tilde{lr}.\text{PP}(i) \in PP'_i\}$, which implies (8.1) holds. \square

Proof of Lemma 8.3.5. We describe a construction of the adversary \mathcal{C} as follows. To start with, \mathcal{C} queries $\text{INIT}()$ oracle and receives svk^* . Also, \mathcal{C} initializes all the states $\text{st}_0, \dots, \text{st}_n$. Then, \mathcal{C} runs \mathcal{A} with access to the oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}, \widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$: Same as in the game $\text{TS-SUF-4}_{\text{ATS}}$, except \mathcal{C} additionally randomly picks an index

$i^* \in HS$ and sets $(svk_{i^*}, ssk_{i^*}) \leftarrow_{\$} (svk^*, \perp)$ instead of generating them by DS.KeyGen. Also, \mathcal{C} initializes H to an empty table.

$\widetilde{\text{PPO}}(i)$ query: Same as in the game TS-SUF-4_{ATS}, except when $i = i^*$, in the execution of $\text{SPP}[H](st_{i^*})$, \mathcal{C} computes $tsig \leftarrow_{\$} \text{SIGN}(tpp)$ instead of generating it by DS.Sign.

$\widetilde{\text{PSIGNO}}(i, lr)$ query: Same as in the game TS-SUF-4_{ATS}.

$\widetilde{\text{RO}}(x)$ query: If $H(x) \neq \perp$, \mathcal{C} returns $H(x)$. Otherwise, \mathcal{C} sets $H(x) \leftarrow_{\$} \mathbb{Z}_p$ and returns $H(x)$.

After receiving the output from \mathcal{A} , denote the event GoodInd as there exists $lr^* \in L$ such that $\text{curSS}_L(lr^*) > 0$, $i^* \in HS \cap lr^*.SS$ and $lr^*.PP(i^*) \notin PP_{i^*}$. If GoodInd does not occur, \mathcal{B} aborts. Otherwise, \mathcal{B} returns $(tpp_{i^*}, tsig_{i^*})$, where $(tpp_{i^*}, tsig_{i^*}) \leftarrow lr^*.PP(i^*)$.

We first show that \mathcal{B} wins the game $\text{SUF-CMA}_{\text{DS}}$ if GoodInd occurs. Since $\text{curSS}_L(lr^*) > 0$, we know for all $i \in lr^*.SS$, $\text{DS.Vf}(svk_i, tpp_i, tsig_i) = 1$ where $(tpp_i, tsig_i) \leftarrow lr^*.PP(i)$. From the simulation, we know $PP_{i^*} = Q$, where Q is defined in the game $\text{SUF-CMA}_{\text{DS}}$. Since $(tpp_{i^*}, tsig_{i^*}) = lr^*.PP(i^*) \notin lr^*.SS$, we know $(tpp_{i^*}, tsig_{i^*})$ is valid for $svk^* = svk_{i^*}$ and $(tpp_{i^*}, tsig_{i^*}) \notin Q$, which implies \mathcal{B} wins the game $\text{SUF-CMA}_{\text{DS}}$.

It is left show that $\Pr[\text{GoodInd}] \geq \frac{1}{n} \Pr[\text{BadLR}]$. We first fix a set $S \in [1..n]$ with size less than t and consider the case when $CS = S$. If BadLR occurs, then there exists $\tilde{lr} \in L$ such that $\text{curSS}_L(\tilde{lr}) > 0$ and $\{i \in HS \cap \tilde{lr}.SS : \tilde{lr}.PP(i) \notin PP_i\} \neq \emptyset$. Since \mathcal{C} perfectly simulates the game TS-SUF-4_{ATS} no matter which i^* is picked, we know the set $\{i \in HS \cap \tilde{lr}.SS : \tilde{lr}.PP(i) \notin PP_i\}$ is independent of the choice of i^* , which implies

$$\begin{aligned} & \Pr[\text{GoodInd} \mid \text{BadLR} \wedge CS = S] \\ &= \Pr\left[i^* \in \{i \in HS \cap \tilde{lr}.SS : \tilde{lr}.PP(i) \notin PP_i\} \mid \text{BadLR} \wedge CS = S\right] \\ &\geq \frac{1}{n - |S|} \geq \frac{1}{n}. \end{aligned}$$

Therefore, we have

$$\begin{aligned} \Pr[\text{GoodInd}] &= \sum_{\substack{S \subseteq [1..n], \\ |S| < t}} \Pr[\text{GoodInd} \mid \text{BadLR} \wedge CS = S] \cdot \Pr[\text{BadLR} \mid CS = S] \\ &\geq \sum_{\substack{S \subseteq [1..n], \\ |S| < t}} \frac{1}{n} \Pr[\text{BadLR} \mid CS = S] = \frac{1}{n} \Pr[\text{BadLR}]. \end{aligned}$$



Chapter 9

SECURITY OF THE FROST SCHEMES

9.1 The FROST Schemes

SCHEME DESCRIPTIONS. This section revisits the security of FROST, first proposed in [91] by Komlo and Goldberg, as a (partially) non-interactive threshold signature scheme. We consider both the original scheme, which we refer to as FROST1, as well as optimized versions, FROST2 and FROST3, from its follow-up work [50, 113]. We give a detailed description of all three schemes in Figures 9.1 and 9.2. The leader state st_0 contains a set $curPP_i$ for each server i representing the set of tokens generated by server i that has not yet been used in a signing request. The state st_i for server i contains a function $mapPP$ that maps each token pp to the randomness that is used to generate pp and $st_i.mapPP(pp) = \perp$ if pp is not generated by server i yet or has already been used in a signing request. All three schemes are echo schemes as defined in Section 8.1. We note that in the description of the protocols, we do not explicitly include the set $st_i.PP$ as required in Section 8.1 since its functionality is already achieved by $st_i.mapPP$. In particular, for pp_i defined in line 9.1, the condition $pp_i \notin st_i.PP$ implies $st_i.mapPP(pp_i) = \perp$ and thus PS returns (\perp, st_i) if $pp_i \notin st_i.PP$.

The coefficient $\lambda_i^{lr.SS}$ is the Lagrange coefficient for the set $lr.SS$, which is defined (for any set $\subseteq [1..n]$) as

$$\lambda_i^S := \prod_{j \in S, i \neq j} \frac{j}{j-i}. \quad (9.1)$$

The algorithm $CompPar$ is a helper algorithm that computes the parameters $R, c, \{d_i\}_{i \in lr.SS}$ used during signing. We stress that the only difference between FROST1 and FROST2 is the way d_i is computed in $CompPar$. The differences between FROST₃ and FROST₁/FROST₂ are in both the generation of lr and the computation of d_i . In FROST1, each d_i is a different hash value for each server i , while in FROST2, d_i 's are the same hash value for all servers. In FROST3, when generating lr , an additional aggregated token (\tilde{R}, \tilde{S}) is computed and put in lr . Then, d_i 's are computed from parts of lr instead of the entire lr . In particular, the hash function takes as input

<p><u>Protocol FROST1,FROST2,FROST3 :</u></p> <p><u>Algorithm Setup(1^κ) :</u> $(\mathbb{G}, p, g) \leftarrow \\$ \text{GGen}(1^\kappa)$ Return (\mathbb{G}, p, g)</p> <p><u>Algorithm KeyGen() :</u> For $i \in [0..t-1]$ do $a_i \leftarrow \\$ \mathbb{Z}_p$ For $i \in [1..n]$ do $\text{sk}_i \leftarrow \\$ \sum_{j=0}^{t-1} i^j \cdot a_j$; $\text{pk}_i \leftarrow g^{\text{sk}_i}$ $\text{pk} \leftarrow g^{a_0}$ $\text{aux} \leftarrow (\text{pk}_1, \dots, \text{pk}_n)$ Return $\text{pk}, \text{aux}, \{\text{sk}_i\}_{i \in [1..n]}$</p> <p><u>Algorithm SPP(st_i) :</u> $r \leftarrow \mathbb{Z}_p$; $s \leftarrow \mathbb{Z}_p$ $pp \leftarrow (g^r, g^s)$ $\text{st}_i.\text{mapPP}(pp) \leftarrow (r, s)$ Return (pp, st_i)</p> <p><u>Algorithm LPP(i, pp, st_0) :</u> $\text{st}_0.\text{curPP}_i \leftarrow \text{st}_0.\text{curPP}_i \cup \{pp\}$ Return st_0</p> <p><u>Algorithm Vf(pk, M, sig) :</u> $(R, z) \leftarrow sig$ $c \leftarrow \text{H}_2(\text{pk}, M, R)$ Return $(g^z = R \cdot \text{pk}^c)$</p>	<p><u>Algorithm PS(lr, i, st_i) :</u> $pp_i \leftarrow lr.\text{PP}(i)$ If $\text{st}_i.\text{mapPP}(pp_i) = \perp$ then Return (\perp, st_i) $(r_i, s_i) \leftarrow \text{st}_i.\text{mapPP}(pp_i)$ $\text{st}_i.\text{mapPP}(pp_i) \leftarrow \perp$ $(R, c, \{d_j\}_{j \in lr.SS})$ $\leftarrow \text{CompPar}(\text{st}_i.\text{vk}, lr)$ $z_i \leftarrow r_i + d_i \cdot s_i + c \cdot \lambda_i^{lr.SS} \cdot \text{st}_i.\text{sk}$ Return $((R, z_i), \text{st}_i)$</p> <p><u>Algorithm Agg(PS, st_0) :</u> $R \leftarrow \perp$; $z \leftarrow 0$ For $(R', z') \in \text{PS}$ do If $R = \perp$ then $R \leftarrow R'$ If $R \neq R'$ then return (\perp, st_0) $z \leftarrow z + z'$ Return $((R, z), \text{st}_0)$</p> <p><u>Algorithm SVf(pk, lr, sig) :</u> $(R^*, z^*) \leftarrow sig$ $(R, c, \{d_j\}_{j \in lr.SS})$ $\leftarrow \text{CompPar}(\text{pk}, lr)$ Return $(R = R^*) \wedge (g^{z^*} = R \cdot \text{pk}^c)$</p>
---	---

Figure 9.1: The common components of FROST1[GGen], FROST2[GGen] and FROST3[GGen], where GGen is a group generation algorithm. The algorithms LR and CompPar are defined in Figure 9.2 Further, n is the number of parties, and t is the threshold of the schemes. PS denotes a set of partial signatures. $\lambda_i^{lr.SS}$ denotes the Lagrange coefficients, which is defined in Equation (9.1).

<p>Protocol $\boxed{\text{FROST1}}$, FROST2, $\boxed{\text{FROST3}}$:</p> <p>Algorithm LR(M, SS, st_0) :</p> <p>If $\exists i \in SS : st_0.curPP_i = \emptyset$ then Return \perp</p> <p>$lr.msg \leftarrow M ; lr.SS \leftarrow SS$</p> <p>For $i \in SS$ do</p> <p> Pick pp_i from $st_0.curPP_i$</p> <p> $(R_i, S_i \leftarrow pp_i)$</p> <p> $lr.PP(i) \leftarrow pp_i$</p> <p> $st_0.curPP_i \leftarrow st_0.curPP_i \setminus \{pp_i\}$</p> <p>$lr.aggPP \leftarrow (\prod_{i \in SS} R_i, \prod_{i \in SS} S_i)$</p> <p>Return (lr, st_0)</p>	<p>Algorithm CompPar(pk, lr) :</p> <p>$M \leftarrow lr.msg$</p> <p>$d \leftarrow H_1(pk, lr.SS, lr.msg, lr.aggPP)$</p> <p>For $i \in lr.SS$ do</p> <p> $d_i \leftarrow H_1(pk, lr, i)$</p> <p> $d_i \leftarrow H_1(pk, lr)$</p> <p> $d_i \leftarrow d$</p> <p> $(R_i, S_i) \leftarrow lr.PP(i)$</p> <p> $R \leftarrow \prod_{i \in lr.SS} R_i S_i^{d_i}$</p> <p> $(\tilde{R}, \tilde{S}) \leftarrow lr.aggPP$</p> <p> $R \leftarrow \tilde{R} \tilde{S}^d$</p> <p>$c \leftarrow H_2(pk, M, R)$</p> <p>Return $(R, c, \{d_i\}_{i \in lr.SS})$</p>
---	--

Figure 9.2: The algorithms LR and CompPar of FROST1[GGen], FROST2[GGen] and FROST3[GGen], where GGen is a group generation algorithm. In particular, FROST1 contains all the dashed boxes, FROST2 contains all the highlighted boxes, and FROST3 contains all the solid boxes. The function $H_i(\cdot)$ is computed as $H(i, \cdot)$ for $i = 1, 2$. PS denotes a set of partial signatures.

only the aggregated token, rather than all tokens in $lr.PP$. We also note that upon receiving lr , each signer $i \in lr.SS$ uses only $lr.PP(i)$ rather than the entire $lr.PP$. Therefore, in practice, the leader can reduce communication complexity by sending just $lr.PP(i)$ along with the other parts of lr .

It is not hard to verify that both schemes satisfy perfect correctness.

OVERVIEW OF OUR RESULTS. Crites, Komlo, and Maller [50] argue that FROST2 improves the signing efficiency of FROST1 as the number of exponentiations for computing the nonce R is reduced from at least t to one, but they only consider TS-UF-0 security of FROST2. Ruffing et al. [113] also only show TS-UF-0 security of FROST3. In this section, we strengthen the prior results (however, in a setting without distributed key generation) by showing that FROST2 is actually TS-SUF-2-secure but not TS-UF-3 secure, and FROST3 is TS-UF-1-secure but not TS-UF-2-

<p>Games AOMDL_{GGen} :</p> <p>cid \leftarrow 0; $\ell \leftarrow$ 0; $T \leftarrow ()$</p> <p>$\{y_i\}_{i \in [\text{cid}]} \leftarrow \mathcal{A}^{\text{CHAL}, \text{DLOG}}()$</p> <p>Return $(\ell < \text{cid}) \wedge (\forall i \in [\text{cid}] : y_i = x_i)$</p>	<p>Oracle DLOG($\{\beta_i \in \mathbb{Z}_p\}_{i \in [\text{cid}]}$) :</p> <p>$\ell \leftarrow \ell + 1$</p> <p>Return $\sum_{i \in [\text{cid}]} \beta_i x_i$</p> <p>Oracle CHAL() :</p> <p>cid \leftarrow cid + 1; $x_{\text{cid}} \leftarrow_{\\$} \mathbb{Z}_p$</p> <p>Return $g^{x_{\text{cid}}}$</p>
---	---

Figure 9.3: The AOMDL game, where \mathbb{G} is a cyclic group with prime order p and generator g .

secure. In contrast, we show FROST1 is TS-SUF-3-secure but not TS-UF-4-secure. Theoretically, our results imply the separations between TS-(S)UF-2 and TS-(S)UF-3 and between TS-(S)UF-3 and TS-(S)UF-4. Practically speaking, our results indicate a separation between the security of FROST1 and FROST2. To complete the picture, a TS-SUF-4 secure variant of FROST1 can be obtained via the general transformation from Theorem 8.3.3, although it is an interesting open question whether a more efficient variant exists.

9.2 Security of FROST2

We first show that FROST2 is TS-SUF-2-secure in the ROM under the AOMDL assumption. The AOMDL assumption, introduced in [105], is formally defined in Figure 9.3. Formally, we show the following theorem.

Theorem 9.2.1. *For any TS-SUF-2 adversary \mathcal{A} making at most q_s queries to PPO and at most q_h queries to RO, there exists an AOMDL adversary \mathcal{B} making at most $2q_s + n$ queries to CHAL such that*

$$\text{Adv}_{\text{FROST2}[\text{GGen}]}^{\text{ts-suf-2}}(\mathcal{A}, \kappa) \leq \sqrt{q \cdot (\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) + 3q^2 \cdot 2^{-\kappa})},$$

where $q = q_s + q_h + 1$. Moreover, \mathcal{B} runs in time roughly equal two times that of \mathcal{A} , plus the time to perform at most $(4n + 2) \cdot q + 2q_s + 2n^2$ exponentiations and group operations.

The previous analysis of FROST2 [50] can be seen as implying TS-SUF-0 security, either in the AGM or under non-standard assumptions (which are, in turn, validated in the AGM). Our result here proves stronger security, without relying on the AGM, but also without considering

FROST's DKG. (We believe our analysis should extend, at least in the AGM, but we omit the added complexity of the DKG in this paper.) The core of the proof is a reduction from OMDL, which will need to use rewinding (via a variant of the Forking Lemma). The main challenge is to ensure that the reduction can simulate properly with a number of queries to DLOG which is smaller than the number of DL challenges. Further below, we are going to show that FROST2 is not TS-UF-3-secure, thus showing the above result is optimal with respect to our hierarchy.

Proof of Theorem 9.2.1. Let \mathcal{A} be an adversary as described in the theorem. Denote the output message-signature pair of \mathcal{A} as $(M^*, sig^* = (R^*, z^*))$. Without loss of generality, we assume \mathcal{A} always queries RO on $H_2(pk, M^*, R^*)$ before \mathcal{A} returns and always queries RO on $H_1(pk, lr)$ prior to the query PSIGNO(i, lr) for some i and lr . (This adds up to q_s additional RO queries, and we let $q = q_h + q_s + 1$.) Denote lr^* as the leader query such that $H_1(pk, lr^*)$ is the first query prior to the query $H_2(pk, M^*, R^*)$ satisfying $SVf(pk, lr^*, sig^*) = \text{true}$. If such lr^* does not exist, lr^* is set to \perp . Denote the event E_1 as

$$\text{Vf}(pk, M^*, sig^*) \wedge (lr^* = \perp \vee \text{curSSL}(lr^*) < t - |CS|).$$

It is clear that if \mathcal{A} wins the game TS-SUF-2_{FROST2}, then E_1 must occur, which implies $\Pr[E_1] \geq \text{Adv}_{\text{FROST2}[\text{GGen}]}^{\text{ts-suf-2}}(\mathcal{A}, \kappa)$. Therefore, the theorem will follow from the following lemma. (We isolate this statement as its own lemma also because it will be helpful in the proof of Theorem 9.3.1 below.) \square

Lemma 9.2.2. *There exists an AOMDL adversary \mathcal{B} making at most $2q_s + t$ queries to CHAL such that*

$$\Pr[E_1] \leq \sqrt{q \cdot (\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) + 3q^2 \cdot 2^{-\kappa})}.$$

Moreover, \mathcal{B} runs in time roughly twice that of \mathcal{A} , plus the time to perform at most $(4n + 2) \cdot q + 2q_s + 2n^2$ exponentiations and group operations.

Before turning to the proof of Lemma 9.2.2, we first introduce the following variant of the forking lemma that will be used within its proof.

Lemma 9.2.3. *Let $q \geq 1$ be an integer, $S \subseteq [1..q]$ be a set, and H be a set. Let \mathcal{A} be a randomized algorithm that on input x, h_1, \dots, h_q outputs a pair (I, Out) , where $I \in \{\perp\} \cup S$ and Out is a side output. Let IG be a randomized algorithm that generates x . The accepting probability of \mathcal{A} is*

Algorithm Fork^A(x) :
 Pick the random coin ρ of \mathcal{A} at random
 $h_1, h'_1, \dots, h_q, h'_q \leftarrow H$
 $(I, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q; \rho)$
 If $I = \perp$ then return \perp
 $(I', \text{Out}') \leftarrow \mathcal{A}(x, h_1, \dots, h_{I-1}, h'_I, \dots, h'_q; \rho)$
 If $I \neq I'$ then return \perp
 Return $(I, \text{Out}, \text{Out}')$

Figure 9.4: The forking algorithm build from \mathcal{A} used in Lemma 9.2.3..

defined as

$$\text{acc}(\mathcal{A}) = \Pr_{x \leftarrow \mathcal{IG}, h_1, \dots, h_q \leftarrow H} [(I, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q) : I \neq \perp] .$$

Consider algorithm Fork^A described in Figure 9.4. The accepting probability of Fork^A is defined as

$$\text{acc}(\text{Fork}^A) = \Pr_{x \leftarrow \mathcal{IG}} [\alpha \leftarrow \mathcal{A}(\text{Fork}^A(x)) : \alpha \neq \perp] .$$

Then, $\text{acc}(\text{Fork}^A) \geq \text{acc}(\mathcal{A})^2/|S|$.

The above lemma slightly extends the generalized Forking Lemma of Bellare and Neven [17] in the sense that if \mathcal{A} can only output index I within a given set $S \subseteq [1..q]$, then the final bound on $\text{acc}(\text{Fork}^A)$ depends only on $|S|$ instead of q . The property allows us to get better bounds in our analysis.

Proof of Lemma 9.2.3. For any $i \in S$, $h_1, \dots, h_{i-1} \in H$, and input x , define

$$Y_i(x, h_1, \dots, h_{i-1}) := \Pr_{h_i, \dots, h_q \leftarrow H} [I = i : (I, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q)] .$$

Then, we have

$$\begin{aligned} \text{acc}(\mathcal{A}) &= \sum_{i \in S} \Pr_{x \leftarrow \mathcal{IG}, h_1, \dots, h_q \leftarrow H} [I = i : (I, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q)] \\ &= \sum_{i \in S} \mathbb{E}_{x \leftarrow \mathcal{IG}, h_1, \dots, h_{i-1} \leftarrow H} [Y_i(x, h_1, \dots, h_{i-1})] . \end{aligned}$$

Thus, we have

$$\text{acc}(\text{Fork}^{\mathcal{A}}) = \sum_{i \in S} \Pr_{x \leftarrow \mathbb{G}, h_1, \dots, h_n, h'_1, \dots, h'_n \leftarrow \mathfrak{H}} [I = I' = i : \quad (9.2)$$

$$(I, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q),$$

$$(I', \text{Out}') \leftarrow \mathcal{A}(x, h_1, \dots, h_{i-1}, h'_i, \dots, h'_n)]$$

$$= \sum_{i \in S} \mathbb{E}_{x \leftarrow \mathbb{G}, h_1, \dots, h_{i-1} \leftarrow \mathfrak{H}} [Y_i(x, h_1, \dots, h_{i-1})^2] \quad (9.3)$$

$$\geq \sum_{i \in S} \left(\mathbb{E}_{x \leftarrow \mathbb{G}, h_1, \dots, h_{i-1} \leftarrow \mathfrak{H}} [Y_{i,j}(x, h_1, \dots, h_{i-1})] \right)^2 \quad (9.4)$$

$$\geq \frac{1}{|S|} \cdot \left(\sum_{i \in S} \mathbb{E}_{x \leftarrow \mathbb{G}, h_1, \dots, h_{i-1} \leftarrow \mathfrak{H}} [Y_i(x, h_1, \dots, h_{i-1})] \right)^2 \quad (9.5)$$

$$= \frac{\text{acc}(\mathcal{A})^2}{|S|}, \quad (9.6)$$

where Equation (9.4) is due to the fact that $\mathbb{E}[X^2] \geq (\mathbb{E}[X])^2$ and Equation (9.5) is due to the fact that $\sum_{i=1}^n a_i^2 \geq \frac{1}{n} (\sum_{i=1}^n a_i)^2$. \square

Proof of Lemma 9.2.2. We first construct an algorithm \mathcal{C} compatible with the syntax in Lemma 9.2.3. The input of \mathcal{C} consists of $(2q_s + t)$ uniformly random group elements $A_0, \dots, A_{t-1}, U_1, V_1, \dots, U_{q_s}, V_{q_s} \in \mathbb{G}$ and uniformly random integers $h_1, \dots, h_{2q} \in \mathbb{Z}_p$. Also, \mathcal{C} can access an oracle DLOG , which on input $\alpha \in \mathbb{Z}_p^{2q_s+t}$ outputs $\text{DLog}_{\mathbb{G},g}(A_0^{\alpha_1} \cdots A_{t-1}^{\alpha_t} U_1^{\alpha_{t+1}} V_1^{\alpha_{t+2}} \cdots U_{q_s}^{\alpha_{t+2q_s-1}} V_{q_s}^{\alpha_{t+2q_s}})$. (We can think of this oracle as part of \mathcal{C} in the context of the Forking Lemma, as \mathcal{C} does not need to be efficient.) To start with, \mathcal{C} initializes all the states $\text{st}_0, \dots, \text{st}_n$. In addition, it initializes counters $\text{ctr}_s, \text{ctr}_h$ to 0 and a function dt to an empty table, which are used to record the DLOG query related to each (U_j, V_j) . \mathcal{C} also initializes $\text{curLR} \leftarrow \emptyset$ to record all leader requests that appears during the game and initializes ctrPP to an empty table, which are used to record the counter corresponding to each token generated by honest parties. We also use a flag BadPPO to denote whether a bad event occurs, which are initially set to `false`. Then, \mathcal{C} runs \mathcal{A} with access to the oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}, \widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$: \mathcal{C} initializes H to an empty table and sets $\text{pk} \leftarrow A_0$, $\text{pk}_i = \prod_{j=0}^{t-1} A_j^{i^j}$ for $i \in [1..n]$, and $\text{sk}_i = \text{DLOG}(\text{pk}_i)$ for $i \in CS$. Finally, \mathcal{C} returns $\text{pk}, \text{aux} = (\text{pk}_1, \dots, \text{pk}_n), \{\text{sk}_i\}_{i \in CS}$.

$\widetilde{\text{RO}}$ query $\text{H}_1(x)$: If $\text{H}_1(x) \neq \perp$, \mathcal{C} returns $\text{H}_1(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, lr)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $\text{H}_1(x) \leftarrow \mathbb{Z}_p$ and returns $\text{H}_1(x)$. Otherwise, \mathcal{C} increases ctr_h by 1,

sets $H_1(x) \leftarrow h_{2\text{ctr}_h-1}$, and adds lr to curLR . Also, \mathcal{C} computes $R \leftarrow \prod_{i \in lr.\text{SS}} R_i S_i^{h_{2\text{ctr}_h-1}}$, where $(R_i, S_i) \leftarrow lr.\text{PP}(i)$. If $H_2(\text{pk}, lr.\text{msg}, R) = \perp$, \mathcal{C} sets $H_2(\text{pk}, lr.\text{msg}, R) = h_{2\text{ctr}_h}$. In addition, define $\text{mapLR}(\text{ctr}_h) := lr$ and set $\text{curLR} \leftarrow \text{curLR} \cup \{lr\}$. Finally, \mathcal{C} returns $H_1(x)$.

$\widetilde{\text{RO}}$ query $H_2(x)$: If $H_2(x) \neq \perp$, \mathcal{C} returns $H_2(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, M, R)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_2(x) \leftarrow_s \mathbb{Z}_p$ and returns $H_2(x)$. Otherwise, \mathcal{C} increases ctr_h by 1 and sets $H_2(x) \leftarrow h_{2\text{ctr}_h}$. Finally, \mathcal{C} returns $H_2(x)$.

$\widetilde{\text{PPO}}(i)$ query: Same as in the game $\text{TS-SUF-2}_{\text{FROST}_2}$, except in the simulation of algorithm SPP, \mathcal{C} first increases ctr_s by 1 and sets $pp \leftarrow (U_{\text{ctr}_s}, V_{\text{ctr}_s})$, $\text{st}_i.\text{mapPP}(pp) \leftarrow (0, 0)$, and $\text{ctrPP}(i, pp) \leftarrow \text{ctr}_s$. In addition, BadPPO is set to true if there exists $lr \in \text{curLR}$ such that $lr.\text{PP}(i) = (U_{\text{ctr}_s}, V_{\text{ctr}_s})$.

$\widetilde{\text{SIGNO}}(i, lr)$ query: Same as in the game $\text{TS-SUF-2}_{\text{FROST}_2}$, except in the simulation of algorithm PS, if $\text{st}_i.\text{mapPP}(pp) \neq \perp$, let $j \leftarrow \text{ctrPP}(i, lr.\text{PP}(i))$, and \mathcal{C} computes $\alpha \in \mathbb{Z}_p^{t+2q_s}$ as

$$\alpha_k := \begin{cases} c \cdot \lambda_k^{lr.\text{SS}} \cdot i^{k-1}, & \text{for } k \in [t], \\ 1, & \text{for } k = t + 2j - 1, \\ d_i, & \text{for } k = t + 2j, \\ 0, & \text{o.w.} \end{cases}$$

and sets $z_i \leftarrow \text{DLOG}(\alpha)$, which equals to $\text{DLOG}_{\mathbb{G}, g} \left(U_j V_j^{d_i} \text{pk}_i^{c\lambda_i^{lr.\text{SS}}} \right)$. In addition, \mathcal{C} sets $\text{dt}(j) \leftarrow (i, k, d_i, c\lambda_i^{lr.\text{SS}}, z_i)$, where k denotes the index such that $H_1(\text{pk}, lr)$ is set to h_{2k-1} during the simulation.

After receiving the output $(M^*, \text{sig}^* = (R^*, z^*))$ from \mathcal{A} , \mathcal{C} returns \perp if $\text{BadPPO} = \text{true}$ or E_1 does not occur. Otherwise, \mathcal{C} finds the index I such that $H_2(\text{pk}, M^*, R^*)$ is set to h_I during the simulation. By our assumption of \mathcal{A} , we know such I must exist. Then, \mathcal{C} returns (I, Out) , where Out consists of all variables received or generated by \mathcal{C} .

ANALYSIS OF \mathcal{C} . To use Lemma 9.2.3, we define $S := \{2k\}_{k \in [1..q]}$ and IG as the algorithm that samples $2q_s + t$ group elements uniformly from \mathbb{G} and outputs them. From the simulation,

we know the output index I of \mathcal{C} is always in S . Also, it is clear that \mathcal{C} simulates the game $\text{TS-SUF-2}_{\text{FROST}_2}$ perfectly when all the inputs of \mathcal{C} are uniformly sampled from their domain, which implies $\text{acc}(\mathcal{C}) \geq \Pr[E_1] - \Pr[\text{BadPPO}]$, where $\Pr[E_1]$ refers to the probability in the original $\text{TS-SUF-2}_{\text{FROST}_2}$ game with \mathcal{A} (as in the lemma statement), whereas $\Pr[\text{BadPPO}]$ is the probability that $\text{BadPPO} = \text{true}$ at the end of \mathcal{C} 's execution. Since every pair U_j, V_j is sampled uniformly from \mathbb{G} , for each $\text{PPO}(i)$ query, the probability BadPPO is set to true is less than $|\text{curLR}|/|G| \leq q_h/p$. Therefore, we have $\Pr[\text{BadPPO}] \leq q_s q_h/p$. By Lemma 9.2.3,

$$\begin{aligned} \text{acc}(\text{Fork}^{\mathcal{C}}) &\geq (\Pr[E_1] - q_s q_h/p)^2/q \geq \Pr[E_1]^2/q - 2\Pr[E_1]q_s q_h/(p \cdot q) \\ &\geq \Pr[E_1]^2/q - 2q_s/p. \end{aligned}$$

CONSTRUCT \mathcal{B} FROM $\text{Fork}^{\mathcal{C}}$. We now give a construct of the OMDL adversary \mathcal{B} using $\text{Fork}^{\mathcal{C}}$, and the available DLOG oracle. To start with, \mathcal{B} queries CHAL oracle $2q_s + t$ times to generate $A_0, \dots, A_{t-1}, U_1, V_1, \dots, U_{q_s}, V_{q_s}$ as the input of $\text{Fork}^{\mathcal{C}}$ and runs $\text{Fork}^{\mathcal{C}}$. Without loss of generality, we can assume all the AOMDL challenges are different, since otherwise, \mathcal{B} can solve them trivially. All DLOG queries from $\text{Fork}^{\mathcal{C}}$ are relayed by \mathcal{B} to DLOG oracle of the game $\text{AOMDL}_{\text{GGen}}$. Denote the event BadHash as any two of the scalars $h_1, h'_1, \dots, h_q, h'_q$ generated in the execution of $\text{Fork}^{\mathcal{C}}$ are same. Since $h_1, h'_1, \dots, h_q, h'_q$ are sampled uniformly from \mathbb{Z}_p , we know $\Pr[\text{BadHash}] \leq 2q^2/p$.

It is left to show that if $\text{Fork}^{\mathcal{C}}$ returns $(I, \text{Out}, \text{Out}')$ and BadHash does not occur, \mathcal{B} can win the game $\text{AOMDL}_{\text{GGen}}$, which implies

$$\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) \geq \text{acc}(\text{Fork}^{\mathcal{C}}) - \Pr[\text{BadHash}] \geq \Pr[E_1]^2/q - 3q^2/p.$$

We directly use the notations in the description of \mathcal{C} to denote the variables in Out and use $(\cdot)'$ to denote the variables in Out' . By the execution of $\text{Fork}^{\mathcal{C}}$, we know $(\text{pk}, M^*, R^*) = (\text{pk}', M^{*'}, R^{*'})$ and $\text{pk} = A_0$. Since $I \in S$, let $k^* = I/2$. It is not hard to see that $\text{mapLR}(k^*) = lr^*$. (If $\text{mapLR}(k^*) = \perp$, lr^* is also \perp .)

We first show how to compute the discrete log of A_0, \dots, A_{t-1} . Denote the discrete log of A_0, \dots, A_{t-1} as a_0, \dots, a_{t-1} and define a polynomial $f(x) := \sum_{i=0}^{t-1} a_i x^i$. Since BadHash does not occur, we have $H_2(\text{pk}, M^*, R^*) = h_I \neq h'_I = H'_2(\text{pk}, M^*, R^*)$. Since $g^{z^*} = R^x A_0^{h_I}$, $g^{z^{*'}} = R^x A_0^{h'_I}$, \mathcal{B} computes $f(0) = a_0 = \frac{z^* - z^{*'}}{h_I - h'_I}$. Define $T_{\text{dt}} := \{j : (i, k, d, c, z) \leftarrow \text{dt}(j), k = k^*\}$. For each $j \in T_{\text{dt}} \cap T_{\text{dt}'}$, let $(i, k, d, c, z) \leftarrow \text{dt}(j)$ and $(i', k', d', c', z') \leftarrow \text{dt}'(j)$, and we have $g^z = U_j V_j^d \text{pk}_i^c$, $g^{z'} = U_j V_j^{d'} \text{pk}_{i'}^{c'}$. Since $\text{BadPPO} = \text{false}$ during both execution of

\mathcal{C} , we know (U_j, V_j) is returned by a query $\text{PPO}(i)$ prior to the query $H_2(\text{pk}, M^*, R^*)$ during the first execution of \mathcal{C} . Since the two executions of \mathcal{C} are exactly the same prior to the query $H_2(\text{pk}, M^*, R^*)$, we know $i' = i$. Also, we know $d = h_k = h_{k^*} = h_{k'} = d'$. Therefore, \mathcal{B} can compute $f(i) = \text{DLog}_{\mathbb{G}, g}(\text{pk}_i) = \frac{z-z'}{c-c'}$. Denote $D := \{i\}_{j \in T_{\text{dt}} \cap T_{\text{dt}'}, (i, k, d, c, z) \leftarrow \text{dt}(j)}$. Since E_1 occurs in the first execution of \mathcal{C} , we know $|T_{\text{dt}}| = |\text{curSS}_L(lr^*)| < t - |CS|$. Therefore, we know $|D| = |T_{\text{dt}} \cap T_{\text{dt}'}| < t - |CS|$. Therefore, \mathcal{B} can pick an arbitrary set $D' \in HS \setminus D$ with size $(t - |CS| - |T_{\text{dt}} \cap T_{\text{dt}'}| - 1)$ and for each $i \in D'$, \mathcal{B} queries DLOG oracle on pk_i . Therefore, \mathcal{B} knows the value of $f(i)$ for $i \in CS \cup D \cup D' \cup \{0\}$. Since $|CS \cup D \cup D' \cup \{0\}| = t$, \mathcal{B} can compute the value of a_0, \dots, a_{t-1} using Lagrange interpolation.

We now show how to compute the discrete log of $U_1, V_1, \dots, U_{q_s}, V_{q_s}$. Denote their discrete log as $u_1, v_1, \dots, u_{q_s}, v_{q_s}$. From the execution of \mathcal{C} , we know $\text{dt}(j) = (i, k, d, c, z) \neq \perp$ if and only if \mathcal{C} queries DLOG on $U_j V_j^d \text{pk}_i^c$. Therefore, denote $\text{DLOG}(U_j V_j^d \text{pk}_i^c)$ as the DLOG query associated with $\text{dt}(j)$. For each $j \in q_s$, there are the following cases.

Case 0: Both $\text{dt}(j)$ and $\text{dt}'(j)$ are \perp . In this case, \mathcal{B} computes u_j, v_j by directly querying oracle DLOG.

Case 1: Exactly one of $\text{dt}(j)$ and $\text{dt}'(j)$ is not \perp . Without loss of generality, assume $\text{dt}(j) = (i, k, d, c, z)$, which implies $g^z = U_j V_j^d \text{pk}_i^c$. \mathcal{B} computes v_j by directly querying oracle DLOG and computes $u_j = z - d \cdot v_j - c \cdot f(i)$.

For all the following cases, both $\text{dt}(j)$ and $\text{dt}'(j)$ are not \perp and we denote $(i, k, d, c, z) \leftarrow \text{dt}(j)$ and $(i', k', d', c', z') \leftarrow \text{dt}'(j)$.

Case 2: $k \neq k'$ or $k = k' > k^*$. In this case, we know $d = h_k \neq h_{k'} = d'$ and $g^z = U_j V_j^d \text{pk}_i^c$, $g^{z'} = U_j V_j^{d'} \text{pk}_{i'}^{c'}$. Therefore, \mathcal{B} computes $v_j = \frac{z - c \cdot f(i) - z' + c' \cdot f(i')}{d - d'}$, $u_j = z - d \cdot v_j - c \cdot f(i)$.

Case 3: $k = k' = k^*$. In this case, \mathcal{B} computes v_j, u_j the same as Case 1.

Case 4: $k = k' < k^*$. \mathcal{B} computes v_j, u_j the same as Case 1. Also, in this case, we have $d = d'$ and $c = c'$. Therefore, \mathcal{B} queries DLOG oracle once in order to simulate the DLOG queries associated with $\text{dt}(j)$ and $\text{dt}'(j)$.

We now count the number of DLOG queries made by \mathcal{B} .

- \mathcal{B} queries DLOG oracle $|CS|$ times queries for simulating query $\text{DLOG}(\text{pk}_i)$ made by \mathcal{C} for each $i \in CS$.
- \mathcal{B} queries DLOG oracle $|D'|$ times queries for computing a_0, \dots, a_{t-1} .
- For each $j \in \mathfrak{q}_s$, \mathcal{B} queries DLOG twice for simulating query associated with $\text{dt}(j)$ and $\text{dt}'(j)$ and computing u_j, v_j in case 0, 1, 2, 4 and queries 3 times in case 3.

Since the condition of case 3 is equivalent to $j \in T_{\text{dt}} \cap T_{\text{dt}'}$, the total number of DLOG queries made by \mathcal{B} is equal to $2\mathfrak{q}_s + |T_{\text{dt}} \cap T_{\text{dt}'}| + |CS| + |D'| = 2\mathfrak{q}_s + t - 1$. Therefore, \mathcal{B} wins the game $\text{AOMDL}_{\text{GGen}}$. \square

9.3 Security of FROST1

In this section, we show that FROST1 is TS-SUF-3-secure in the ROM under the OMDL assumption. Formally, we show the following theorem.

Theorem 9.3.1. *Let GGen be a group generation algorithm. For any TS-SUF-3 adversary \mathcal{A} making at most \mathfrak{q}_s queries to PPO and at most \mathfrak{q}_h queries to RO, there exists an AOMDL adversary \mathcal{B} making at most $2\mathfrak{q}_s + t$ queries to CHAL such that*

$$\text{Adv}_{\text{FROST1}[\text{GGen}]}^{\text{ts-suf-3}}(\mathcal{A}, \kappa) \leq 4n \cdot q \cdot \sqrt{\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) + 6q \cdot 2^{-\kappa}},$$

where $q = \mathfrak{q}_s + \mathfrak{q}_h + 1$. Moreover, \mathcal{B} runs in time roughly equal two times that of \mathcal{A} , plus the time to perform at most $6n \cdot q + 4\mathfrak{q}_s + 2n^2$ exponentiations and group operations.

The proof here follows a similar pattern than that of Theorem 9.2.1, but will be more complex. In particular, the lesser tight bound is due to the fact that we need to consider an additional bad event, which we upper bound via a different reduction from OMDL. As we explain in detail below, this reduction will make use of a looser Forking Lemma, which is a variant of the ‘‘Local Forking Lemma’’ [15], which only resamples a single random oracle output when rewinding. The extra looseness is due to needing to ensure an extra condition when rewinding.

Proof of Theorem 9.3.1. Let \mathcal{A} be the adversary described in the theorem. Denote the output message-signature pair of \mathcal{A} as $(M^*, \text{sig}^* = (R^*, z^*))$. Without loss of generality, we assume \mathcal{A} always queries RO on $H_2(\text{pk}, M^*, R^*)$ before \mathcal{A} returns and always queries RO on $H_1(\text{pk}, lr, i)$

prior to the query $\text{PSIGNO}(i, lr)$ for some i and lr . (This adds up to q_s additional RO queries, and we let $q = q_h + q_s + 1$.) Denote lr^* as the leader query such that $H_1(\text{pk}, lr^*, i)$ is the first RO query prior to the $H_2(\text{pk}, M^*, R^*)$ query for some i satisfying $\text{SVf}(\text{pk}, lr^*, sig^*) = \text{true}$. If such lr^* does not exist, lr^* is set to \perp . Denote the event E_1 as

$$\text{Vf}(\text{pk}, M^*, sig^*) \wedge (lr^* = \perp \vee \text{curSS}_L(lr^*) < t - |CS|).$$

Denote the event E_2 as

$$\text{Vf}(\text{pk}, M^*, sig^*) \wedge lr^* \neq \perp \wedge \text{curSS}_L(lr^*) \neq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}.$$

If \mathcal{A} wins the game $\text{TS-SUF-3}_{\text{FROST1}}$ and $lr^* \neq \perp$, we know either $\text{curSS}_L(lr^*) < t - |CS|$ or $\text{curSS}_L(lr^*) \neq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$. Therefore, if \mathcal{A} wins the game $\text{TS-SUF-3}_{\text{FROST1}}$, then either E_1 or E_2 occurs, which implies

$$\text{Adv}_{\text{FROST1}[\text{GGen}]}^{\text{ts-suf-3}}(\mathcal{A}, \kappa) \leq \Pr[E_1] + \Pr[E_2] \leq 2 \max\{\Pr[E_1], \Pr[E_2]\}.$$

Thus, we conclude the theorem with the following two lemmas.

Lemma 9.3.2. *There exists an OMDL adversary \mathcal{B} making at most $2q_s + t$ queries to CHAL such that*

$$\Pr[E_1] \leq \sqrt{q \cdot (\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) + 3q^2(n+1)^2 \cdot 2^{-\kappa})},$$

Moreover, \mathcal{B} runs in time roughly equal two times that of \mathcal{A} , plus the time to perform at most $6n \cdot q + 4q_s + 2n^2$ exponentiations and group operations.

Lemma 9.3.3. *There exists an OMDL adversary \mathcal{B} making at most $2q_s$ queries to CHAL such that*

$$\Pr[E_2] \leq n \cdot q \sqrt{2(\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) + 2^{-\kappa})}.$$

Moreover, \mathcal{B} runs in time roughly equal two times that of \mathcal{A} , plus the time to perform at most $6n \cdot q + 4q_s + 2n^2$ exponentiations and group operations.

This completes the proof of the theorem, subject to proofs of the lemmas that we discuss next. □

The proof of Lemma 9.3.2 is almost the same as Lemma 9.2.2, so we omit the full proof. The only difference is that \mathcal{C} takes as input $h_1, \dots, h_{(n+1)q}$ in order to simulate all RO queries. For a RO query $H_1(\text{pk}, lr, i)$, \mathcal{C} first enumerates all $i' \in [n]$ and assigns $h_{(\text{ctr}_h - 1)(n+1) + i'}$ to $H_1(\text{pk}, lr, i')$.

Then, \mathcal{C} computes the nonce R for lr and assigns $h_{\text{ctr}_h(n+1)}$ to $H_2(\text{pk}, lr.\text{msg}, R)$ if it is not assigned any value yet. Similarly, for a new RO query $H_1(\text{pk}, M, R)$, its value is set to $h_{\text{ctr}_h(n+1)}$. The rest follows by similar analysis.

To prove Lemma 9.3.3, we need the following variant of the forking lemma, which extends the local forking lemma of [15]. The difference is that forking is happening on two indices I, J (leading to I', J' in the forking) while in [15] there is a single I (and corresponding I' in the forking). The difference between a local forking ([15] and Lemma 9) versus classical ([17] and Lemma 5) is that in the former only one point is resampled in forking while in the latter it is all points following the fork.

Lemma 9.3.4. *Let $q \geq 1$ be an integer and H and Q be two sets. Let \mathcal{A} be a randomized algorithm that on input x, h_1, \dots, h_q outputs a tuple (I, J, Out) , where $I \in \{\perp\} \cup [1..q]$, $J \in Q$, and Out is a side output. Let IG be a randomized algorithm that generates x . The accepting probability of \mathcal{A} is defined as*

$$\text{acc}(\mathcal{A}) := \Pr_{x \leftarrow \text{IG}, h_1, \dots, h_q \leftarrow H} [(I, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q) : I \neq \perp].$$

Consider algorithm $\text{Fork}_2^{\mathcal{A}}$ described in Figure 9.5. The accepting probability of $\text{Fork}_2^{\mathcal{A}}$ is defined as

$$\text{acc}(\text{Fork}_2^{\mathcal{A}}) := \Pr_{x \leftarrow \text{IG}} [\alpha \leftarrow \text{Fork}_2^{\mathcal{A}}(x) : \alpha \neq \perp].$$

Then, $\text{acc}(\text{Fork}_2^{\mathcal{A}}) \geq \text{acc}(\mathcal{A})^2 / (q \cdot |Q|)$.

Proof. Denote $\hat{h}_i = (h_1, \dots, h_{i-1}, h_{i+1}, \dots, h_q) \in H^{q-1}$. For any $i \in [1..q]$, $j \in Q$, $\hat{h}_i \in H^{q-1}$, and input x , define

$$Y_{i,j}(x, \hat{h}_i) := \Pr_{h_i \leftarrow H} [I = i, J = j : (I, J, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q)].$$

Then, we have

$$\begin{aligned} \text{acc}(\mathcal{A}) &= \sum_{i=1}^q \sum_{j \in Q} \Pr_{x \leftarrow \text{IG}, h_1, \dots, h_q \leftarrow H} [I = i, J = j : (I, J, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q)] \\ &= \sum_{i=1}^q \sum_{j \in Q} \mathbb{E}_{x \leftarrow \text{IG}, \hat{h}_i \leftarrow H^{q-1}} [Y_{i,j}(x, \hat{h}_i)]. \end{aligned}$$

Algorithm Fork₂^A(x) :
 Pick the random coin ρ of \mathcal{A} at random
 $h_1, \dots, h_q \leftarrow H$
 $(I, J, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q; \rho)$
 If $I = \perp$ then return \perp
 $h'_I \leftarrow H$
 $(I', J', \text{Out}') \leftarrow \mathcal{A}(x, h_1, \dots, h_{I-1}, h'_I, h_{I+1}, \dots, h_q; \rho)$
 If $I \neq I'$ or $J \neq J'$ then return \perp
 Return $(I, J, \text{Out}, \text{Out}')$

Figure 9.5: The forking algorithm build from \mathcal{A} used in Lemma 9.3.4.

Thus, we have

$$\text{acc}(\text{Fork}^{\mathcal{A}}) = \sum_{i=1}^q \sum_{j \in Q} \Pr_{x \leftarrow \text{IG}, h_1, \dots, h_n, h'_i \leftarrow \mathfrak{H}} [I = I' = i, J = J' = j : \quad (9.7)$$

$$(I, J, \text{Out}) \leftarrow \mathcal{A}(x, h_1, \dots, h_q),$$

$$(I', J', \text{Out}') \leftarrow \mathcal{A}(x, h_1, \dots, h_{i-1}, h'_i, h_{i+1}, h_q)]$$

$$= \sum_{i=1}^q \sum_{j \in Q} \mathbb{E}_{x \leftarrow \text{IG}, \hat{h}_i \leftarrow \mathfrak{H}^{q-1}} [Y_{i,j}(x, \hat{h}_i)^2] \quad (9.8)$$

$$\geq \sum_{i=1}^q \sum_{j \in Q} \left(\mathbb{E}_{x \leftarrow \text{IG}, \hat{h}_i \leftarrow \mathfrak{H}^{q-1}} [Y_{i,j}(x, \hat{h}_i)] \right)^2 \quad (9.9)$$

$$\geq \frac{1}{q \cdot |Q|} \cdot \left(\sum_{i=1}^q \sum_{j \in Q} \mathbb{E}_{x \leftarrow \text{IG}, \hat{h}_i \leftarrow \mathfrak{H}^{q-1}} [Y_{i,j}(x, \hat{h}_i)] \right)^2 \quad (9.10)$$

$$= \frac{\text{acc}(\mathcal{A})^2}{q \cdot |Q|}, \quad (9.11)$$

where (9.9) is due to the fact that $\mathbb{E}[X^2] \geq (\mathbb{E}[X])^2$ and (9.10) is due to the fact that $\sum_{i=1}^n a_i^2 \geq \frac{1}{n} (\sum_{i=1}^n a_i)^2$. \square

Proof of Lemma 9.3.3. We first construct an algorithm \mathcal{C} following the syntax of the algorithm described in Lemma 9.3.4. The input of \mathcal{C} consists of $2q_s$ uniformly random group elements $U_1, V_1, \dots, U_{q_s}, V_{q_s} \in \mathbb{G}$ and uniformly random vectors $h_1, \dots, h_{n,q} \in (\mathbb{Z}_p)$. Similarly to the

proof of Lemma 9.2.2, \mathcal{C} can access DLOG oracle and at the beginning, initializes all the states st_0, \dots, st_n as in the game $\text{TS-SUF-3}_{\text{FROST1}}$, and initializes the counters ctr_s, ctr_h to 0 and the function dt to an empty table. \mathcal{C} also initializes ctr_{PP} to an empty table, which are used to record the counter corresponding to each token generated by honest parties. Then, \mathcal{C} runs \mathcal{A} with access to the oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}, \widetilde{\text{RO}}$, which are simulated as follows. In the following description, we use i to denote the index of parties, j to denote the index of $U_1, V_1, \dots, U_{q_s}, V_{q_s}$, and k to denote the index of $h_1, \dots, h_{n \cdot q}$.

$\widetilde{\text{INIT}}(CS)$: \mathcal{C} initializes H to an empty table and samples a_0, \dots, a_{t-1} uniformly from \mathbb{Z}_p . Define $f(x) := \sum_{i=0}^{t-1} a_i x^i$. Then, \mathcal{C} sets $pk \leftarrow g^{f(0)}$, $pk_i \leftarrow g^{f(i)}$ for $i \in [1..n]$, and $sk_i \leftarrow f(i)$ for $i \in CS$. Finally, \mathcal{C} returns $pk, aux = (pk_1, \dots, pk_n), \{sk_i\}_{i \in CS}$.

$\widetilde{\text{RO}}$ query $H_1(x)$: If $H_1(x) \neq \perp$, \mathcal{C} returns $H_1(x)$. Otherwise, \mathcal{C} parses x as $(\tilde{pk}, lr, \tilde{i})$ for some $\tilde{i} \in [1..n]$. If the parsing fails or $\tilde{pk} \neq pk$, \mathcal{C} sets $H_1(x) \leftarrow_{\$} \mathbb{Z}_p$ and returns $H_1(x)$. Otherwise, \mathcal{C} increases ctr_h by 1 and sets $H_1(pk, lr, i) \leftarrow h_{n(ctr_h-1)+i}$ for each $i \in [1..n]$. In addition, define $\text{mapLR}(ctr_h) := lr$. Then, \mathcal{C} computes $R \leftarrow \prod_{i \in lr.SS} R_i S_i^{d_i}$, where $(R_i, S_i) \leftarrow lr.PP(i)$ and $d_i = H_1(pk, lr, i)$. If $H_2(pk, lr.msg, R) = \perp$, \mathcal{C} sets $H_2(pk, lr.msg, R) \leftarrow_{\$} \mathbb{Z}_p$. Finally, \mathcal{C} returns $H_1(x)$.

$\widetilde{\text{RO}}$ query $H_2(x)$: If $H_2(x) \neq \perp$, \mathcal{C} returns $H_2(x)$. Otherwise, \mathcal{C} sets $H_2(x) \leftarrow_{\$} \mathbb{Z}_p$ and returns $H_2(x)$.

$\widetilde{\text{PPO}}(i)$ query: Same as in the game $\text{TS-SUF-3}_{\text{FROST1}}$, except in the simulation of algorithm SPP, \mathcal{C} first increases ctr_s by 1 and sets $pp \leftarrow (U_{ctr_s}, V_{ctr_s})$, $st_i.\text{mapPP}(pp) \leftarrow (0, 0)$, and $ctr_{PP}(i, pp) \leftarrow ctr_s$.

$\widetilde{\text{PSIGNO}}(i, lr)$ query: Same as in the game $\text{TS-SUF-3}_{\text{FROST1}}$, except in the simulation of algorithm PS, if $st_i.\text{mapPP}(pp) \neq \perp$, let $j \leftarrow ctr_{PP}(i, lr.PP(i))$, and \mathcal{C} computes $\alpha \in \mathbb{Z}_p^{t+2q_s}$ as

$$\alpha_k := \begin{cases} 1, & \text{for } k = 2j - 1, \\ d_i, & \text{for } k = 2j, \\ 0, & \text{o.w.} \end{cases}$$

and sets $z_i \leftarrow \text{DLOG}(\alpha) + c\lambda_i^{lr.SS} \cdot f(i)$, where $\text{DLOG}(\alpha) = \text{DLog}_{\mathbb{G},g}(U_j V_j^{d_i})$. In addition, \mathcal{C} sets $\text{dt}(j) \leftarrow (k, d_i, z_i - c\lambda_i^{lr.SS} \cdot f(i))$, where k denotes the index such that $H_1(\text{pk}, lr, i)$ is set to h_k during the simulation.

After receiving the output $(M^*, \text{sig}^* = (R^*, z^*))$ from \mathcal{A} , \mathcal{C} returns (\perp, \perp, \perp) if E_2 does not occur. Otherwise, we know $\text{curSS}_L(lr^*) > 0$ and $\text{curSS}_L(lr^*) \neq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$. Therefore, there exists k^* and i^* such that $\text{mapLR}(k^*) = lr^*$ and $i^* \in \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\} \setminus \text{curSS}_L(lr^*)$. (Since $\text{curSS}_L(lr^*) \subseteq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$, we must have $\{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\} \setminus \text{curSS}_L(lr^*) \neq \emptyset$.) Since $i^* \in \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$, there exists $j^* \in [1..q_s]$ such that $lr^*.PP(i^*) = (U_{j^*}, V_{j^*})$. If $\text{dt}(j^*) = \perp$, \mathcal{C} sets $J \leftarrow \perp$. Otherwise, let $(k, d, z) \leftarrow \text{dt}(j^*)$ and \mathcal{C} sets $J = k$. Then, \mathcal{C} returns $(n(k^* - 1) + i^*, J, \text{Out})$, where Out consists of all variables received or generated by \mathcal{C} , including i^*, j^*, k^*, lr^* .

ANALYSIS OF \mathcal{C} . To use Lemma 9.3.4, we define IG as the algorithm that samples $2q_s$ group elements uniformly from \mathbb{G} and outputs them. The output J is either \perp or in $[1..(n \cdot q)]$. It is not hard to see that \mathcal{C} simulates the game $\text{TS-SUF-3}_{\text{FROST1}}$ perfectly when all the inputs of \mathcal{C} are uniformly sampled from their domain, which implies $\text{acc}(\mathcal{C}) \geq \Pr[E_2]$, where $\Pr[E_2]$ refers to the probability in the original $\text{TS-SUF-3}_{\text{FROST1}}$ game with \mathcal{A} (as in the lemma statement). By Lemma 9.3.4,

$$\text{acc}(\text{Fork}_2^{\mathcal{C}}) \geq \frac{\Pr[E_2]^2}{n \cdot q(n \cdot q + 1)} \leq \frac{\Pr[E_2]^2}{2n^2q^2}.$$

CONSTRUCT \mathcal{B} FROM $\text{Fork}_2^{\mathcal{C}}$. We now give a construct of the OMDL adversary \mathcal{B} using $\text{Fork}_2^{\mathcal{C}}$. To start with, \mathcal{B} queries CHAL oracle $2q_s$ times to generate $U_1, V_1, \dots, U_{q_s}, V_{q_s}$ as the input of $\text{Fork}_2^{\mathcal{C}}$ and runs $\text{Fork}_2^{\mathcal{C}}$. Without loss of generality, we can assume all the OMDL challenges are different, since otherwise, \mathcal{B} can solve them trivially. All DLOG queries from $\text{Fork}_2^{\mathcal{C}}$ are relayed by \mathcal{B} to DLOG oracle of the game $\text{OMDL}_{\mathbb{G}\text{Gen}}$. Denote the event BadHash as $h_I \neq h'_I$, where I are outputted by the first execution of \mathcal{C} . Since h_I, I are independent of h'_I , we know $\Pr[\text{BadHash}] \leq 1/p$.

It is left to show that if $\text{Fork}_2^{\mathcal{C}}$ returns $(I, J, \text{Out}, \text{Out}')$ and BadHash does not occur, \mathcal{B} can win the game $\text{OMDL}_{\mathbb{G}\text{Gen}}$, which implies

$$\text{Adv}_{\mathbb{G}\text{Gen}}^{\text{omdl}}(\mathcal{B}, \kappa) \geq \text{acc}(\text{Fork}_2^{\mathcal{C}}) - \Pr[\text{BadHash}] \geq \frac{\Pr[E_2]^2}{2n^2q^2} - 1/p.$$

We directly use the notations in the description of \mathcal{C} to denote the variables in Out and use $(\cdot)'$ to denote the variables in Out' . We first show how to compute the discrete log of U_j, V_j for $j \neq j^*$. Denote u_j, v_j as the discrete log of U_j, V_j . There are the following cases.

Case 0: Both $\text{dt}(j)$ and $\text{dt}'(j)$ are \perp . In this case, \mathcal{B} computes u_j, v_j by directly querying oracle $\text{DLOG}(U_j)$ and $\text{DLOG}(V_j)$.

Case 1: Exactly one of $\text{dt}(j)$ and $\text{dt}'(j)$ is not \perp . Without loss of generality, assume $\text{dt}(j) = (k, d, z)$, which implies $g^z = U_j V_j^d$. \mathcal{B} computes v_j by directly querying oracle $\text{DLOG}(V_j)$ and computes $u_j = z - d \cdot v_j$.

For all the following cases, both $\text{dt}(j)$ and $\text{dt}'(j)$ are not \perp and we denote $(k, d, z) \leftarrow \text{dt}(j)$ and $(k', d', z') \leftarrow \text{dt}'(j)$.

Case 2: $d \neq d'$. In this case, \mathcal{B} computes $v_j = \frac{z-z'}{d-d'}$, $u_j = z - d \cdot v_j$.

Case 3: $d = d'$. In this case, \mathcal{B} computes v_j, u_j the same as Case 1. Also, since $d = d'$, \mathcal{B} queries DLOG oracle only once in order to answer queries $\text{DLOG}(U_j V_j^d)$ and $\text{DLOG}(U_j V_j^{d'})$ from $\text{Fork}^{\mathcal{C}}$.

From the execution of \mathcal{C} , we know $\text{dt}(j) = (k, d, z) \neq \perp$ if and only if \mathcal{C} queries DLOG on $(U_j V_j^d)$. Therefore, denote $\text{DLOG}(U_j V_j^d)$ as the DLOG query associated with $\text{dt}(j)$. For all the above cases, \mathcal{B} queries DLOG oracle twice for simulating DLOG queries associated with $\text{dt}(j)$ and $\text{dt}'(j)$ and computing u_j, v_j .

We now show how to compute u_{j^*} and v_{j^*} . From the execution of $\text{Fork}_2^{\mathcal{C}}$, we know $\text{pk} = \text{pk}'$ and $\text{mapLR}(k) = \text{mapLR}'(k)$ for all $k \leq I$, which implies $lr^* = \text{mapLR}(I) = \text{mapLR}'(I) = lr^{*'}$. Since E_2 occurs in both executions of \mathcal{C} , we know $\text{SVf}(\text{pk}, lr^*, (R^*, z^*)) = \text{true}$ and $\text{SVf}(\text{pk}, lr^{*'}, (R^{*'}, z^{*'})) = \text{true}$ are valid. Therefore, $g^{z^*} = R^* g^{a_0 c}$, $R^* = \sum_{i \in lr^*.SS} R_i S_i^{d_i}$, $g^{z^{*'}} = R^{*'} g^{a_0 c'}$, $R^{*'} = \sum_{i \in lr^{*}.SS} R_i S_i^{d'_i}$, where $(R_i, S_i) = lr.PP(i)$, $c = H_2(\text{pk}, M^*, R^*)$, $c' = H_2(\text{pk}, M^*, R^{*'})$, and $d_i = H_1(\text{pk}, lr^*, i)$, $d'_i = H_1(\text{pk}, lr^{*'}, i)$. Since for each $i \neq i^*$ we have $d_i = h_{n(k^*-1)+i} = d'_i$, we have $g^{z^* - z^{*'}} = \frac{R^*}{R^{*'}} g^{a_0(c-c')} = S_{i^*}^{d_{i^*} - d'_{i^*}} g^{a_0(c-c')}$. Therefore, \mathcal{C} can compute $v_{j^*} = \frac{z^* - z^{*'} - a_0(c-c')}{d_{i^*} - d'_{i^*}}$. If $J = \perp$, \mathcal{B} computes u_{j^*} by querying $\text{DLOG}(U_{j^*})$ directly. In this case, \mathcal{B} queries DLOG only once to compute u_{j^*} and v_{j^*} . If $J \neq \perp$, let $(k, d, z) \leftarrow \text{dt}(j^*) =$ and $(k', d', z') \leftarrow \text{dt}'(j^*)$. Then, \mathcal{B} computes $u_{j^*} = z - d \cdot v_{j^*}$. Since $i^* \notin S_2(lr^*)$, we know $k \neq I$. (Otherwise, suppose $k = I$. Since $I = n(k^* - 1) + i^*$ and $\text{mapLR}(k^*) = lr^*$, we know a $\text{PSIGNO}(i^*, lr^*)$ is made and does not return \perp during the simulation, which implies $i^* \in S_2(lr^*)$.) Thus, we have $k' = J = k \neq I$ and $d = h_J = d'$, which means \mathcal{B} only needs to query DLOG once to simulate the DLOG queries associated with $\text{dt}(j)$ and $\text{dt}'(j)$. Therefore, the total number of DLOG queries made by \mathcal{B} is equal to $2q_s - 1$, which implies \mathcal{B} wins the game OMDL_{Gen} . \square

9.4 Security of FROST3

We show that FROST3 is TS-UF-1-secure in the ROM under the AOMDL assumption. Formally, we show the following theorem.

Theorem 9.4.1. *Let GGen be a group generation algorithm. For any TS-UF-1 adversary \mathcal{A} making at most q_s queries to PPO and at most q_h queries to RO, there exists an AOMDL adversary \mathcal{B} making at most $2q_s + n$ queries to CHAL such that*

$$\text{Adv}_{\text{FROST3}[\text{GGen}]}^{\text{ts-uf-1}}(\mathcal{A}, \kappa) \leq \sqrt{q \cdot (\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) + 3q^2 \cdot 2^{-\kappa})},$$

where $q = q_s + q_h + 1$. Moreover, \mathcal{B} runs in time roughly equal two times that of \mathcal{A} , plus the time to perform at most $(4n + 2) \cdot q + 2q_s + 2n^2$ exponentiations and group operations.

The proof is almost identical to that of FROST2. The only difference is that in FROST3, the size of T_{dt} can be bounded only by $\text{curSS}_M(M^*)$ —which is itself bounded by $t - |CS|$ due to the winning condition of TS-UF-1—rather than by $\text{curSS}_L(lr)$. The reason is that in FROST3, multiple distinct lr 's may correspond to (M^*, R^*) , i.e., $lr.\text{msg} = M^*$ and $\text{CompPar}(\text{pk}, lr) \rightarrow (R^*, \dots)$. Therefore, T_{dt} cannot be bounded by any single lr . We will see this leads to a concrete attack for TS-UF-2-seucurity in Section 9.5 We provide the proof below for completeness.

Proof of Theorem 9.4.1. Let \mathcal{A} be an adversary as described in the theorem. Denote the output message-signature pair of \mathcal{A} as $(M^*, \text{sig}^* = (R^*, z^*))$. Without loss of generality, we assume \mathcal{A} always queries RO on $H_2(\text{pk}, M^*, R^*)$ before \mathcal{A} returns and always queries RO on $H_1(\text{pk}, lr)$ prior to the query $\text{PSIGNO}(i, lr)$ for some i and lr . (This adds up to q_s additional RO queries, and we let $q = q_h + q_s + 1$.)

We first construct an algorithm \mathcal{C} compatible with the syntax in Lemma 9.2.3. The input of \mathcal{C} consists of $(2q_s + t)$ uniformly random group elements $A_0, \dots, A_{t-1}, U_1, V_1, \dots, U_{q_s}, V_{q_s} \in \mathbb{G}$ and uniformly random integers $h_1, \dots, h_{2q} \in \mathbb{Z}_p$. Also, \mathcal{C} can access an oracle DLOG, which on input $\alpha \in \mathbb{Z}_p^{2q_s+t}$ outputs $\text{DLog}_{\mathbb{G},g}(A_0^{\alpha_1} \dots A_{t-1}^{\alpha_t} U_1^{\alpha_{t+1}} V_1^{\alpha_{t+2}} \dots U_{q_s}^{\alpha_{t+2q_s-1}} V_{q_s}^{\alpha_{t+2q_s}})$. (We can think of this oracle as part of \mathcal{C} in the context of the Forking Lemma, as \mathcal{C} does not need to be efficient.) To start with, \mathcal{C} initializes all the states $\text{st}_0, \dots, \text{st}_n$. In addition, it initializes counters $\text{ctr}_s, \text{ctr}_h$ to 0 and a function dt to an empty table, which are used to record the DLOG query related to each (U_j, V_j) . \mathcal{C} also initializes $\text{curLR} \leftarrow \emptyset$ to record all leader requests that appears during the

game and initializes ctrPP to an empty table, which are used to record the counter corresponding to each token generated by honest parties. We also use a flag BadPPO to denote whether a bad event occurs, which are initially set to false . Then, \mathcal{C} runs \mathcal{A} with access to the oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}, \widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$: \mathcal{C} initializes H to an empty table and sets $\text{pk} \leftarrow A_0$, $\text{pk}_i = \prod_{j=0}^{t-1} A_j^{i^j}$ for $i \in [1..n]$, and $\text{sk}_i = \text{DLOG}(\text{pk}_i)$ for $i \in CS$. Finally, \mathcal{C} returns $\text{pk,aux} = (\text{pk}_1, \dots, \text{pk}_n), \{\text{sk}_i\}_{i \in CS}$.

$\widetilde{\text{RO}}$ query $H_1(x)$: If $H_1(x) \neq \perp$, \mathcal{C} returns $H_1(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, lr)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_1(x) \leftarrow_s \mathbb{Z}_p$ and returns $H_1(x)$. Otherwise, \mathcal{C} increases ctr_h by 1, sets $H_1(x) \leftarrow h_{2\text{ctr}_h-1}$, and adds lr to curLR . Also, \mathcal{C} computes $R \leftarrow \prod_{i \in lr.SS} R_i S_i^{h_{2\text{ctr}_h-1}}$, where $(R_i, S_i) \leftarrow lr.PP(i)$. If $H_2(\text{pk}, lr.\text{msg}, R) = \perp$, \mathcal{C} sets $H_2(\text{pk}, lr.\text{msg}, R) = h_{2\text{ctr}_h}$. In addition, define $\text{mapLR}(\text{ctr}_h) := lr$ and set $\text{curLR} \leftarrow \text{curLR} \cup \{lr\}$. Finally, \mathcal{C} returns $H_1(x)$.

$\widetilde{\text{RO}}$ query $H_2(x)$: If $H_2(x) \neq \perp$, \mathcal{C} returns $H_2(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, M, R)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_2(x) \leftarrow_s \mathbb{Z}_p$ and returns $H_2(x)$. Otherwise, \mathcal{C} increases ctr_h by 1 and sets $H_2(x) \leftarrow h_{2\text{ctr}_h}$. Finally, \mathcal{C} returns $H_2(x)$.

$\widetilde{\text{PPO}}(i)$ query: Same as in the game $\text{TS-SUF-2}_{\text{FROST}_2}$, except in the simulation of algorithm SPP , \mathcal{C} first increases ctr_s by 1 and sets $pp \leftarrow (U_{\text{ctr}_s}, V_{\text{ctr}_s})$, $\text{st}_i.\text{mapPP}(pp) \leftarrow (0, 0)$, and $\text{ctrPP}(i, pp) \leftarrow \text{ctr}_s$. In addition, BadPPO is set to true if there exists $lr \in \text{curLR}$ such that $lr.PP(i) = (U_{\text{ctr}_s}, V_{\text{ctr}_s})$.

$\widetilde{\text{PSIGNO}}(i, lr)$ query: Same as in the game $\text{TS-SUF-2}_{\text{FROST}_2}$, except in the simulation of algorithm PS , if $\text{st}_i.\text{mapPP}(pp) \neq \perp$, let $j \leftarrow \text{ctrPP}(i, lr.PP(i))$, and \mathcal{C} computes $\alpha \in \mathbb{Z}_p^{t+2q_s}$ as

$$\alpha_k := \begin{cases} c \cdot \lambda_k^{lr.SS} \cdot i^{k-1}, & \text{for } k \in [t], \\ 1, & \text{for } k = t + 2j - 1, \\ d_i, & \text{for } k = t + 2j, \\ 0, & \text{o.w.} \end{cases}$$

and sets $z_i \leftarrow \text{DLOG}(\alpha)$, which equals to $\text{DLOG}_{\mathbb{G},g} \left(U_j V_j^{d_i} \text{pk}_i^{c\lambda_i^{lr,SS}} \right)$. In addition, \mathcal{C} sets $\text{dt}(j) \leftarrow (i, k, d_i, c\lambda_i^{lr,SS}, z_i)$, where k denotes the index such that $\text{H}_1(\text{pk}, lr)$ is set to h_{2k-1} during the simulation.

After receiving the output $(M^*, \text{sig}^* = (R^*, z^*))$ from \mathcal{A} , \mathcal{C} returns \perp if $\text{BadPPO} = \text{true}$ or E_1 does not occur. Otherwise, \mathcal{C} finds the index I such that $\text{H}_2(\text{pk}, M^*, R^*)$ is set to h_I during the simulation. By our assumption of \mathcal{A} , we know such I must exist. Then, \mathcal{C} returns (I, Out) , where Out consists of all variables received or generated by \mathcal{C} .

ANALYSIS OF \mathcal{C} . To use Lemma 9.2.3, we define $S := \{2k\}_{k \in [1..q]}$ and IG as the algorithm that samples $2q_s + t$ group elements uniformly from \mathbb{G} and outputs them. From the simulation, we know the output index I of \mathcal{C} is always in S . Also, it is clear that \mathcal{C} simulates the game $\text{TS-UF-1}_{\text{FROST}_3}$ perfectly when all the inputs of \mathcal{C} are uniformly sampled from their domain, which implies $\text{acc}(\mathcal{C}) \geq \text{Adv}_{\text{FROST}_3}^{\text{ts-uf-1}}(\mathcal{A}, \kappa) - \Pr[\text{BadPPO}]$, where $\Pr[E_1]$ refers to the probability in the original $\text{TS-SUF-2}_{\text{FROST}_2}$ game with \mathcal{A} (as in the lemma statement), whereas $\Pr[\text{BadPPO}]$ is the probability that $\text{BadPPO} = \text{true}$ at the end of \mathcal{C} 's execution. Since every pair U_j, V_j is sampled uniformly from \mathbb{G} , for each $\text{PPO}(i)$ query, the probability BadPPO is set to true is less than $|\text{curLR}|/|G| \leq q_h/p$. Therefore, we have $\Pr[\text{BadPPO}] \leq q_s q_h/p$. By Lemma 9.2.3,

$$\text{acc}(\text{Fork}^{\mathcal{C}}) \geq \text{Adv}_{\text{FROST}_3}^{\text{ts-uf-1}}(\mathcal{A}, \kappa)^2/q - 2q_s/p.$$

CONSTRUCT \mathcal{B} FROM $\text{Fork}^{\mathcal{C}}$. We now give a construct of the OMDL adversary \mathcal{B} using $\text{Fork}^{\mathcal{C}}$, and the available DLOG oracle. To start with, \mathcal{B} queries CHAL oracle $2q_s + t$ times to generate $A_0, \dots, A_{t-1}, U_1, V_1, \dots, U_{q_s}, V_{q_s}$ as the input of $\text{Fork}^{\mathcal{C}}$ and runs $\text{Fork}^{\mathcal{C}}$. Without loss of generality, we can assume all the AOMDL challenges are different, since otherwise, \mathcal{B} can solve them trivially. All DLOG queries from $\text{Fork}^{\mathcal{C}}$ are relayed by \mathcal{B} to DLOG oracle of the game $\text{AOMDL}_{\text{GGen}}$. Denote the event BadHash as any two of the scalars $h_1, h'_1, \dots, h_q, h'_q$ generated in the execution of $\text{Fork}^{\mathcal{C}}$ are same. Since $h_1, h'_1, \dots, h_q, h'_q$ are sampled uniformly from \mathbb{Z}_p , we know $\Pr[\text{BadHash}] \leq 2q^2/p$.

It is left to show that if $\text{Fork}^{\mathcal{C}}$ returns $(I, \text{Out}, \text{Out}')$ and BadHash does not occur, \mathcal{B} can win the game $\text{AOMDL}_{\text{GGen}}$, which implies

$$\text{Adv}_{\text{GGen}}^{\text{aomdl}}(\mathcal{B}, \kappa) \geq \text{acc}(\text{Fork}^{\mathcal{C}}) - \Pr[\text{BadHash}] \geq \text{Adv}_{\text{FROST}_3}^{\text{ts-uf-1}}(\mathcal{A}, \kappa)^2/q - 3q^2/p.$$

We directly use the notations in the description of \mathcal{C} to denote the variables in Out and use $(\cdot)'$ to denote the variables in Out' . By the execution of $\text{Fork}^{\mathcal{C}}$, we know $(\text{pk}, M^*, R^*) = (\text{pk}', M^{*'}, R^{*'})$

and $\text{pk} = A_0$. Since $I \in S$, let $k^* = I/2$. It is not hard to see that $\text{mapLR}(k^*) = lr^*$. (If $\text{mapLR}(k^*) = \perp$, lr^* is also \perp .)

We first show how to compute the discrete log of A_0, \dots, A_{t-1} . Denote the discrete log of A_0, \dots, A_{t-1} as a_0, \dots, a_{t-1} and define a polynomial $f(x) := \sum_{i=0}^{t-1} a_i x^i$. Since BadHash does not occur, we have $H_2(\text{pk}, M^*, R^*) = h_I \neq h'_I = H'_2(\text{pk}, M^*, R^*)$. Since $g^{z^*} = R^x A_0^{h_I}$, $g^{z^{*'}} = R^x A_0^{h'_I}$, \mathcal{B} computes $f(0) = a_0 = \frac{z^* - z^{*'}}{h_I - h'_I}$. Define $T_{\text{dt}} := \{j : (i, k, d, c, z) \leftarrow \text{dt}(j), k = k^*\}$. For each $j \in T_{\text{dt}} \cap T_{\text{dt}'}$, let $(i, k, d, c, z) \leftarrow \text{dt}(j)$ and $(i', k', d', c', z') \leftarrow \text{dt}'(j)$, and we have $g^z = U_j V_j^d \text{pk}_i^c$, $g^{z'} = U_j V_j^{d'} \text{pk}_{i'}^{c'}$. Since $\text{BadPPO} = \text{false}$ during both execution of \mathcal{C} , we know (U_j, V_j) is returned by a query $\text{PPO}(i)$ prior to the query $H_2(\text{pk}, M^*, R^*)$ during the first execution of \mathcal{C} . Since the two executions of \mathcal{C} are exactly the same prior to the query $H_2(\text{pk}, M^*, R^*)$, we know $i' = i$. Also, we know $d = h_k = h_{k^*} = h_{k'} = d'$. Therefore, \mathcal{B} can compute $f(i) = \text{DLog}_{\mathbb{G}, g}(\text{pk}_i) = \frac{z - z'}{c - c'}$. Denote $D := \{i\}_{j \in T_{\text{dt}} \cap T_{\text{dt}'}, (i, k, d, c, z) \leftarrow \text{dt}(j)}$. Since E_1 occurs in the first execution of \mathcal{C} , we know $|T_{\text{dt}}| = |\text{curSSM}(M^*)| < t - |CS|$. Therefore, we know $|D| = |T_{\text{dt}} \cap T_{\text{dt}'}| < t - |CS|$. Therefore, \mathcal{B} can pick an arbitrary set $D' \in HS \setminus D$ with size $(t - |CS| - |T_{\text{dt}} \cap T_{\text{dt}'}| - 1)$ and for each $i \in D'$, \mathcal{B} queries DLOG oracle on pk_i . Therefore, \mathcal{B} knows the value of $f(i)$ for $i \in CS \cup D \cup D' \cup \{0\}$. Since $|CS \cup D \cup D' \cup \{0\}| = t$, \mathcal{B} can compute the value of a_0, \dots, a_{t-1} using Lagrange interpolation.

We now show how to compute the discrete log of $U_1, V_1, \dots, U_{q_s}, V_{q_s}$. Denote their discrete log as $u_1, v_1, \dots, u_{q_s}, v_{q_s}$. From the execution of \mathcal{C} , we know $\text{dt}(j) = (i, k, d, c, z) \neq \perp$ if and only if \mathcal{C} queries DLOG on $U_j V_j^d \text{pk}_i^c$. Therefore, denote $\text{DLOG}(U_j V_j^d \text{pk}_i^c)$ as the DLOG query associated with $\text{dt}(j)$. For each $j \in q_s$, there are the following cases.

Case 0: Both $\text{dt}(j)$ and $\text{dt}'(j)$ are \perp . In this case, \mathcal{B} computes u_j, v_j by directly querying oracle $\text{DLOG}(U_j)$ and $\text{DLOG}(V_j)$.

Case 1: Exactly one of $\text{dt}(j)$ and $\text{dt}'(j)$ is not \perp . Without loss of generality, assume $\text{dt}(j) = (i, k, d, c, z)$, which implies $g^z = U_j V_j^d \text{pk}_i^c$. \mathcal{B} computes v_j by directly querying oracle $\text{DLOG}(V_j)$ and computes $u_j = z - d \cdot v_j - c \cdot f(i)$.

For all the following cases, both $\text{dt}(j)$ and $\text{dt}'(j)$ are not \perp and we denote $(i, k, d, c, z) \leftarrow \text{dt}(j)$ and $(i', k', d', c', z') \leftarrow \text{dt}'(j)$.

Case 2: $k \neq k'$ or $k = k' > k^*$. In this case, we know $d = h_k \neq h_{k'} = d'$ and $g^z = U_j V_j^d \text{pk}_i^c$, $g^{z'} = U_j V_j^{d'} \text{pk}_{i'}^{c'}$. Therefore, \mathcal{B} computes $v_j = \frac{z - c \cdot f(i) - z' + c' \cdot f(i')}{d - d'}$, $u_j = z - d \cdot v_j - c \cdot f(i)$.

Case 3: $k = k' = k^*$. In this case, \mathcal{B} computes v_j, u_j the same as Case 1.

Case 4: $k = k' < k^*$. \mathcal{B} computes v_j, u_j the same as Case 1. Also, in this case, we have $d = d'$ and $c = c'$. Therefore, \mathcal{B} queries DLOG oracle once in order to simulate the DLOG queries associated with $\text{dt}(j)$ and $\text{dt}'(j)$.

We now count the number of DLOG queries made by \mathcal{B} .

- \mathcal{B} queries DLOG oracle $|CS|$ times queries for simulating query $\text{DLOG}(\text{pk}_i)$ made by \mathcal{C} for each $i \in CS$.
- \mathcal{B} queries DLOG oracle $|D'|$ times queries for computing a_0, \dots, a_{t-1} .
- For each $j \in \mathfrak{q}_s$, \mathcal{B} queries DLOG twice for simulating query associated with $\text{dt}(j)$ and $\text{dt}'(j)$ and computing u_j, v_j in case 0, 1, 2, 4 and queries 3 times in case 3.

Since the condition of case 3 is equivalent to $j \in T_{\text{dt}} \cap T_{\text{dt}'}$, the total number of DLOG queries made by \mathcal{B} is equal to $2\mathfrak{q}_s + |T_{\text{dt}} \cap T_{\text{dt}'}| + |CS| + |D'| = 2\mathfrak{q}_s + t - 1$. Therefore, \mathcal{B} wins the game $\text{AOMDL}_{\text{GGen}}$. \square

9.5 Attacks for the FROST Schemes

FROST1 IS NOT TS-UF-4 SECURE. Consider the setting where $n = 20$ and $t = 3$ and the adversary \mathcal{A} for the game $\text{TS-UF-4}_{\text{FROST1}}$ described in Figure 9.6. We now show that $\text{Adv}_{\text{FROST1}}^{\text{ts-uf-4}}(\mathcal{A}, \kappa) = 1$. From the execution of **PSIGNO**, we know $g^{z_1} = R_1 S_1^{d_{11}} \text{pk}_{11}^{\lambda_{11}^{\{11,15,20\}} \cdot c}$. The key observation here is that $\lambda_{11}^{\{11,15,20\}} = \frac{15 \cdot 20}{(15-11)(20-11)} = \frac{25}{3} = \frac{5 \cdot 10}{(5-11)(10-11)} = \lambda_{11}^{\{5,10,11\}}$. Therefore,

$$\begin{aligned} g^z &= R_1 S_1^{d_{11}} g^{r_2+r_3+s_2 \cdot d_{15}+s_3 \cdot d_{20}} \text{pk}_{11}^{\lambda_{11}^{\{11,15,20\}} \cdot c} \text{pk}_5^{\lambda_5^{\{5,10,11\}} \cdot c} \text{pk}_{10}^{\lambda_{10}^{\{5,10,11\}} \cdot c} \\ &= R g^{c \cdot \sum_{i \in \{5,10,11\}} \lambda_i^{\{5,10,11\}} \cdot \text{sk}_i} = R \cdot \text{pk}^c, \end{aligned}$$

which implies $(M, (R, z))$ is valid for pk . Also, it is clear that $\text{curSS}_L(lr) = \{11\}$ and $HS \cap lr.SS = \{11, 15, 20\}$, which implies the condition $\text{tf}_4(lr)$ does not hold. Therefore, \mathcal{A} wins the game $\text{TS-UF-4}_{\text{FROST1}}$ with probability 1.

The reason why the attack is possible for FROST1 is because the honest server 11 replies to the leader request lr with tokens $lr.PP(15)$ and $lr.PP(20)$ not generated by the honest servers 15 and

Adversary $\mathcal{A}^{\text{INIT,PPO,PSIGNO,RO}}$:

$CS \leftarrow \{5, 10\}$; $(\text{pk}, \text{aux}, \{\text{sk}_5, \text{sk}_{10}\}) \leftarrow_{\$} \text{INIT}(CS)$

$(R_1, S_1) \leftarrow_{\$} \text{PPO}(11)$; $s_2, r_2, s_3, r_3 \leftarrow_{\$} \mathbb{Z}_p$

$lr.\text{msg} \leftarrow M$; $lr.\text{SS} \leftarrow \{11, 15, 20\}$

$lr.\text{PP}(11) \leftarrow (R_1, S_1)$; $lr.\text{PP}(15) \leftarrow (g^{r_2}, g^{s_2})$; $lr.\text{PP}(20) \leftarrow (g^{r_3}, g^{s_3})$

$z_1 \leftarrow \text{PSIGNO}(11, lr)$

For $i \in \{11, 15, 20\}$ do $d_i \leftarrow \text{RO}(1, \text{pk}, lr, i)$

$R \leftarrow R_1 S_1^{d_{11}} g^{r_2 + r_3 + s_2 \cdot d_{15} + s_3 \cdot d_{20}}$; $c \leftarrow \text{RO}(2, \text{pk}, R, M)$

$z \leftarrow z_1 + r_2 + r_3 + s_2 \cdot d_{15} + s_3 \cdot d_{20} + c(\lambda_5^{\{5,10,11\}} \cdot \text{sk}_5 + \lambda_{10}^{\{5,10,11\}} \cdot \text{sk}_{10})$

Return $(M, (R, z))$

Figure 9.6: Adversary \mathcal{A} that wins the game $\text{TS-UF-4}_{\text{FROST1}}$, where M is a fixed message.

20 but by the adversary instead. Therefore, the attack is prevented by the general transformation from TS-SUF-3 security to TS-SUF-4 security described in Figure 8.6 since after the transformation an honest server replies to a leader request only when all the tokens within the request are authenticated by the corresponding servers, and it is not possible for the adversary to generate authenticated tokens on behalf of honest servers.

FROST2 IS NOT TS-UF-3 SECURE. Consider the setting where $n = 4$ and $t = 3$ and the adversary \mathcal{A} for the game $\text{TS-UF-3}_{\text{FROST2}}$ described in Figure 9.7. We now show that $\text{Adv}_{\text{FROST2}}^{\text{ts-uf-3}}(\mathcal{A}, \kappa) = 1$. From the execution of PSIGNO , we know $g^{z_1} = R_1 S_1^d \text{pk}_1^{\lambda_1^{\{1,2,3\}} \cdot c}$. Therefore,

$$\begin{aligned} g^z &= R_1^\gamma S_1^{d \cdot \gamma} \text{pk}_1^{\gamma \cdot \lambda_1^{\{1,2,3\}} \cdot c} \text{pk}_3^{\lambda_3^{\{1,3,4\}} \cdot c} \text{pk}_4^{\lambda_4^{\{1,3,4\}} \cdot c} \\ &= R g^{c \cdot \sum_{i \in \{1,3,4\}} \lambda_i^{\{1,3,4\}} \cdot \text{sk}_i} = R \cdot \text{pk}^c, \end{aligned}$$

which implies $(M, (R, z))$ is valid for pk . Also, it is clear that $\text{curSS}_L(lr) = \{1\}$ and $\{i \in HS \cap lr.\text{SS} : lr.\text{PP}(i) \in \text{PP}_i\} = \{1, 2\}$, which implies the condition $\text{tf}_3(lr)$ does not hold. Therefore, \mathcal{A} wins the game $\text{TS-UF-3}_{\text{FROST2}}$ with probability 1.

FROST3 IS NOT TS-UF-2 SECURE. Consider the setting where $n = 2$ and $t = 2$ and the adversary \mathcal{A} for the game $\text{TS-UF-2}_{\text{FROST3}}$ described in Figure 9.8. We now show that $\text{Adv}_{\text{FROST3}}^{\text{ts-uf-2}}(\mathcal{A}, \kappa) = 1$. From the execution of PSIGNO , we know $g^{z_1} = R_1 S_1^d \text{pk}_1^{\lambda_1^{\{1,2\}} \cdot c}$ and $g^{z_2} = R_2 S_2^d \text{pk}_2^{\lambda_2^{\{1,2\}} \cdot c}$.

Adversary $\mathcal{A}^{\text{INIT,PPO,PSIGNO,RO}}$:

$CS \leftarrow \{3, 4\}$; $(\text{pk}, \text{aux}, \{\text{sk}_3, \text{sk}_4\}) \leftarrow \text{INIT}(CS)$
 $(R_1, S_1) \leftarrow \text{PPO}(1)$; $(R_2, S_2) \leftarrow \text{PPO}(2)$; $\gamma \leftarrow \lambda_1^{\{1,3,4\}} / \lambda_1^{\{1,2,3\}}$
 $lr.\text{msg} \leftarrow M$; $lr.\text{SS} \leftarrow \{1, 2, 3\}$
 $lr.\text{PP}(1) \leftarrow (R_1, S_1)$; $lr.\text{PP}(2) \leftarrow (R_2, S_2)$
 $lr.\text{PP}(3) \leftarrow (R_1^{\gamma^{-1}} R_2^{-1}, S_1^{\gamma^{-1}} S_2^{-1})$
 $z_1 \leftarrow \text{PSIGNO}(1, lr)$
 $d \leftarrow \text{RO}(1, \text{pk}, lr)$; $R \leftarrow R_1^\gamma S_1^{\gamma \cdot d}$; $c \leftarrow \text{RO}(2, \text{pk}, R, M)$
 $z \leftarrow \gamma \cdot z_1 + c(\lambda_3^{\{1,3,4\}} \cdot \text{sk}_3 + \lambda_4^{\{1,3,4\}} \cdot \text{sk}_4)$
Return $(M, (R, z))$

Figure 9.7: Adversary \mathcal{A} that wins the game $\text{TS-UF-3}_{\text{FROST}_2}$, where M is a fixed message.

Therefore,

$$g^z = R_1 S_1^d \text{pk}_1^{\lambda_1^{\{1,2\}} \cdot c} R_2 S_2^d \text{pk}_2^{\lambda_2^{\{1,2\}} \cdot c} = R_1 R_2 (S_1 S_2)^d (\text{pk}_1^{\lambda_1^{\{1,2\}}} \text{pk}_2^{\lambda_2^{\{1,2\}}})^c = R \cdot \text{pk}^c ,$$

which implies $(M, (R, z))$ is valid for pk . Also, it is clear that $\text{curSS}_L(lr_1) = \{1\}$ and $\text{curSS}_L(lr_2) = \{2\}$, which implies the condition $\text{tf}_2(lr)$ does not hold. Therefore, \mathcal{A} wins the game $\text{TS-UF-2}_{\text{FROST}_3}$ with probability 1.

Adversary $\mathcal{A}^{\text{INIT}, \text{PPO}, \text{PSIGNO}, \text{RO}}$:

$CS \leftarrow \emptyset ; (\text{pk}, \text{aux}, \perp) \leftarrow_{\$} \text{INIT}(CS)$

$(R_1, S_1) \leftarrow_{\$} \text{PPO}(1) ; (R_2, S_2) \leftarrow_{\$} \text{PPO}(2)$

$lr_1.\text{msg} \leftarrow M ; lr_1.\text{SS} \leftarrow \{1, 2\} ; lr_1.\text{aggPP} \leftarrow (R_1 R_2, S_1 S_2)$

$lr_1.\text{PP}(1) \leftarrow (R_1, S_1) ; lr_1.\text{PP}(2) \leftarrow_{\$} (gR_2, gS_2)$

$lr_2.\text{msg} \leftarrow M ; lr_2.\text{SS} \leftarrow \{1, 2\} ; lr_2.\text{aggPP} \leftarrow (R_1 R_2, S_1 S_2)$

$lr_2.\text{PP}(1) \leftarrow (gR_1, gS_1) ; lr_2.\text{PP}(2) \leftarrow_{\$} (R_2, S_2)$

$z_1 \leftarrow \text{PSIGNO}(1, lr_1) ; z_2 \leftarrow \text{PSIGNO}(2, lr_2)$

$d \leftarrow \text{RO}(1, \text{pk}, lr.\text{SS}, lr.\text{msg}, lr.\text{aggPP}) ; R \leftarrow R_1 R_2 (S_1 S_2)^d ; c \leftarrow \text{RO}(2, \text{pk}, R, M)$

$z \leftarrow z_1 + z_2$

Return $(M, (R, z))$

Figure 9.8: Adversary \mathcal{A} that wins the game $\text{TS-UF-2}_{\text{FROST}_3}$, where M is a fixed message.

Chapter 10

PARTIALLY NON-INTERACTIVE THRESHOLD SIGNATURES FROM LINEAR HASH FUNCTIONS

In this chapter, we first give the definition of linear hash functions and algebraic one-more preimage resistance (AOMPR) of a linear hash function family. Then, we show that AOMPR is implied by collision resistance of the linear hash function family, and that such linear hash function family can be instantiated from pairing-free groups based on the DL assumption. Finally, we give new threshold signature constructions based on linear hash functions.

10.1 Linear Hash Functions

The notion of linear hash functions is introduced in [80, 81], which is in turn adapted from [10]. We adapt the definition from [80] by additionally requiring the scalar set \mathcal{S} to be a field and \mathcal{D} and \mathcal{R} to be \mathcal{S} -modules, which is necessary for the reduction from collision resistance to AOMPR and for our constructions in Section 10.4 to work.

Definition 10.1.1. *A linear hash function family LHF is a pair of algorithms (PGen, F) such that*

- a) *PGen is a randomized algorithm that takes as input the security parameter 1^κ and returns the system parameter par that defines three sets $\mathcal{S} = \mathcal{S}(\text{par})$, $\mathcal{D} = \mathcal{D}(\text{par})$ and $\mathcal{R} = \mathcal{R}(\text{par})$, where \mathcal{S} is a field, and \mathcal{D} and \mathcal{R} are \mathcal{S} -modules. Moreover, we require $|\mathcal{S}| \geq 2^\kappa$, $|\mathcal{D}| \geq 2^\kappa$, and $|\mathcal{R}| \geq 2^\kappa$.*
- b) *F is a deterministic function that takes as input the system parameter par and an element $x \in \mathcal{D}$ and returns an element in \mathcal{R} such that $\text{F}(\text{par}, \cdot) : \mathcal{D} \rightarrow \mathcal{R}$ is a epimorphism of \mathcal{S} -modules. Moreover, F is not a monomorphism, which is equivalent to there exists $z^* \in \mathcal{D}$ such that $z^* \neq 0$ and $\text{F}(\text{par}, z^*) = 0$. For simplicity, we omit par from the input of F from now on.*

Game $\text{CR}_{\text{LHF}}^{\mathcal{A}}(\kappa)$:
 $par \leftarrow \text{PGen}(1^\kappa)$
 $(x_1, x_2) \leftarrow_{\$} \mathcal{A}(par)$
 If $x_1 \neq x_2$ and $F(x_1) = F(x_2)$ then
 Return 1
 Return 0

Figure 10.1: The CR security game for a linear hash family $\text{LHF} = (\text{PGen}, F)$.

COLLISION RESISTANCE. Collision resistance of linear hash functions is analogous to collision resistance of cryptographic hash functions, which ensures that it is hard to find two distinct inputs that map to the same output. The $\text{CR}_{\text{LHF}}^{\mathcal{A}}$ game is defined in Figure 10.1. The corresponding advantage of \mathcal{A} is defined as $\text{Adv}_{\text{LHF}}^{\text{cr}}(\mathcal{A}, \kappa) := \Pr[\text{CR}_{\text{LHF}}^{\mathcal{A}} = 1]$.

10.2 Algebraic One-more Preimage Resistance

We introduce the notion of algebraic one-more preimage resistance (AOMPR) for linear hash functions, which is formally defined via the game $\text{AOMPR}_{\text{LHF}}^{\mathcal{A}}$, as described in Figure 10.2. It guarantees that any adversary given a description of a linear hash function $(\mathcal{S}, \mathcal{D}, \mathcal{R}, F)$ cannot invert $q + 1$ challenges X_1, \dots, X_{q+1} , where $X_i = F(x_i)$ for $x_i \leftarrow_{\$} \mathcal{D}$, by making at most q queries to the PI oracle that, on any input $Y \in \mathcal{R}$ that is an affine combination of the challenges, outputs an element in the preimage of Y . It is syntactically analogous to the algebraic one-more discrete logarithm (AOMDL) problem [105], where the adversary wants to compute the discrete logarithms of $q + 1$ random challenges in \mathbb{G} by making at most q queries to the DLOG oracle, which outputs the discrete logarithm of the input Y only when Y is an affine combination of the challenges and the combination is known to the adversary.

The following theorem, our main result on AOMPR, shows that AOMPR of a linear hash function family is implied by its collision resistance.

Theorem 10.2.1. *For any linear hash function family LHF and any AOMPR adversary \mathcal{A} making at most q queries to CHAL, there exists an adversary \mathcal{B} for the $\text{CR}_{\text{LHF}}^{\text{LHF}}$ game running in a similar*

<p><u>Game AOMPR$_{\text{LHF}}^{\mathcal{A}}(\kappa)$:</u></p> <p>$par \leftarrow_{\\$} \text{PGen}(1^\kappa)$</p> <p>$cid \leftarrow 0 ; \ell \leftarrow 0$</p> <p>$\{y_i\}_{i \in [cid]} \leftarrow \mathcal{A}^{\text{CHAL}, \text{PI}}(par)$</p> <p>If $\ell \geq cid$ then return 0</p> <p>If $\forall i \in [cid] F(y_i) = X_i$ then</p> <p style="padding-left: 2em;">Return 1</p> <p>Return 0</p>	<p><u>Oracle CHAL() :</u></p> <p>$cid \leftarrow cid + 1$</p> <p>$x_{cid} \leftarrow_{\\$} \mathcal{D} ; X_{cid} \leftarrow F(x_{cid})$</p> <p>Return X_{cid}</p> <p><u>Oracle PI($Y, \alpha, \{\beta_i\}_{i \in [cid]}$) :</u></p> <p>Require: $Y = F(\alpha) + \sum_{i \in [cid]} \beta_i X_i$</p> <p>$\ell \leftarrow \ell + 1$</p> <p>Return $\alpha + \sum_{i \in [cid]} \beta_i x_i$</p>
--	---

Figure 10.2: The AOMPR game for a linear hash function family $\text{LHF} = (\text{PGen}, F)$. For the inputs of PI, X is in \mathcal{R} , α is in \mathcal{D} , and each β_i is in \mathcal{S} .

running time as \mathcal{A} such that $\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{A}, \kappa) \leq 2\text{Adv}_{\text{LHF}}^{\text{cr}}(\mathcal{B}, \kappa)$.

Proof of Theorem 10.2.1. Given an adversary \mathcal{A} for the $\text{AOMPR}^{\text{LHF}}$ game, without loss of generality, we assume that \mathcal{A} is deterministic, queries CHAL exactly q times, and queries PI exactly $q - 1$ times. The construction of \mathcal{B} is straightforward. After receiving par from the CR^{LHF} game, \mathcal{B} runs \mathcal{A} on input par by simulating the oracles CHAL and PI exactly the same as in the $\text{AOMPR}^{\text{LHF}}$ game. After \mathcal{A} outputs $\{y_i\}_{i \in [q]}$, if

$$\exists i \in [q] \text{ such that } F(y_i) = X_i \text{ and } y_i \neq x_i, \quad (10.1)$$

where x_i and X_i are generated in the oracle CHAL, then \mathcal{B} outputs (x_i, y_i) . Otherwise, \mathcal{B} aborts.

ANALYSIS OF \mathcal{B} . Denote the event $\text{WIN}_{\mathcal{B}}$ as after \mathcal{A} returns, the condition (10.1) holds. If $\text{WIN}_{\mathcal{B}}$ occurs, \mathcal{B} wins the CR^{LHF} game since $F(x_i) = X_i = F(y_i)$, which implies $\text{Adv}_{\text{LHF}}^{\text{cr}}(\mathcal{B}, \kappa) = \Pr[\text{WIN}_{\mathcal{B}}]$.

It is left to show that $\Pr[\text{WIN}_{\mathcal{B}}] \geq \frac{1}{2}\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{A}, \kappa)$. Since \mathcal{A} is deterministic, the execution of \mathcal{A} is fixed given the pair (par, \mathbf{x}) , where $\mathbf{x} \in \mathcal{D}^q$ denotes the randomness generated in the oracle CHAL. Denote the event $\text{WIN}_{\mathcal{A}}$ as \mathcal{A} wins the $\text{AOMPR}^{\text{LHF}}$ game simulated by \mathcal{B} . Since \mathcal{B} simulate the game perfectly, we know $\Pr[\text{WIN}_{\mathcal{A}}] = \text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{A}, \kappa)$. For each par , denote

$$\mathcal{W}_{\mathcal{A}} := \{\mathbf{x} \mid \text{WIN}_{\mathcal{A}} \text{ occurs given } (par, \mathbf{x})\},$$

$$\mathcal{W}_{\mathcal{B}} := \{\mathbf{x} \mid \text{WIN}_{\mathcal{B}} \text{ occurs given } (par, \mathbf{x})\}.$$

Claim 10.2.2. For each par , there exists a bijection $\Phi : \mathcal{W}_A \rightarrow \mathcal{W}_B$ such that for any $\mathbf{x} \in \mathcal{W}_A$, we have $\mathbf{x} \in \mathcal{W}_B \vee \Phi(\mathbf{x}) \in \mathcal{W}_B$.

From the above claim, we can conclude the proof since

$$\begin{aligned} \Pr[\text{WIN}_B] &= \Pr[\mathbf{x} \in \mathcal{W}_B] = \frac{1}{2} (\Pr[\mathbf{x} \in \mathcal{W}_B] + \Pr[\Phi(\mathbf{x}) \in \mathcal{W}_B]) \\ &\geq \frac{1}{2} \Pr[\mathbf{x} \in \mathcal{W}_B \vee \Phi(\mathbf{x}) \in \mathcal{W}_B] \geq \frac{1}{2} \Pr[\mathbf{x} \in \mathcal{W}_A] \\ &= \frac{1}{2} \Pr[\text{WIN}_A] = \frac{1}{2} \text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{A}, \kappa). \end{aligned}$$

□

Proof of Claim 10.2.2. We construct Φ as follows. For each $\mathbf{x} \in \mathcal{W}_A$, consider the execution of \mathcal{A} given (par, \mathbf{x}) . Denote $B \in \mathcal{S}^{(q-1) \times q}$ as the query matrix of the execution, which is defined as follows.

Definition 10.2.3. Given an execution of an adversary \mathcal{A} for the AOMPR game, where \mathcal{A} makes q queries to CHAL and ℓ queries to PI, define the query matrix of the execution as $B \in \mathcal{S}^{\ell \times q}$ such that

$$B_{i,j} = \begin{cases} \beta_i^{(j)}, & i \in [\text{cid}^{(j)}] \\ 0, & o.w. \end{cases},$$

where $\beta_i^{(j)}$ and $\text{cid}^{(j)}$ are the values of β_i and cid when \mathcal{A} makes the j -th query to PI.

We now define

$$\Phi(\mathbf{x}) := \mathbf{x} + \mathbf{u}^{(B)} z^*,$$

where $z^* \in \mathcal{D}$ and $\mathbf{u}^{(B)} \in \mathcal{S}^q$ are defined in the following claim.

Claim 10.2.4. There exists $z^* \in \mathcal{D}$ such that $F(z^*) = 0$ and for any matrix $A \in \mathcal{S}^{\ell \times q}$ where $0 < \ell < q$, there exists a vector $\mathbf{u}^{(A)} \in \mathcal{S}^q$ and $i \in [q]$ such that

$$A\mathbf{u}^{(A)} = 0 \wedge \exists i \in [q] : u_i^{(A)} z^* \neq 0. \quad (10.2)$$

Proof of Claim 10.2.4. Since F is not a monomorphism from \mathcal{D} to \mathcal{R} , there exists a non-zero element $z^* \in \mathcal{D}$ such that $F(z^*) = 0$. Since \mathcal{S} is a field and A has rank at most $\ell < q$, there exists a non-zero vector $\mathbf{u}^{(A)} \in \mathcal{S}^q$ such that $A\mathbf{u}^{(A)} = 0$. Also, since $\mathbf{u}^{(A)}$ is non-zero, there exists $i \in [q]$ such that $u_i^{(A)} \neq 0$, and since \mathcal{S} is a field and $z^* \neq 0$, we have $u_i^{(A)} z^* \neq 0$. □

ANALYSIS OF Φ . For simplicity, we use \mathbf{u} to denote $\mathbf{u}^{(B)}$ in the following analysis. We first show that the executions of \mathcal{A} given (par, \mathbf{x}) and given $(par, \Phi(\mathbf{x}))$ are identical. Since $F(\Phi(\mathbf{x})) = F(\mathbf{x}) + \mathbf{u} \cdot F(z^*) = F(\mathbf{x}) + \mathbf{u} \cdot 0 = F(\mathbf{x})$, the challenges output by CHAL are the same in the two executions. For the j -th query to PI, suppose the prior views of \mathcal{A} are identical. Then, \mathcal{A} must make the same query $\left(X^{(j)}, \alpha^{(j)}, \{\beta_i^{(j)}\}_{i \in [\text{cid}^{(j)}]}\right)$ in both executions. Since $B\mathbf{u} = 0$, we have $\alpha^{(j)} + \sum_{i \in [\text{cid}^{(j)}]} \beta_i^{(j)} x_i = \alpha^{(j)} + \left(\boldsymbol{\beta}^{(j)}\right)^T \mathbf{x} = \alpha^{(j)} + \left(\boldsymbol{\beta}^{(j)}\right)^T (\mathbf{x} + \mathbf{u}z^*) = \alpha^{(j)} + \sum_{i \in [\text{cid}^{(j)}]} \beta_i^{(j)} (\Phi(\mathbf{x}))_i$, where $\boldsymbol{\beta}^{(j)}$ denotes the j -th row of B . Therefore, \mathcal{A} receives the same value from PI in both executions. By induction, the views of \mathcal{A} are identical in both executions and thus \mathcal{A} outputs the same values in both executions, which implies $\Phi(\mathbf{x}) \in \mathcal{W}_A$ and thus Φ is a map from \mathcal{W}_A to \mathcal{W}_A .

Then, it is not hard to see that $\mathbf{x} \in \mathcal{W}_B \vee \Phi(\mathbf{x}) \in \mathcal{W}_B$. Since the executions of \mathcal{A} given \mathbf{x} and $\Phi(\mathbf{x})$ are identical, the outputs y_1, \dots, y_q of \mathcal{A} are also identical in the two executions. Since there exists $i \in [q]$ such that $u_i z^* \neq 0$, we have either $y_i \neq x_i$ or $y_i \neq x_i + u_i \cdot z^*$, which means WIN_B occurs either in the execution given \mathbf{x} or $\Phi(\mathbf{x})$.

It is left to show that Φ is a bijection. Since both the domain and range of Φ are \mathcal{W}_A , which is a finite set, it is enough to show that Φ is an injection. For any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{W}_A$ such that $\Phi(\mathbf{x}_1) = \Phi(\mathbf{x}_2)$, since the execution of \mathcal{A} given \mathbf{x}_1 is identical to that given $\Phi(\mathbf{x}_1)$ and the execution of \mathcal{A} given \mathbf{x}_2 is identical to that given $\Phi(\mathbf{x}_2)$, we know the executions of \mathcal{A} given \mathbf{x}_1 and \mathbf{x}_2 are identical, which implies the query matrix B in the two executions are identical. Therefore, we have $\Phi(\mathbf{x}_1) = \mathbf{x}_1 + \mathbf{u}z^*$ and $\Phi(\mathbf{x}_2) = \mathbf{x}_2 + \mathbf{u}z^*$ for the same $\mathbf{u} \in \mathcal{S}^q$, which implies $\mathbf{x}_1 = \mathbf{x}_2$. This shows that Φ is an injection. \square

10.3 Instantiations From the Discrete Logarithm Problem

Following the instantiation from [80], a linear hash function family GLHF is instantiated from a group generation algorithm GGen as follows.

- On input 1^κ , PGen runs GGen(1^κ) and receives a group description (\mathbb{G}, p, g) . Then, PGen uniformly samples $Z \in \mathbb{G}$ and returns $\kappa \leftarrow (\mathbb{G}, p, g, Z)$.
- Given $\kappa = (\mathbb{G}, p, g, Z)$, define $\mathcal{S} := \mathbb{Z}_p$, $\mathcal{D} := \mathbb{Z}_p^2$, $\mathcal{R} := \mathbb{G}$. Also, for any $(x_1, x_2) \in \mathbb{Z}_p^2$, define $F(x_1, x_2) := g^{x_1} Z^{x_2}$.
- The operation over \mathcal{D} is defined as follows. For any $(x_1, y_1), (x_2, y_2) \in \mathcal{D}$ and $s \in \mathcal{S}$, $(x_1, y_1) + (x_2, y_2) = (x_1 + x_2, y_1 + y_2)$ and $s \cdot (x_1, y_1) = (sx_1, sy_1)$.

- The operation over \mathcal{R} is defined as follows. For any $x_1, x_2 \in \mathcal{R}$ and $s \in \mathcal{S}$, $x_1 + x_2 = x_1 x_2$, $s \cdot x_1 = x_1^s$, where $x_1 x_2$ and x_1^s are the group operations of \mathbb{G} .

The following theorem shows that GLHF is a linear hash function family and collision resistance of GLHF is implied by the discrete logarithm assumption. Hauck et al. [80] shows similar statements, and we also give the proof below for completeness.

Lemma 10.3.1. *For any group generation algorithm GGen , $\text{GLHF}[\text{GGen}]$ is a linear hash function family (Definition 10.1.1). Moreover, for any adversary \mathcal{A} for the $\text{CR}^{\text{GLHF}[\text{GGen}]}$ game, there exists an adversary \mathcal{B} for the $\text{DLog}^{\text{GGen}}$ game such that $\text{Adv}_{\text{GLHF}[\text{GGen}]}^{\text{cr}}(\mathcal{A}, \kappa) \leq \text{Adv}_{\text{GGen}}^{\text{dlog}}(\mathcal{B}, \kappa)$.*

Proof of Lemma 10.3.1. Given $\text{par} = (\mathbb{G}, p, g, Z)$ output from $\text{PGen}(1^\kappa)$ that defines $(\mathcal{S}, \mathcal{D}, \mathcal{R}, \text{F})$, we need to show

1. \mathcal{D} and \mathcal{R} are \mathcal{S} -modules.
2. F is an epimorphism from \mathcal{D} to \mathcal{R} but not a monomorphism.
3. The collision resistance of RLHF is implied by the RSA assumption.

PART 1. It is clear that $\mathcal{D} = \mathbb{G} \times \mathbb{G}$ and $\mathcal{R} = \mathbb{G}$ are \mathbb{Z}_p -modules.

PART 2. It is easy to verify F is a homomorphism of \mathcal{S} -modules, since for any $b \in \mathcal{S}$ and $(x_1, y_1), (x_2, y_2) \in \mathcal{D}$,

$$\begin{aligned} \text{F}((x_1, y_1) + b \cdot (x_2, y_2)) &= \text{F}(x_1 + bx_1, y_1 + by_2) \\ &= g^{x_1+bx_2} Z^{y_1+by_2} \\ &= g^{x_1} Z^{x_1} (g^{x_2} Z^{y_2})^b \\ &= \text{F}(x_1, y_1) + b\text{F}(x_2, y_2). \end{aligned}$$

F is epimorphism since for any $X \in \mathcal{R}$, we have $\text{F}(x, 0) = X$, where x denotes the discrete log of X to the base g . Denote z as the discrete log of Z to base g , and F is not a monomorphism, since $\text{F}(z, -1) = g^z Z^{-1} = g^0$.

PART 3. For any adversary \mathcal{A} for the CR^{GLHF} game, we construct \mathcal{B} for the DLog game as follows. After receiving (\mathbb{G}, p, g, Z) , \mathcal{B} runs \mathcal{A} with input (\mathbb{G}, p, g, Z) . If \mathcal{A} wins the CR^{GLHF} by outputting $(x_1, y_1), (x_2, y_2) \in \mathcal{D}$, such that $(x_1, y_1) \neq (x_2, y_2)$ and $\text{F}(x_1, y_1) = \text{F}(x_2, y_2)$, we have $g^{x_1} Z^{y_1} = g^{x_2} Z^{y_2}$. Therefore, \mathcal{B} can compute $z = \frac{y_2 - y_1}{x_1 - x_2}$ and we have $g^z = Z$. \square

10.4 Threshold Signatures from LHF

OUR SCHEMES. Figure 10.3 shows the protocols FROST1-H and FROST2-H that are transformed from FROST1 and FROST2, respectively. In addition to the general transformation, we need to pick an injection $x_{(\cdot)} : [n] \rightarrow \mathcal{S}$. The choice of $x_{(\cdot)}$ can be arbitrary, and the corresponding Lagrange coefficient for a set of index $S \subseteq [n]$ and $i \in S$ is defined as $\lambda_i^S := \prod_{j \in S \setminus \{i\}} \frac{x_j}{x_i - x_j}$. We analyse the correctness of the scheme in Section 10.4.1 Also, similar to the multi-signature case, we optimize the schemes by sampling key shares from $\mathcal{D}_{\text{key}} \subseteq \mathcal{D}$ and setting the hash range to be $\mathcal{S}_{\text{hash}} \subseteq \mathcal{S}$.

The following theorems show that, under the AOMPR assumption, FROST2-H is TS-SUF-2-secure and FROST1-H is TS-SUF-3-secure in the random oracle model. We prove the theorems using the same techniques from Chapter 9. Here, we briefly highlight the differences:

- We need to show that \mathcal{B} simulates the TS-SUF-2 and TS-SUF-3 games perfectly when no bad event occurs and that the bad events occur with a negligible probability when the secret key is sampled from \mathcal{D}_{key} instead of \mathbb{Z}_p , and the randomness r_j is sampled from \mathcal{D} instead of \mathbb{Z}_p .
- We need to show that \mathcal{B} can compute a preimage for each challenge (Claim 10.4.5 and Claim 10.4.8) instead of the discrete logarithm to the base element. More precisely, the problem can be described as follows. Denote the challenges by $U_1, \dots, U_\ell \in \mathcal{R}$. After the interaction with \mathcal{A} , \mathcal{B} computes a matrix $A \in \mathcal{S}^{\ell \times \ell}$ and a vector $\mathbf{b} \in \mathcal{D}^\ell$ such that $A \cdot \mathbf{U} = \mathbf{F}(\mathbf{b})$, we need to show that A has full rank and thus \mathcal{B} can compute a vector $\mathbf{u} = A^{-1}\mathbf{b}$ such that $\mathbf{F}(\mathbf{u}) = \mathbf{U}$.

Theorem 10.4.1. *For any TS-SUF-2 adversary \mathcal{A} game making at most q_s queries to PPO and q_h queries to RO, there exists an AOMPR adversary \mathcal{B} making at most $2q_s + t$ queries to CHAL running in time roughly equal two times that of \mathcal{A} such that*

$$\text{Adv}_{\text{FROST2-H[LHF]}}^{\text{ts-suf-2}}(\mathcal{A}, \kappa) \leq \sqrt{q \cdot (\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{B}, \kappa) + (5q^2)/2^\kappa)},$$

where $q = q_h + q_s + 1$.

Theorem 10.4.2. *For any TS-SUF-3 adversary \mathcal{A} making at most q_s queries to PPO and q_h queries to RO, there exists an AOMPR adversary \mathcal{B} making at most $2q_s + t$ queries to CHAL running in time roughly equal two times that of \mathcal{A} such that*

$$\text{Adv}_{\text{FROST1-H[LHF]}}^{\text{ts-suf-3}}(\mathcal{A}, \kappa) \leq 4n \cdot q \cdot \sqrt{(\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{B}, \kappa) + 6q/2^\kappa)},$$

<p><u>Algorithm Setup(1^κ) :</u> $par \leftarrow_s \text{PGen}(1^\kappa)$ Return par</p> <p><u>Algorithm KeyGen() :</u> For $i \in [0..t-1]$ do $a_i \leftarrow_s \mathcal{D}_{\text{key}}$ For $i \in [n]$ do $sk_i \leftarrow_s \sum_{j=0}^{t-1} a_j \cdot x_i^j$ $pk_i \leftarrow F(sk_i)$ $pk \leftarrow F(a_0)$ $aux \leftarrow (pk_1, \dots, pk_n)$ Return $pk, aux, \{sk_i\}_{i \in [1..n]}$</p> <p><u>Algorithm SPP(st_i) :</u> $r \leftarrow_s \mathcal{D} ; s \leftarrow_s \mathcal{D}$ $pp \leftarrow (F(r), F(s))$ $st_i.\text{mapPP}(pp) \leftarrow (r, s)$ Return (pp, st_i)</p> <p><u>Algorithm Vf(pk, m, sig) :</u> $(R, s) \leftarrow sig$ $c \leftarrow H_2(pk, m, R)$ Return $(F(s) = R + c \cdot pk)$</p>	<p><u>Algorithm LPP(i, pp, st_0) :</u> $st_0.\text{curPP}_i \leftarrow st_0.\text{curPP}_i \cup \{pp\}$ Return st_0</p> <p><u>Algorithm LR(M, SS, st_0) :</u> If $\exists i \in SS : st_0.\text{curPP}_i = \emptyset$ then Return \perp $lr.\text{msg} \leftarrow M ; lr.SS \leftarrow SS$ For $i \in SS$ do Pick pp_i from $st_0.\text{curPP}_i$ $lr.PP(i) \leftarrow pp_i$ $st_0.\text{curPP}_i \leftarrow st_0.\text{curPP}_i \setminus \{pp_i\}$ Return (lr, st_0)</p> <p><u>Algorithm PS(lr, i, st_i) :</u> $pp_i \leftarrow lr.PP(i)$ If $st_i.\text{mapPP}(pp_i) = \perp$ then Return (\perp, st_i) $(r_i, s_i) \leftarrow st_i.\text{mapPP}(pp_i)$ $st_i.\text{mapPP}(pp_i) \leftarrow \perp$ $(R, c, \{d_j\}_{j \in lr.SS})$ $\leftarrow \text{CompPar}(st_i.pk, lr)$ $z_i \leftarrow r_i + d_i \cdot s_i + c \cdot \lambda_i^{lr.SS} \cdot st_i.sk$ Return $((R, z_i), st_i)$</p>	<p><u>Algorithm CompPar(pk, lr) :</u> $m \leftarrow lr.\text{msg} ; (R^*, s^*) \leftarrow sig$ For $i \in lr.SS$ do <div style="border: 1px solid black; padding: 2px; display: inline-block;"> $d_i \leftarrow H_1(pk, lr, i)$ </div> <div style="border: 1px dashed black; padding: 2px; display: inline-block;"> $d_i \leftarrow H_1(pk, lr)$ </div> $(R_i, S_i) \leftarrow lr.PP(i)$ $R \leftarrow \sum_{i \in lr.SS} (R_i + d_i S_i)$ $c \leftarrow H_2(pk, M, R)$ Return $(R, c, \{d_i\}_{i \in lr.SS})$</p> <p><u>Algorithm Agg(PS, st_0) :</u> $R \leftarrow \perp ; z \leftarrow 0$ For $(R', z') \in PS$ do If $R = \perp$ then $R \leftarrow R'$ If $R \neq R'$ then return (\perp, st_0) $z \leftarrow z + z'$ Return $((R, z), st_0)$</p> <p><u>Algorithm SVf(pk, lr, sig) :</u> $(R^*, z^*) \leftarrow sig$ $(R, c, \{d_j\}_{j \in lr.SS})$ $\leftarrow \text{CompPar}(st_i.pk, lr)$ Return $(R = R^*)$ $\wedge (F(z^*) = R + c \cdot pk)$</p>
---	---	---

Figure 10.3: The protocol FROST1-H[LHF] and FROST1-H[LHF], where LHF = (PGen, F) is a linear hash function family. The protocol FROST1-H contains all but the dashed box, and the protocol FROST2-H contains all but the solid box. Further, n is the number of parties, and t is the threshold of the schemes. $x_{(\cdot)}$ is an injection from $[n]$ to \mathcal{S} and $\lambda_i^{lr.SS}$ denotes the Lagrange coefficient which is computed as $\lambda_i^{lr.SS} := \prod_{j \in S \setminus \{i\}} \frac{x_j}{x_j - x_i}$. \mathcal{D}_{key} is a subset of \mathcal{D} such that F is a bijection between \mathcal{D}_{key} and \mathcal{S} .

where $q = q_h + q_s + 1$.

To instantiate FROST1-H and FROST2-H in pairing-free groups, we set $\mathcal{D}_{\text{key}} := \{(x, 0) : x \in \mathbb{Z}\}$ and $\mathcal{S}_{\text{hash}} := \mathcal{S}$. It is clear that $\text{char}(\mathcal{S}) = p \geq 2^\kappa$, F is a bijection from \mathcal{D}_{key} to \mathcal{R} , and $|\mathcal{S}_{\text{hash}}| = |\mathcal{S}| \geq 2^\kappa$. Also, for instantiating FROST1-H and FROST2-H, we set $x_i := i$.

By combining Theorem 10.2.1 and Lemma 10.3.1 with the above theorems, we show the security of FROST1-H and FROST2-H instantiated from GLHF (defined in Section 10.3) under the discrete logarithm assumption in the random oracle model.

10.4.1 Correctness of FROST1-H and FROST2-H

For the correctness of the schemes, we just need to show Shamir's secret sharing scheme works over \mathcal{D} . More precisely, we want to show for any $S \subseteq [n]$ with size t , we have

$$\sum_{i \in S} \lambda_i^S \text{sk}_i = a_0, \quad (10.3)$$

where

$$\text{sk}_i = \sum_{j=0}^{t-1} a_j x_i^j, \quad \lambda_i^S = \prod_{j \in S \setminus \{i\}} \frac{x_j}{x_j - x_i},$$

for each $i \in S$ and some $a_0, \dots, a_{t-1} \in \mathcal{D}$.

Claim 10.4.3. For any $S \subseteq [n]$ with size t and any integer $0 \leq k < t$, we have

$$\sum_{i \in S} \lambda_i^S x_i^k = \begin{cases} 1, & k = 0 \\ 0, & k \geq 1 \end{cases}. \quad (10.4)$$

With the above claim, we can show (10.3) since

$$\begin{aligned} \sum_{i \in S} \lambda_i^S \text{sk}_i &= \sum_{i \in S} \lambda_i^S \sum_{j=0}^{t-1} a_j x_i^j = \sum_{j=0}^{t-1} a_j \sum_{i \in S} \lambda_i^S x_i^j \\ &= \sum_{j=0}^{t-1} a_j \cdot \mathbf{1}\{j = 0\} = a_0. \end{aligned}$$

Proof of Claim 10.4.3. Define $f(x) = x^k$. By the polynomial interpolation over the field \mathcal{S} , we have $\sum_{i \in S} \lambda_i^S f(x_i) = f(0)$, which proves claim. \square

10.4.2 Proof of Theorem 10.4.1

Let \mathcal{A} be an adversary as described in the theorem. Denote the output message-signature pair of \mathcal{A} as $(m^*, sig^* = (R^*, z^*))$. Without loss of generality, we assume \mathcal{A} always queries RO on $H_2(pk, m^*, R^*)$ before \mathcal{A} returns and always queries RO on $H_1(pk, lr)$ prior to the query $PSIGNO(i, lr)$ for some i and lr . (This adds up to q_s additional RO queries, and we let $q = q_h + q_s + 1$.) Denote lr^* as the leader query such that $H_1(pk, lr^*)$ is the first query prior to the query $H_2(pk, m^*, R^*)$ satisfying $SVf(pk, lr^*, sig^*) = \text{true}$. If such lr^* does not exist, lr^* is set to \perp . Denote the event E_1 as

$$\text{Vf}(pk, m^*, sig^*) \wedge (lr^* = \perp \vee \text{curSSL}(lr^*) < t - |CS|) .$$

It is clear that if \mathcal{A} wins the game $\text{TS-SUF-2}^{\text{FROST2-H[LHF]}}$, then E_1 must occur, which implies $\Pr[E_1] \geq \text{Adv}_{\text{FROST2-H[LHF]}}^{\text{ts-suf-2}}(\mathcal{A})$. Therefore, the theorem will follow from the following lemma. (We isolate this statement as its own lemma also because it will be helpful in the proof of Theorem 10.4.2 below.)

Lemma 10.4.4. *There exists an adversary \mathcal{B} for the $\text{AOMPR}^{\text{LHF}}$ game making at most $2q_s + t$ queries to CHAL such that*

$$\Pr[E_1] \leq \sqrt{q \cdot (\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{B}) + 5q^2/2^\kappa)} .$$

Moreover, \mathcal{B} runs in time roughly twice that of \mathcal{A} .

Proof of Lemma 10.4.4. We first construct an algorithm \mathcal{C} compatible with the syntax in Lemma 9.2.3 and then construct \mathcal{B} from $\text{Fork}^{\mathcal{C}}$. The input of \mathcal{C} consists of par that defines a linear hash function $(\mathcal{S}, \mathcal{D}, \mathcal{R}, F)$ and uniformly random elements $h_1, \dots, h_{2q} \in \mathcal{S}_{\text{hash}}$. Also, \mathcal{C} can access oracles CHAL and PI defined the same as those in the $\text{AOMPR}^{\text{LHF}}$ game. (We can think of the oracles as part of the input of \mathcal{C} in the context of the Forking Lemma.) For simplicity, when \mathcal{C} makes a query $(X, \alpha, \{\beta_i\})$ to PI, we omit the coefficients $\alpha, \{\beta_i\}$ which are clear from the context. To start with, \mathcal{C} makes $2q_s + t$ queries to CHAL and denotes the challenges as $A_0, \dots, A_{t-1}, U_1, V_1, \dots, U_{q_s}, V_{q_s} \in \mathcal{D}$. Then, \mathcal{C} initializes all the states st_0, \dots, st_n . In addition, it initializes counters $\text{ctr}_s, \text{ctr}_h$ to 0 and a function dt to an empty table, which are used to record the PI query related to each (U_j, V_j) . \mathcal{C} also initializes $\text{curLR} \leftarrow \emptyset$ to record all leader requests that appears during the game and initializes ctrPP to an empty table, which are used to record the counter corresponding to each token generated by honest parties. We also use a flag BadPPO to denote whether a

bad event occurs, which are initially set to `false`. Then, \mathcal{C} runs \mathcal{A} with access to the oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}, \widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$: \mathcal{C} initializes H to an empty table and sets $\overline{\text{pk}} \leftarrow A_0, \overline{\text{pk}}_i = \prod_{j=0}^{t-1} A_j \cdot x_i^j$ for $i \in [n]$, and $\overline{\text{sk}}_i \leftarrow \text{PI}(\overline{\text{pk}}_i)$ for $i \in CS$. \mathcal{C} samples $\tilde{a}_i \leftarrow_{\$} \mathcal{D}_{\text{key}}$ for $i \in [0..(t-1)]$ and sets $\text{sk}_i \leftarrow \sum_{j=0}^{t-1} \tilde{a}_j x_i^j$ for $i \in CS$. Then, \mathcal{C} computes a polynomial $f(x) = \sum_{i=0}^{t-1} \mu_i x^i$ such that $\mu_i \in \mathcal{D}$ for $i \in [0..(t-1)]$, $f(x_i) = \text{sk}_i - \overline{\text{sk}}_i$ for $i \in CS$, and $f(x_i) = 0$ for $i \in S'$, where $S' \subseteq [n]$ denotes the set of the first $(t - |CS|)$ honest parties.¹ \mathcal{C} sets $\text{pk}_i \leftarrow \overline{\text{pk}}_i + F(f(x_i))$ for $i \in [n]$. Finally, \mathcal{C} returns $(\text{pk}, \text{aux} = (\text{pk}_1, \dots, \text{pk}_n), \{\text{sk}_i\}_{i \in CS})$.

$\widetilde{\text{RO}}$ query $H_1(x)$: If $H_1(x) \neq \perp$, \mathcal{C} returns $H_1(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, lr)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_1(x) \leftarrow_{\$} \mathcal{S}_{\text{hash}}$ and returns $H_1(x)$. Otherwise, \mathcal{C} increases ctr_h by 1, sets $H_1(x) \leftarrow h_{2\text{ctr}_h-1}$, and adds lr to curLR . Also, \mathcal{C} computes $R \leftarrow \sum_{i \in lr.\text{SS}} (R_i + h_{2\text{ctr}_h-1} \cdot S_i)$, where $(R_i, S_i) \leftarrow lr.\text{PP}(i)$. If $H_2(\text{pk}, lr.\text{msg}, R) = \perp$, \mathcal{C} sets $H_2(\text{pk}, lr.\text{msg}, R) = h_{2\text{ctr}_h}$. In addition, define $\text{mapLR}(\text{ctr}_h) := lr$ and set $\text{curLR} \leftarrow \text{curLR} \cup \{lr\}$. Finally, \mathcal{C} returns $H_1(x)$.

$\widetilde{\text{RO}}$ query $H_2(x)$: If $H_2(x) \neq \perp$, \mathcal{C} returns $H_2(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, m, R)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_2(x) \leftarrow_{\$} \mathcal{S}_{\text{hash}}$ and returns $H_2(x)$. Otherwise, \mathcal{C} increases ctr_h by 1 and sets $H_2(x) \leftarrow h_{2\text{ctr}_h}$. Finally, \mathcal{C} returns $H_2(x)$.

$\widetilde{\text{PPO}}(i)$ query: Same as in the game $\text{TS-SUF-2}^{\text{FROST}^2\text{-H}}$, except in the simulation of algorithm SPP, \mathcal{C} first increases ctr_s by 1 and sets $pp \leftarrow (U_{\text{ctr}_s}, V_{\text{ctr}_s})$, $\text{st}_i.\text{mapPP}(pp) \leftarrow (0, 0)$, and $\text{ctrPP}(i, pp) \leftarrow \text{ctr}_s$. In addition, BadPPO is set to `true` if there exists $lr \in \text{curLR}$ such that $lr.\text{PP}(i) = (U_{\text{ctr}_s}, V_{\text{ctr}_s})$.

$\widetilde{\text{PSIGNO}}(i, lr)$ query: Same as in the game $\text{TS-SUF-2}^{\text{FROST}^2\text{-H}}$, except in the simulation of algorithm PS, if $\text{st}_i.\text{mapPP}(pp) \neq \perp$, \mathcal{C} sets

$$z_i \leftarrow \text{PI} \left(U_j + d_i V_j + c \cdot \lambda_i^{lr.\text{SS}} \cdot \text{pk}_i \right),$$

where $j \leftarrow \text{ctrPP}(i, lr.\text{PP}(i))$. In addition, \mathcal{C} sets $\text{dt}(j) \leftarrow (i, k, d_i, c\lambda_i^{lr.\text{SS}}, z_i)$, where k denotes the index such that $H_1(\text{pk}, lr)$ is set to h_{2k-1} during the simulation.

¹Since the degree of f is t and we fix t points of f , f is fixed and we can compute the coefficients of f by solving a linear equation.

After receiving the output $(m^*, sig^* = (R^*, z^*))$ from \mathcal{A} , \mathcal{C} returns \perp if $\text{BadPPO} = \text{true}$ or E_1 does not occur. Otherwise, \mathcal{C} finds the index I such that $H_2(\text{pk}, m^*, R^*)$ is set to h_I during the simulation. By our assumption of \mathcal{A} , we know such I must exist. Then, \mathcal{C} returns (I, Out) , where Out consists of all variables received or generated by \mathcal{C} .

ANALYSIS OF \mathcal{C} . To use Lemma 9.2.3, we define $S := \{2k\}_{k \in [1..q]}$ and IG as the algorithm that runs $\text{PGen}(1^\kappa)$ and outputs par . From the simulation, we know the output index I of \mathcal{C} is always in S .

It is not hard to see \mathcal{C} simulates the game $\text{TS-SUF-}2^{\text{FROST}^2\text{-H}}$ perfectly, which implies $\text{acc}(\mathcal{C}) \geq \Pr[E_1] - \Pr[\text{BadPPO}]$, where $\Pr[E_1]$ refers to the probability in the original $\text{TS-SUF-}2^{\text{FROST}^2\text{-H[LHF]}}$ game with \mathcal{A} (as in the lemma statement), whereas $\Pr[\text{BadPPO}]$ is the probability that $\text{BadPPO} = \text{true}$ at the end of \mathcal{C} 's execution. Since (U_j, V_j) is sampled uniformly from \mathcal{R} . Therefore, for each $\text{PPO}(i)$ query, the probability BadPPO is set to true is less than $|\text{curLR}|/|\mathcal{R}| \leq q_h/2^\kappa$. Therefore, we have $\Pr[\text{BadPPO}] \leq q_s q_h / 2^\kappa$. By Lemma 9.2.3,

$$\begin{aligned} \text{acc}(\text{Fork}^{\mathcal{C}}) &\geq (\Pr[E_1] - q_s q_h / 2^\kappa)^2 / q \geq \Pr[E_1]^2 / q - 2\Pr[E_1] q_s q_h / (2^\kappa \cdot q) \\ &\geq \Pr[E_1]^2 / q - q / 2^\kappa. \end{aligned}$$

CONSTRUCT \mathcal{B} FROM $\text{Fork}^{\mathcal{C}}$. We now give a construct of the AOMPR adversary \mathcal{B} using $\text{Fork}^{\mathcal{C}}$, and the available CHAL and PI oracles. To start with, \mathcal{B} receives par from the $\text{AOMPR}^{\text{LHF}}$ game and runs $\text{Fork}^{\mathcal{C}}(par)$. All the CHAL queries from the first execution of \mathcal{C} are relayed by \mathcal{B} to its own CHAL oracle, and for all the CHAL queries from the second execution of \mathcal{C} , \mathcal{B} answers them with the same challenges as the first execution. All the PI queries from $\text{Fork}^{\mathcal{C}}$ are relayed by \mathcal{B} to its own PI oracle. Without loss of generality, we can assume all the challenges are different, since otherwise, \mathcal{B} can solve them trivially. Denote the event BadHash as $\{h_1, \dots, h_{2q}\} \cap \{h'_1, \dots, h'_{2q}\} \neq \emptyset$, where $h_1, h'_1, \dots, h_{2q}, h'_{2q}$ generated in the execution of $\text{Fork}^{\mathcal{C}}$ are same. Since the hash values are sampled uniformly from $\mathcal{S}_{\text{hash}}$, we know $\Pr[\text{BadHash}] \leq 4q^2 / |\mathcal{S}_{\text{hash}}| \leq 4q^2 / 2^\kappa$. Then, we can conclude the proof with the following claim, which implies

$$\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{B}) \geq \text{acc}(\text{Fork}^{\mathcal{C}}) - \Pr[\text{BadHash}] \geq \Pr[E_1]^2 / q - 5q^2 / 2^\kappa.$$

□

Claim 10.4.5. \mathcal{B} can win the game $\text{AOMPR}_{\text{LHF}}$, if $\text{Fork}^{\mathcal{C}}$ returns $(I, \text{Out}, \text{Out}')$ and BadHash does not occur.

Proof. We directly use the notations in the description of \mathcal{C} to denote the variables in Out and use $(\cdot)'$ to denote the variables in Out' . The total number of the CHAL queries is $t + 2q_s$ and the corresponding challenges are $A_0, \dots, A_{t-1}, U_1, V_1, \dots, U_{q_s}, V_{q_s}$.

We first show how to compute a_0, \dots, a_{t-1} such that $F(a_i) = A_i$. By the execution of $\text{Fork}^{\mathcal{C}}$, we know $(\text{pk}, m^*, R^*) = (\text{pk}', m'^*, R'^*)$ and $\text{pk} = A_0$. Since $I \in S$, let $k^* = I/2$. It is not hard to see that $\text{mapLR}(k^*) = lr^*$. (If $\text{mapLR}(k^*) = \perp$, lr^* is also \perp .) Since BadHash does not occur, we have $H_2(\text{pk}, m^*, R^*) = h_I \neq h'_I = H_2(\text{pk}, m^*, R^*)$. Since $F(z^*) = R^* + h_I A_0$, $F(z'^*) = R^* + h'_I A_0$, we have $F(z^* - z'^*) = (h_I - h'_I)A_0$ and therefore \mathcal{B} computes $a_0 \leftarrow \frac{z^* - z'^*}{h_I - h'_I}$.

Define $T_{\text{dt}} := \{j : (i, k, d, c, z) \leftarrow \text{dt}(j), k = k^*\}$. For each $j \in T_{\text{dt}} \cap T_{\text{dt}'}$, let $(i, k, d, c, z) \leftarrow \text{dt}(j)$ and $(i', k', d', c', z') \leftarrow \text{dt}'(j)$, and we have $F(z) = U_j + dV_j + c \cdot \text{pk}_i$, $F(z') = U_j + d'V_j + c' \cdot \text{pk}_{i'}$, $c = h_I \lambda_i^{lr^* \cdot \text{SS}}$, and $c' = h'_I \lambda_{i'}^{lr'^* \cdot \text{SS}}$. Since $\text{BadPPO} = \text{false}$ during both execution of \mathcal{C} , we know (U_j, V_j) is returned by a query $\text{PPO}(i)$ prior to the query $H_2(\text{pk}, m^*, R^*)$ during the first execution of \mathcal{C} . Since the two executions of \mathcal{C} are exactly the same prior to the query $H_2(\text{pk}, m^*, R^*)$, we know $i' = i$. Also, we know $d = h_{2k-1} = h_{2k^*-1} = h_{2k'-1} = d'$, which implies $F(z - z') = \lambda_{i'}^{lr'^* \cdot \text{SS}}(h_I - h'_I) \cdot \text{pk}_i$. Since $h_I \neq h'_I$, \mathcal{B} computes $y_i \leftarrow \frac{z - z'}{\lambda_{i'}^{lr'^* \cdot \text{SS}}(h_I - h'_I)}$, which satisfies $F(y_i) = \text{pk}_i$. Denote $D := \{i\}_{j \in T_{\text{dt}} \cap T_{\text{dt}'}, (i, k, d, c, z) \leftarrow \text{dt}(j)}$. Since E_1 occurs in the first execution of \mathcal{C} , we know $|T_{\text{dt}}| = |\text{curSSL}(lr^*)| < t - |CS|$. Therefore, we know $|D| = |T_{\text{dt}} \cap T_{\text{dt}'}| < t - |CS|$. Therefore, \mathcal{B} can pick an arbitrary set $D' \in HS \setminus D$ with size $(t - |CS| - |T_{\text{dt}} \cap T_{\text{dt}'}| - 1)$ and for each $i \in D'$, \mathcal{B} sets $y_i \leftarrow \text{PI}(\text{pk}_i)$. Denote $D_{\text{tot}} = CS \cup D \cup D'$, and we have $|D_{\text{tot}}| = t - 1$ and for each $i \in D_{\text{tot}}$, \mathcal{B} knows y_i such that $F(y_i) = \text{pk}_i$. Since $\text{pk}_i = F(f(x_i)) + A_0 + \sum_{j \in [t-1]} A_j \cdot x_i^j$, denote $D_{\text{tot}} = \{i_1, \dots, i_{t-1}\}$ and we have

$$M \cdot \begin{pmatrix} A_1 \\ \dots \\ A_{t-1} \end{pmatrix} = \begin{pmatrix} F(y_{i_1} - f(x_{i_1}) - a_0) \\ \dots \\ F(y_{i_{t-1}} - f(x_{i_{t-1}}) - a_0) \end{pmatrix}, \text{ where } M = \begin{pmatrix} x_{i_1} & \dots & x_{i_1}^{t-1} \\ \vdots & \ddots & \vdots \\ x_{i_{t-1}} & \dots & x_{i_{t-1}}^{t-1} \end{pmatrix}. \quad (10.5)$$

Since M is a Vandermonde matrix, we know M has full rank and thus \mathcal{B} can compute a_1, \dots, a_{t-1} from (10.5) such that $F(a_i) = A_i$ for $i \in [t-1]$. Further, for $i \in [n] \setminus D_{\text{tot}}$, \mathcal{B} computes $y_i \leftarrow f(x_i) + \sum_{j \in [0..(t-1)]} a_j \cdot x_i^j$, which satisfies $F(y_i) = \text{pk}_i$.

We now show how to compute $u_1, v_1, \dots, u_{q_s}, v_{q_s}$ such that $F(u_i) = U_i$ and $F(v_i) = V_i$. From the execution of \mathcal{C} , we know $\text{dt}(j) = (i, k, d, c, z) \neq \perp$ if and only if \mathcal{C} queries PI on $U_j + dV_j + c \cdot \text{pk}_i$. Therefore, denote $U_j + dV_j + c \cdot \text{pk}_i$ as the PI query associated with $\text{dt}(j)$. For each $j \in q_s$, there are the following cases.

Case 0: Both $\text{dt}(j)$ and $\text{dt}'(j)$ are \perp . In this case, \mathcal{B} computes u_j and v_j by directly querying oracle PI on U_j and V_j .

Case 1: Exactly one of $\text{dt}(j)$ and $\text{dt}'(j)$ is not \perp . Without loss of generality, assume $\text{dt}(j) = (i, k, d, c, z)$, which implies $F(z) = U_j + dV_j + c \cdot \text{pk}_i$. \mathcal{B} computes v_j by directly querying oracle PI and computes $u_j \leftarrow z - d \cdot v_j - c \cdot y_i$.

For all the following cases, both $\text{dt}(j)$ and $\text{dt}'(j)$ are not \perp and we denote $(i, k, d, c, z) \leftarrow \text{dt}(j)$ and $(i', k', d', c', z') \leftarrow \text{dt}'(j)$.

Case 2: $k \neq k'$ or $k = k' > k^*$. In this case, we know $d = h_{2k-1} \neq h_{2k'-1} = d'$ and $F(z) = U_j + dV_j + c \cdot \text{pk}_i$, $F(z') = U_j + d'V_j + c' \cdot \text{pk}_{i'}$. Therefore, \mathcal{B} computes $v_j \leftarrow \frac{z - c \cdot y_i - z' + c' \cdot y_{i'}}{d - d'}$, and $u_j \leftarrow z - d \cdot v_j - c \cdot y_i$.

Case 3: $k = k' = k^*$. In this case, \mathcal{B} computes v_j, u_j the same as Case 1.

Case 4: $k = k' < k^*$. \mathcal{B} computes v_j, u_j the same as Case 1. Also, in this case, we have $d = d'$ and $c = c'$. Therefore, \mathcal{B} queries PI oracle only once in order to simulate the PI queries associated with $\text{dt}(j)$ and $\text{dt}'(j)$.

We now count the number of PI queries made by \mathcal{B} .

- \mathcal{B} queries PI oracle $|CS|$ times queries for simulating query $\text{PI}(\text{pk}_i)$ made by \mathcal{C} for each $i \in |CS|$.
- \mathcal{B} queries PI oracle $|D'|$ times queries for computing a_0, \dots, a_{t-1} .
- For each $j \in \mathfrak{q}_s$, \mathcal{B} queries PI twice for simulating query associated with $\text{dt}(j)$ and $\text{dt}'(j)$ and computing u_j and v_j in case 0, 1, 2, 4 and queries 3 times in case 3.

Since the condition of case 3 is equivalent to $j \in T_{\text{dt}} \cap T_{\text{dt}'}$, the total number of PI queries made by \mathcal{B} is equal to $2\mathfrak{q}_s + |T_{\text{dt}} \cap T_{\text{dt}'}| + |CS| + |D'| = 2\mathfrak{q}_s + t - 1$. Therefore, \mathcal{B} wins the game $\text{AOMPR}_{\text{LHF}}$. \square

10.4.3 Proof of Theorem 10.4.2

Let \mathcal{A} be the adversary described in the theorem. Denote the output message-signature pair of \mathcal{A} as $(m^*, sig^* = (R^*, z^*))$. Without loss of generality, we assume \mathcal{A} always queries RO on $H_2(pk, m^*, R^*)$ before \mathcal{A} returns and always queries RO on $H_1(pk, lr, i)$ prior to the query PSIGNO(i, lr) for some i and lr . (This adds up to q_s additional RO queries, and we let $q = q_h + q_s + 1$.) Denote lr^* as the leader query such that $H_1(pk, lr^*, i)$ is the first RO query prior to the $H_2(pk, m^*, R^*)$ query for some i satisfying $SVf(pk, lr^*, sig^*) = \text{true}$. If such lr^* does not exist, lr^* is set to \perp . Denote the event E_1 as

$$\text{Vf}(pk, m^*, sig^*) \wedge (lr^* = \perp \vee \text{curSSL}(lr^*) < t - |CS|).$$

Denote the event E_2 as

$$\text{Vf}(pk, m^*, sig^*) \wedge lr^* \neq \perp \wedge \text{curSSL}(lr^*) \neq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}.$$

If \mathcal{A} wins the game TS-SUF-3^{FROST1-H} and $lr^* \neq \perp$, we know either $\text{curSSL}(lr^*) < t - |CS|$ or $\text{curSSL}(lr^*) \neq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$. Therefore, if \mathcal{A} wins the game TS-SUF-3^{FROST1-H}, then either E_1 or E_2 occurs, which implies

$$\text{Adv}_{\text{FROST1-H}[\text{LHF}]}^{\text{ts-suf-3}}(\mathcal{A}) \leq \Pr[E_1] + \Pr[E_2] \leq 2 \max\{\Pr[E_1], \Pr[E_2]\}.$$

Thus, we conclude the theorem with the following two lemmas.

Lemma 10.4.6. *There exists an adversary \mathcal{B} for the AOMPR^{LHF} game making at most $2q_s + t$ queries to CHAL such that*

$$\Pr[E_1] \leq \sqrt{q \cdot (\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{B}) + 3q^2(n+1)^2/2^\kappa)},$$

Moreover, \mathcal{B} runs in time roughly equal two times that of \mathcal{A} .

Lemma 10.4.7. *There exists an adversary \mathcal{B} for the AOMPR^{LHF} making at most $2q_s$ queries to CHAL such that*

$$\Pr[E_2] \leq n \cdot q \sqrt{2(\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{B}) + 1/2^\kappa)}.$$

Moreover, \mathcal{B} runs in time roughly equal two times that of \mathcal{A} .

This completes the proof of the theorem, subject to proofs of the lemmas that we discuss next.

The proof of Lemma 10.4.6 is almost the same as Lemma 10.4.4, so we omit the full proof. The only difference is that \mathcal{C} takes as input $h_1, \dots, h_{(n+1)q}$ in order to simulate all RO queries. For a RO query $H_1(\text{pk}, lr, i)$, \mathcal{C} first enumerates all $i' \in [n]$ and assigns $h_{(\text{ctr}_h-1)(n+1)+i'}$ to $H_1(\text{pk}, lr, i')$. Then, \mathcal{C} computes the nonce R for lr and assigns $h_{\text{ctr}_h(n+1)}$ to $H_2(\text{pk}, lr.\text{msg}, R)$ if it is not assigned any value yet. Similarly, for a new RO query $H_1(\text{pk}, M, R)$, its value is set to $h_{\text{ctr}_h(n+1)}$. The rest follows by a similar analysis.

Proof of Lemma 10.4.7. We first construct an algorithm \mathcal{C} following the syntax of the algorithm described in Lemma 9.3.4 and then construct \mathcal{B} from $\text{Fork}^{\mathcal{C}}$. The input of \mathcal{C} consists of par that defines a linear hash function $(\mathcal{S}, \mathcal{D}, \mathcal{R}, F)$ and uniform random elements $h_1, \dots, h_{n \cdot q} \in \mathcal{S}_{\text{hash}}$. Similarly to the proof of Lemma 10.4.4, \mathcal{C} can oracles CHAL and PI oracle and at the beginning, start with, \mathcal{C} makes $2q_s$ queries to CHAL and denotes the challenges as $U_1, V_1, \dots, U_{q_s}, V_{q_s} \in \mathcal{D}$. Then, \mathcal{C} initializes all the states st_0, \dots, st_n as in the game $\text{TS-SUF-3}^{\text{FROST}^1\text{-H}}$, and initializes the counters $\text{ctr}_s, \text{ctr}_h$ to 0 and the function dt to an empty table. \mathcal{C} also initializes ctrPP to an empty table, which are used to record the counter corresponding to each token generated by honest parties. Then, \mathcal{C} runs \mathcal{A} with access to the oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}, \widetilde{\text{RO}}$, which are simulated as follows. In the following description, we use i to denote the index of parties, j to denote the index of $U_1, V_1, \dots, U_{q_s}, V_{q_s}$, and k to denote the index of $h_1, \dots, h_{n \cdot q}$.

$\widetilde{\text{INIT}}(CS)$: \mathcal{C} initializes H to an empty table and samples a_0, \dots, a_{t-1} uniformly from \mathcal{D}_{key} . Define $f(x) := \sum_{i=0}^{t-1} a_i x^i$. Then, \mathcal{C} sets $\text{pk} \leftarrow F(a_0)$, $\text{pk}_i \leftarrow F(f(x_i))$ for $i \in [n]$, and $\text{sk}_i \leftarrow f(i)$ for $i \in CS$. Finally, \mathcal{C} returns $\text{pk}, \text{aux} = (\text{pk}_1, \dots, \text{pk}_n), \{\text{sk}_i\}_{i \in CS}$.

$\widetilde{\text{RO}}$ query $H_1(x)$: If $H_1(x) \neq \perp$, \mathcal{C} returns $H_1(x)$. Otherwise, \mathcal{C} parses x as $(\tilde{\text{pk}}, lr, \tilde{i})$ for some $\tilde{i} \in [1..n]$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_1(x) \leftarrow_{\$} \mathcal{S}_{\text{hash}}$ and returns $H_1(x)$. Otherwise, \mathcal{C} increases ctr_h by 1 and sets $H_1(\text{pk}, lr, i) \leftarrow h_{n(\text{ctr}_h-1)+i}$ for each $i \in [n]$. In addition, let $\text{mapLR}(\text{ctr}_h) := lr$. Then, \mathcal{C} computes $R \leftarrow \sum_{i \in lr.\text{SS}} (R_i + S_i^{d_i})$, where $(R_i, S_i) \leftarrow lr.\text{PP}(i)$ and $d_i = H_1(\text{pk}, lr, i)$. If $H_2(\text{pk}, lr.\text{msg}, R) = \perp$, \mathcal{C} sets $H_2(\text{pk}, lr.\text{msg}, R) \leftarrow_{\$} \mathbb{Z}_p$. Finally, \mathcal{C} returns $H_1(x)$.

$\widetilde{\text{RO}}$ query $H_2(x)$: If $H_2(x) = \perp$, \mathcal{C} sets $H_2(x) \leftarrow_{\$} \mathcal{S}_{\text{hash}}$. Then, \mathcal{C} returns $H_2(x)$.

$\widetilde{\text{PPO}}(i)$ query: Same as in the game $\text{TS-SUF-3}^{\text{FROST1-H}}$, except in the simulation of algorithm SPP, \mathcal{C} first increases ctr_s by 1 and sets $pp \leftarrow (U_{\text{ctr}_s}, V_{\text{ctr}_s})$, $\text{st}_i.\text{mapPP}(pp) \leftarrow (0, 0)$, and $\text{ctrPP}(i, pp) \leftarrow \text{ctr}_s$.

$\widetilde{\text{PSIGNO}}(i, lr)$ query: Same as in the game $\text{TS-SUF-3}^{\text{FROST1-H}}$, except in the simulation of algorithm PS, if $\text{st}_i.\text{mapPP}(pp) \neq \perp$, \mathcal{C} computes

$$z_i \leftarrow \text{PI}(U_j + d_i V_j) + c \cdot \lambda_i^{\text{lr}.SS} \cdot f(i),$$

where $j \leftarrow \text{ctrPP}(i, pp)$. In addition, \mathcal{C} sets $\text{dt}(j) \leftarrow (k, d_i, z_i - c\lambda_i^{\text{lr}.SS} \cdot f(i))$, where k denotes the index such that $H_1(\text{pk}, lr, i)$ is set to h_k during the simulation.

After receiving the output $(m^*, \text{sig}^* = (R^*, z^*))$ from \mathcal{A} , \mathcal{C} returns (\perp, \perp, \perp) if E_2 does not occur. Otherwise, we know $\text{curSSL}(lr^*) > 0$ and $\text{curSSL}(lr^*) \neq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$. Therefore, there exists k^* and i^* such that $i^* \in \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\} \setminus \text{curSSL}(lr^*)$ and $\text{mapLR}(k^*) = lr^*$. (Since $\text{curSSL}(lr^*) \subseteq \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$, we must have $\{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\} \setminus \text{curSSL}(lr^*) \neq \emptyset$.) Since $i^* \in \{i \in HS \cap lr^*.SS : lr^*.PP(i) \in PP_i\}$, there exists $j^* \in [1..q_s]$ such that $lr^*.PP(i^*) = (U_{j^*}, V_{j^*})$. If $\text{dt}(j^*) = \perp$, \mathcal{C} sets $J \leftarrow \perp$. Otherwise, let $(k, d, z) \leftarrow \text{dt}(j^*)$ and \mathcal{C} sets $J = k$. Then, \mathcal{C} returns $(n(k^* - 1) + i^*, J, \text{Out})$, where Out consists of all variables received or generated by \mathcal{C} , including i^*, j^*, k^*, lr^* .

ANALYSIS OF \mathcal{C} . To use Lemma 9.3.4, we define IG as the algorithm that runs $\text{PGen}(1^\kappa)$ and outputs par . The output J is either \perp or in $[1..(n \cdot q)]$. It is not hard to see that \mathcal{C} simulates the game $\text{TS-SUF-3}^{\text{FROST1-H}}$ perfectly, which implies $\text{acc}(\mathcal{C}) \geq \Pr[E_2]$, where $\Pr[E_2]$ refers to the probability in the original $\text{TS-SUF-3}^{\text{FROST1-H}}$ game with \mathcal{A} (as in the lemma statement). By Lemma 9.3.4,

$$\text{acc}(\text{Fork}_2^{\mathcal{C}}) \geq \frac{\Pr[E_2]^2}{n \cdot q(n \cdot q + 1)} \leq \frac{\Pr[E_2]^2}{2n^2q^2}.$$

CONSTRUCT \mathcal{B} FROM $\text{Fork}_2^{\mathcal{C}}$. We now give a construct of the AOMPR adversary \mathcal{B} using $\text{Fork}_2^{\mathcal{C}}$, and the available CHAL and PI oracles. To start with, \mathcal{B} receives par from the $\text{AOMPR}^{\text{LHF}}$ game and runs $\text{Fork}_2^{\mathcal{C}}(par)$. All the CHAL queries from the first execution of \mathcal{C} are relayed by \mathcal{B} to its own CHAL oracle, and for all the CHAL queries from the second execution of \mathcal{C} , \mathcal{B} answers them with the same challenges as the first execution. All the PI queries from $\text{Fork}_2^{\mathcal{C}}$ are relayed by \mathcal{B}

to its own PI oracle. Without loss of generality, we can assume all the challenges are different, since otherwise, \mathcal{B} can solve them trivially. Denote the event BadHash as $h_I \neq h'_I$, where I are outputted by the first execution of \mathcal{C} . Since h_I, I are independent of h'_I , we know $\Pr[\text{BadHash}] \leq 1/|\mathcal{S}_{\text{hash}}| \leq 1/2^\kappa$. Then, we can conclude the proof with the following claim, which implies

$$\text{Adv}_{\text{LHF}}^{\text{aompr}}(\mathcal{B}) \geq \text{acc}(\text{Fork}_2^{\mathcal{C}}) - \Pr[\text{BadHash}] \geq \frac{\Pr[E_2]^2}{2n^2q^2} - 1/2^\kappa .$$

□

Claim 10.4.8. \mathcal{B} can win the game $\text{AOMPR}_{\text{LHF}}$, if $\text{Fork}^{\mathcal{C}}$ returns $(I, \text{Out}, \text{Out}')$ and BadHash does not occur.

Proof. We directly use the notations in the description of \mathcal{C} to denote the variables in Out and use $(\cdot)'$ to denote the variables in Out'. The total number of the CHAL queries is $2q_s$ and the corresponding challenges are $U_1, V_1, \dots, U_{q_s}, V_{q_s}$.

We now show how to compute u_j, v_j for each $j \in [q_s]$ such that $F(u_i) = U_i$ and $F(v_i) = V_i$. There are the following cases.

Case 0: Both $\text{dt}(j)$ and $\text{dt}'(j)$ are \perp . In this case, \mathcal{B} computes u_j, v_j by directly querying oracle $\text{PI}(U_j)$ and $\text{PI}(V_j)$.

Case 1: Exactly one of $\text{dt}(j)$ and $\text{dt}'(j)$ is not \perp . Without loss of generality, assume $\text{dt}(j) = (k, d, z)$, which implies $F(z) = U_j + dV_j$. \mathcal{B} computes v_j by directly querying oracle $\text{PI}(V_j)$ and computes $u_j \leftarrow z - d \cdot v_j$.

For all the following cases, both $\text{dt}(j)$ and $\text{dt}'(j)$ are not \perp and we denote $(k, d, z) \leftarrow \text{dt}(j)$ and $(k', d', z') \leftarrow \text{dt}'(j)$.

Case 2: $d \neq d'$. In this case, \mathcal{B} computes $v_j = \frac{z-z'}{d-d'}$, $u_j = z - d \cdot v_j$.

Case 3: $d = d'$. In this case, \mathcal{B} computes v_j, u_j the same as Case 1. Also, since $d = d'$, \mathcal{B} queries PI oracle only once in order to answer queries $\text{PI}(U_j + dV_j)$ and $\text{PI}(U_j + d'V_j)$ from $\text{Fork}^{\mathcal{C}}$.

From the execution of \mathcal{C} , we know $\text{dt}(j) = (k, d, z) \neq \perp$ if and only if \mathcal{C} queries PI on $(U_j + dV_j)$. Therefore, denote $(U_j + dV_j)$ as the PI query associated with $\text{dt}(j)$. For all the above cases, \mathcal{B}

queries PI oracle twice for simulating PI queries associated with $\text{dt}(j)$ and $\text{dt}'(j)$ and computing u_j, v_j .

We now show how to compute u_{j^*} and v_{j^*} . From the execution of $\text{Fork}_2^{\mathcal{C}}$, we know $\text{pk} = \text{pk}'$ and $\text{mapLR}(k) = \text{mapLR}'(k)$ for all $k \leq I$, which implies $lr^* = \text{mapLR}(I) = \text{mapLR}'(I) = lr^{*'}$. Since E_2 occurs in both executions of \mathcal{C} , we know $\text{SVf}(\text{pk}, lr^*, (R^*, z^*)) = \text{true}$ and $\text{SVf}(\text{pk}, lr^*, (R^{*'}, z^{*'})) = \text{true}$ are valid. Therefore, $F(z^*) = R^* + F(a_0c)$, $R^* = \sum_{i \in lr^*} \text{SS}(R_i + d_i S_i)$, $g^{z^{*'}} = R^{*'} + F(a_0c')$, $R^{*'} = \sum_{i \in lr^{*'}} \text{SS}(R_i + d'_i S_i)$, where $(R_i, S_i) = lr.\text{PP}(i)$, $c = H_2(\text{pk}, M^*, R^*)$, $c' = H_2(\text{pk}, M^*, R^{*'})$, and $d_i = H_1(\text{pk}, lr^*, i)$, $d'_i = H_1(\text{pk}, lr^*, i)$. Since for each $i \neq i^*$ we have $d_i = h_{n(k^*-1)+i} = d'_i$, we have

$$F(z^* - z^{*'}) = R^* - R^{*'} + F(a_0(c - c')) = (d_{i^*} - d'_{i^*})S_{i^*} + F(a_0(c - c')) .$$

Therefore, \mathcal{C} can compute $v_{j^*} = \frac{z^* - z^{*'} - a_0(c - c')}{d_{i^*} - d'_{i^*}}$. If $J = \perp$, \mathcal{B} computes u_{j^*} by querying $\text{PI}(U_{j^*})$ directly. In this case, \mathcal{B} queries PI only once to compute u_{j^*} and v_{j^*} . If $J \neq \perp$, let $(k, d, z) \leftarrow \text{dt}(j^*)$ and $(k', d', z') \leftarrow \text{dt}'(j^*)$. Then, \mathcal{B} computes $u_{j^*} = z - d \cdot v_{j^*}$. Since $i^* \notin S_2(lr^*)$, we know $k \neq I$. (Otherwise, suppose $k = I$. Since $I = n(k^* - 1) + i^*$ and $\text{mapLR}(k^*) = lr^*$, we know a $\text{PSIGNO}(i^*, lr^*)$ is made and does not return \perp during the simulation, which implies $i^* \in S_2(lr^*)$.) Thus, we have $k' = J = k \neq I$ and $d = h_J = d'$, which means \mathcal{B} only needs to query PI once to simulate the PI queries associated with $\text{dt}(j^*)$ and $\text{dt}'(j^*)$. Therefore, the total number of PI queries made by \mathcal{B} is equal to $2q_s - 1$, which implies \mathcal{B} wins the game $\text{AOMDL}^{\text{LHF}}$. \square

Part III

LATTICE-BASED THRESHOLD SIGNATURES

Chapter 11

THE ALGEBRAIC ONE-MORE MISIS PROBLEM AND LATTICE-BASED THRESHOLD SIGNATURES: INTRODUCTION

One-more assumptions enable proofs of security for several interactive protocols, most notably identification schemes, threshold signatures, multi-signatures, and blind signatures. The perhaps most prominent example is the *one-more discrete logarithm assumption* (OMDL) [16], which requires the hardness of computing Q discrete logarithm instances given access to an oracle that allows the adversary to compute $Q - 1$ discrete logarithms of *arbitrary* group elements. It has been used throughout several security proofs (e.g., cf. [16, 65, 105, 14]), where the oracle allows the reduction to simulate secret-key dependent behavior of honest parties without knowing the secret key. This notwithstanding, the main point of controversy is that an assumption such as OMDL is very strong—for instance, Koblitz and Menezes [89] point out that in certain groups, it is easier to break OMDL than to solve the standard discrete logarithm problem. The only available route justifying its plausible hardness on standard elliptic curves is a proof [12] in the generic-group model (GGM) [118, 101].

PROVABLY-HARD ONE-MORE PROBLEMS. A few recent works [122, 9] follow a different path: While they still leverage the power of one-more problems as an intermediate *interface* to design modular security proofs, they also instantiate the underlying mathematical structure to support a proof that the one-more problem is indeed hard based on more standard assumptions, such as the hardness of (standard) discrete log/RSA (in [122]) or of DDH (in [9]).

This is the angle pursued by this paper—we seek one-more problems which are sufficiently *expressive* to enable useful security proofs, while also enjoying *provable* hardness from standard assumptions. This paper deals specifically with *lattice-based cryptography*. Espitau, Katsumata, and Takemure (EKT) [60] recently introduced a one-more problem they refer to as *algebraic one-more MLWE* (AOM-MLWE) and show that its hardness yields the security of a two-round threshold signature. They also establish the *selective* hardness of AOM-MLWE from the standard MLWE and MSIS assumptions, which is however unfortunately not sufficient to support their application to threshold signatures. Another example—less relevant for this work, however—is the one-more

ISIS problem [7], used in the construction of lattice-based blind signatures, which has also only been validated via cryptanalysis.

OUR CONTRIBUTIONS, IN A NUTSHELL. We introduce a new variant of AOM-MLWE which we refer to as *Algebraic One-More MISIS* (AOM-MISIS, for short),¹ and for which we show two different types of results:

- **Provability.** We show that the hardness of AOM-MISIS follows from the hardness of MSIS and MLWE, with suitable parameters. In fact, the hardness of AOM-MISIS *implies* the hardness of AOM-MLWE, and thus our result carries over to AOM-MLWE, providing the first reduction of AOM-MLWE to standard assumptions, which was left as a main open question in [60].
- **Expressivity.** We show the security of the EKT threshold signature scheme from [60], as well as of the recent scheme by Chairattana-Apirom, Tessaro, and Zhu [41], assuming the hardness of AOM-MISIS. In turn, this establishes the security of these schemes from MSIS and MLWE. We obtain either better concrete security bounds compared to [41], or proofs under weaker assumptions compared to [60] for slightly larger concrete parameter sets. We also give a proof that the EKT scheme satisfies the strongest notion of security in the hierarchy of Bellare *et al.* [14].

ALGEBRAIC ONE-MORE MISIS. The definition of AOM-MISIS relies on the cyclotomic ring $R = \mathbb{Z}[X]/(X^N + 1)$, where N is a power of two, as well as the associated ring $R_q = R/qR \cong \mathbb{Z}_q[X]/(X^N + 1)$ for an odd prime q . The problem is defined via the following game:

- **Input.** The adversary is initially given a matrix $A = [\bar{A} | \mathbb{I}_k] \in R_q^{k \times m}$, where $\bar{A} \leftarrow_{\$} R_q^{k \times (m-k)}$, as well as Q instances $\mathbf{t}_i \leftarrow A \mathbf{s}_i$, for $i \in [Q]$, where $\mathbf{s}_i \leftarrow_{\$} \mathcal{D}_{\sigma_i}^m$ is “small”, and sampled from an m -dimensional discrete Gaussian with parameter σ_i . The instance number Q is a parameter of the game.
- **Oracle access.** The adversary can also (adaptively) query an oracle PI which takes as input a Q -dimensional vector $\mathbf{b} \in R^Q$, and returns $\sum_{i \in [Q]} b_i \mathbf{s}_i$.

¹Our naming is due to the fact that we prefer to think of our problem as a one-more version of Inhomogenous SIS (SIS), rather than of LWE.

To win, the adversary needs to output both $\hat{\mathbf{b}} = (\hat{b}_1, \dots, \hat{b}_Q) \in R^Q$ and a “short” solution $\hat{\mathbf{s}} \in R^m$ such that

$$\sum_{i \in [Q]} \hat{b}_i \mathbf{t}_i = A \hat{\mathbf{s}} .$$

To exclude trivial winning strategies, $\hat{\mathbf{b}}$ must however not be in the span of the vectors $\mathbf{b}_1, \mathbf{b}_2, \dots$ queried to PI. We in fact require something stronger, namely that in order to win, the adversary also needs to additionally output $\mathbf{u} \in R^Q$, along with $\hat{\mathbf{b}}$ and $\hat{\mathbf{s}}$, and the following two properties need to be satisfied:

- The vector \mathbf{u} is orthogonal to all vectors $\mathbf{b}_1, \mathbf{b}_2, \dots$ queried to PI, but not to $\hat{\mathbf{b}}$, i.e., $\mathbf{b}_i^\top \mathbf{u} = 0$, but $\hat{\mathbf{b}}^\top \mathbf{u} \neq 0$. Here, orthogonality is defined with respect to the ring R .
- Both vectors $(\hat{b}_1 \cdot \sigma_1, \dots, \hat{b}_Q \cdot \sigma_Q)$ and $(u_1/\sigma_1, \dots, u_Q/\sigma_Q)$ are appropriately small, when interpreted as real vectors. Both conditions are necessary, as otherwise a simple attack exists.

Our main result is a concrete reduction showing that the AOM-MISIS problem is indeed hard if the Module LWE (MLWE) and Module SIS (MSIS) assumptions also hold with suitable parameters, related to the parameters of the above game. Our proof relies on a novel *generalization* of the technique by Tessaro and Zhu [122], which was originally used to obtain hard one-more problems based on linear hash functions. The origin of this technique goes back to the analysis of Okamoto signatures [107], and was used also in [80]. Some prior works [122, 81] can be interpreted as attempts to adapt this approach to the lattice setting in limited ways, but in our context, they would lead to worse parameters and restrictions on the generality of our game. We discuss further details in Section 11.1 below.

We note that the hardness of AOM-MISIS implies the hardness of AOM-MLWE, and thus we resolve the main open problem from [60], which only provided an analysis for *selective* adversaries issuing their oracle queries beforehand. In fact, the main difference between AOM-MISIS and AOM-MLWE is the winning condition. We require outputting a single short solution for a suitable non-trivial linear combination of the challenges, whereas AOM-MLWE requires outputting Q short solutions for all challenges, given $Q - 1$ access to a similar oracle, although with different conditions on the queries. These changes are what in part enables simpler reductions for threshold signatures. A more detailed discussion is provided in Section 12.2.

APPLICATIONS TO THRESHOLD SIGNATURES. We first recall that in a *t-out-of-n threshold signature scheme* [58, 59], n potential signers each hold a secret share of a secret signing key, with an associated public verification key. Any subset of (at least) t of these signers is able to jointly produce a signature, via interaction, whereas an adversary that controls fewer than t signers should not be able, on its own, to come up with a valid signature.

We leverage AOM-MISIS to obtain tighter analyses for state-of-the-art *two-round* lattice-based threshold signatures, based on the MSIS/MLWE assumptions, and without assuming the hardness of any ad-hoc one-more problem. We focus on two threshold signatures, which we refer to as CATZ [41] and EKT [60]. There are several small differences between the two constructions, despite the fact that they instantiate similar ideas. Both can be seen as a natural threshold version of the Fiat-Shamir-with-abort paradigm [97] that also underlies DILITHIUM [100], albeit dispensing with the actual abort, and using a sufficiently large modulus instead. In particular, EKT is a threshold version of Raccoon [55], submitted to the additional NIST call for post-quantum signatures [103]. They are natural lattice analogues of FROST [91, 14, 49], a very lightweight threshold Schnorr signature.

Taking some liberty from their formal description, in both schemes, as a result of the second round, the i -th signer produces an affine signature share

$$z_i = r_i + c \cdot \lambda_i ss_i,$$

in a suitable algebraic structure, where ss_i is the i -th signer's key share, c is a hash value associated with the signature, and λ_i is a linear reconstruction coefficient associated with the secret sharing scheme. With S being the set of signers, the final signature has format (\mathbf{R}, \mathbf{z}) , where $\mathbf{R} = A \sum_{i \in S} r_i$, (A is a public matrix) and $\mathbf{z} = \sum_{i \in S} z_i$. However, the two schemes differ in the following aspects:

- In CATZ, to avoid leakage of the secret shares, $\lambda_i ss_i$ needs to remain sufficiently small and thus a secret sharing scheme with small reconstruction coefficients is needed. This incurs a significant cost due to the larger share size when compared to Shamir secret sharing. The security proof in [41] is fairly involved, and gives a direct reduction to MSIS.
- To enable the use of Shamir secret sharing, EKT changes the initial setup to ensure that any pair of signers shares a secret key, and these keys are used, in each execution of the signing protocol, to generate (pseudo)random masks $(msk_i)_{i \in S}$ such that $\sum_{i \in S} msk_i = 0$. The i -th

signer then sends $z_i + \text{msk}_i$ in the second round instead of z_i , and this ensures that only the *sum* of the z_i 's is leaked. The security proof in [60] relies on the direct use of the AOM-MLWE assumption, and since the constructed adversary is inherently not selective, they need to rely on a conjecture about the hardness of AOM-MLWE.

In Sections 13.1 and 13.2, we give new security analyses for both schemes based on direct reductions from AOM-MISIS, which in turn yield concrete security proofs from MLWE and MSIS. Our result for CATZ yields better parameters than the analysis of [41]. The result about EKT, in addition to now basing security on MLWE/MSIS alone, also shows a strong security property for this scheme, namely TS-UF-4 in the hierarchy of Bellare et al. [14]. The concrete efficiency numbers derived from our bounds are somewhat worse than those from EKT, but their concrete analysis relies on their own cryptanalysis of AOM-MLWE, whereas we use standard parameters for MLWE and MSIS. Closing the gap between the cryptanalysis-driven parameter choices and those derived from our security reduction remains an interesting open problem.

The recent work on Ringtail [33] also proposes a threshold signature scheme similar in spirit to EKT and with a proof of security under MLWE. In Ringtail, unlike EKT, signers need to know already in the first round the set of involved signers S , and thus, unlike EKT and CATZ, Ringtail is not partially non-interactive in the sense of Bellare *et al.* [14]. Ringtail also needs to authenticate first-round messages. The authors of Ringtail mention removing these restrictions, while preserving comparable efficiency and provability from MLWE/MSIS, is an open problem, which we resolve here based on the EKT scheme.

Ringtail's security bound degrades with the number of random oracle queries, whereas ours degrades with the number of signing sessions. We note that the latter is a system parameter that can be enforced, whereas the former scales with the running time of the adversary.

REMARKS ON CATZ. While the concrete efficiency of CATZ is significantly worse than that of EKT, we see value in the CATZ construction because it does not rely on pairwise masks. CATZ closely resembles the structure of the original FROST protocol, making it a promising starting point for achieving identifiable aborts (constructing efficient partially non-interactive lattice-based threshold signatures with identifiable aborts is a big open problem in the field). In contrast, achieving identifiable aborts in EKT seems less likely without relying on heavyweight NIZK proofs, due to its reliance on pairwise masks.

Also, the security argument of CATZ works for a class of linear secret sharing schemes rather

than a specific scheme. Its main source of inefficiency stems from the underlying secret sharing scheme, and its efficiency could be significantly improved if better instantiations of the underlying secret sharing scheme are proposed.

OTHER RELATED WORK. We note here that there are other approaches to threshold signatures. First off, *Fully-Homomorphic Encryption* (FHE) generically yields round-optimal threshold signatures [30, 29, 8]. These require however the homomorphic evaluation of the signing algorithm, and thus come with a substantial computational overhead. Earlier works [51, 48] proposed two-round n -out-of- n threshold signatures derived from constructions for the related notion of multi-signatures. Gur *et al.* [79] proposed two-round construction based on linearly homomorphic encryption (LHE) which supports arbitrary thresholds. Both rounds are message-dependent. More recently, Pino *et al.* [54] propose a more efficient lattice-based threshold signature scheme that does not rely on FHE or the aforementioned heavy primitives, but the drawback is that the protocol has three message-dependent rounds. Recent work [87] also gives a five-round threshold signature with adaptive security based on MLWE/MSIS, whereas a threshold version of Falcon [112] was presented in [61]. The latter scheme is designed for robustness but requires four message-dependent rounds and incurs quadratic communication complexity in the threshold.

11.1 Technical Overview

The main goal of this section is to provide a detailed overview of our main result establishing the hardness of AOM-MISIS based on MSIS and MLWE. We stress that our reduction will involve concrete parameters, but we keep the discussion in this section somewhat qualitative on this front.

At a technical level, it helps to see AOM-MISIS as a lattice-based analogue of the AOMPR framework proposed by Tessaro and Zhu (TZ) [122] to define one-more problems for linear hash functions. Their framework does not cover lattice problems, however—we aim to adapt it to lattices, and this presents a number of challenges, which we explain in this section. We will also discuss how prior works have already tried to overcome such challenges, how those approaches are not sufficient for our purposes, and what we do instead.

We focus first on a variant of AOM-MISIS we denote as wAOM-MISIS, where the adversary does not output the relaxation vector $\hat{\mathbf{b}}$ and is instead asked to output Q solutions $\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_Q$ with small norm (instead of a single solution) such that $\mathbf{t}_i = A\hat{\mathbf{s}}_i$ for all $i \in [Q]$. Also, the constraint $\hat{\mathbf{b}}^T \mathbf{u} \neq 0$ is now changed to $\mathbf{u} \neq 0$ instead, while $\mathbf{b}_i^T \mathbf{u} = 0$ still must hold for the queries to PI.

This problem is not easier than AOM-MISIS: If the adversary wins the wAOM-MISIS game, there exists an index $i \in [Q]$ such that $u_i \neq 0$, and thus the adversary can win the original game by outputting $(\hat{s}_i, e_i, \mathbf{u})$, where e_i is the i -th unit vector. This is also the idea underlying the proof that AOM-MISIS hardness implies the hardness of AOM-MLWE in Section 12.2.

Reduction ideas from [122].

We now provide a self-contained review of TZ's proof framework in the context of wAOM-MISIS, and explain where it fails. In particular, given an adversary \mathcal{A} for the wAOM-MISIS game, we build a simple MSIS adversary \mathcal{B} that takes a random matrix $A \in R_q^{k \times m}$ as input and outputs a vector \mathbf{z} with small norm such that $A\mathbf{z} = 0$. To start with, \mathcal{B} runs \mathcal{A} by simulating the wAOM-MISIS game faithfully using the matrix A and sampling \mathbf{s}_i by itself. After receiving $(\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_Q, \mathbf{u})$, if \mathcal{A} wins, \mathcal{B} finds an index $i \in [Q]$ such that $\hat{\mathbf{s}}_i \neq \mathbf{s}_i$ and outputs $\mathbf{z} \leftarrow \hat{\mathbf{s}}_i - \mathbf{s}_i$. Otherwise, \mathcal{B} aborts.

ANALYSIS OF \mathcal{B} . It is not hard to see that \mathcal{B} wins the MSIS game if such an index i indeed exists, since $A\mathbf{z} = A\hat{\mathbf{s}}_i - A\mathbf{s}_i = 0$ and $\|\mathbf{z}\| \leq \|\mathbf{s}_i\| + \|\hat{\mathbf{s}}_i\|$ is bounded given both $\|\mathbf{s}_i\|$ and $\|\hat{\mathbf{s}}_i\|$ are bounded. Here $\|\cdot\|$ denotes the ℓ_2 -norm. The hard part is to show that such i exists w.h.p. under the assumption that \mathcal{A} wins w.h.p., and this is what we discuss next.

Following the same lines as TZ, we can assume w.l.o.g. that \mathcal{A} is deterministic, and thus the view $\text{View}_{\mathcal{A}}(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ of \mathcal{A} (assuming the matrix A is fixed) is completely determined by the challenge secrets $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Also, we denote by \mathbf{D} the set of secrets $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ such that \mathcal{A} wins the wAOM-MISIS game.

While this is not the case for us, suppose for now that \mathbf{D} is finite and each $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ is sampled uniformly from \mathbf{D} . The first step is to find a *derangement* Φ of \mathbf{D} , i.e., a permutation of \mathbf{D} without fixed points, such that for any $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \in \mathbf{D}$, $\text{View}_{\mathcal{A}}(\mathbf{s}_1, \dots, \mathbf{s}_Q) = \text{View}_{\mathcal{A}}(\mathbf{s}'_1, \dots, \mathbf{s}'_Q)$ with $(\mathbf{s}'_1, \dots, \mathbf{s}'_Q) = \Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Therefore, \mathcal{A} also produces the same output $(\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_n, \mathbf{u})$ in both cases. Since $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \neq (\mathbf{s}'_1, \dots, \mathbf{s}'_Q)$, there exists an index $i \in [Q]$ such that $\mathbf{s}_i \neq \mathbf{s}'_i$ and thus either $\mathbf{s}_i \neq \hat{\mathbf{s}}_i$ or $\mathbf{s}'_i \neq \hat{\mathbf{s}}_i$, which means \mathcal{B} wins the MSIS game in at least one of the cases. Since Φ is a permutation of \mathbf{D} , for at least half of elements in \mathbf{D} lead to \mathcal{B} winning. Therefore, \mathcal{B} wins with at least half of the probability that \mathcal{A} wins. Crucially, note that Φ does not need to be efficient for this argument to succeed, as the algorithm \mathcal{B} described above is the actual reduction.

CONSTRUCTING Φ . The first idea is to define Φ such that

$$\Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q) := (\mathbf{s}_1 + u_1 \cdot \Delta, \dots, \mathbf{s}_Q + u_Q \cdot \Delta), \quad (11.1)$$

where $\Delta \in R^m$ is a non-zero vector such that $A\Delta = 0$ (we will discuss below how to ensure such Δ exists), and \mathbf{u} is the vector output by \mathcal{A} in an execution with the secret value $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. It is not hard to check that $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ and $\Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ produce the same view. First off, since $A(\mathbf{s}_i + u_i \cdot \Delta) = A\mathbf{s}_i + u_i \cdot (A\Delta) = A\mathbf{s}_i$, the input of \mathcal{A} is identical in both cases. Further, for each PI query \mathbf{b} made by \mathcal{A} , since \mathbf{u} is a non-zero vector satisfying $\sum_{i \in [Q]} u_i b_i = 0$, we have $\sum_{i \in [Q]} b_i (\mathbf{s}_i + u_i \cdot \Delta) = \sum_{i \in [Q]} b_i \mathbf{s}_i + (\sum_{i \in [Q]} b_i u_i) \cdot \Delta = \sum_{i \in [Q]} b_i \mathbf{s}_i$, which means the response received by \mathcal{A} is identical in both cases.

ISSUES AND PRIOR APPROACHES. The issue with the above approach is that in our setting the secrets are not sampled uniformly from a set. Two prior works [122, 81] have overcome this issue by sampling \mathbf{s}_i uniformly from a bounded box $\mathcal{B}_{\beta_i}^m := \{\mathbf{x} \in R^m \mid \|\mathbf{x}\| \leq \beta_i\}$. Still, the challenge now is that Φ is no longer a permutation over a subset of $\mathcal{B}_{\beta_1}^m \times \dots \times \mathcal{B}_{\beta_Q}^m$, since $\|\mathbf{s}_i + u_i \cdot \Delta\|$ can be larger than β_i even though $\|\mathbf{s}_i\| \leq \beta_i$.

To overcome this issue, Hauck *et al.* [81]² proposed to set β_i to be very large such that the total fraction of $\mathbf{s}_i \in \mathcal{B}_{\beta_i}^m$ such that $\mathbf{s}_i + u_i \cdot \Delta$ falls outside $\mathcal{B}_{\beta_i}^m$ is negligible in the security parameter κ . Then, one can still show \mathcal{B} wins with nearly at least half of the winning probability of \mathcal{A} . However, this requires $\sigma_i = \Omega(2^\kappa \|u_i \Delta\|)$, which results in very inefficient constructions using this approach.

Chairattana-Apirom *et al.* [41] overcame the 2^κ barrier with a different approach. They let \mathbf{s}_1 be sampled from $\mathcal{B}_{\beta_1}^m$, whereas \mathbf{s}_i is sampled from $\mathcal{D}_{\sigma_i}^m$. They use this in the security proof of their threshold signature scheme based on the MSIS assumption, but their technique can be massaged into a reduction from MSIS to a restricted version of the wAOM-MISIS game, where the PI queries are restricted, but still sufficient for their application.

Their proof is in some sense a probabilistic relaxation of TZ's, in that for any $\mathbf{s}_1 \in R^m$, they consider a *random variable* $\text{View}_{\mathcal{A}}(\mathbf{s}_1)$ which represents the view of \mathcal{A} given the first secret is \mathbf{s}_1 , whereas the remaining secrets are sampled afresh from discrete Gaussian distributions. If for any $\mathbf{s}_1 \neq \mathbf{s}'_1 \in R^m$, such that $A\mathbf{s}_1 = A\mathbf{s}'_1$, we can show that $\text{View}_{\mathcal{A}}(\mathbf{s}_1)$ and $\text{View}_{\mathcal{A}}(\mathbf{s}'_1)$ are identically distributed, then we can once again carry out the same argument as above. They in fact generalize this argument by showing that if the Rényi divergence of $\text{View}_{\mathcal{A}}(\mathbf{s}_1)$ from $\text{View}_{\mathcal{A}}(\mathbf{s}'_1)$ is small, one

²We note that [81] precedes [122], but deals with a lattice-analogue of a linear-hash function based framework introduced by [80] similar in spirit to that of [122].

can still show that the winning probabilities of \mathcal{B} and \mathcal{A} are sufficiently related. Finally, they show the Rényi divergence is indeed bounded for the special types of queries they consider. Doing so, they ensure that σ_i depends linearly on \sqrt{Q} instead of 2^κ .

Still, their reduction fails if \mathcal{A} issues general PI queries, and this is because the Rényi divergence of $\text{View}_{\mathcal{A}}(s_1)$ from $\text{View}_{\mathcal{A}}(s'_1)$ is no longer bounded. For example, the first secret (either s_1 or s'_1) can be recovered from an PI query with input $(1, 0, \dots, 0)$. Moreover, they need to ensure that almost all s_1 's have one partner $s'_i \neq s_i \in \mathcal{B}_{\beta_1}^m$ such that $As_1 = As'_1$, and this requires $\beta_1 \geq 2^{\kappa/N} q^{k/m}$,³ which leads to inefficient parameters in their threshold signature.

11.1.1 Step 1: Generalizing TZ's argument

Our goal is to provide a reduction that supports an adversary \mathcal{A} without any extra restrictions on its queries to PI, and furthermore that ensures hardness even for sufficiently small parameter choices in wAOM-MISIS. For this reason, we take a different route and generalize TZ's argument to a setting where each s_i is sampled (independently) from an arbitrary distribution \mathcal{P}_i , rather than uniformly from a finite set.

Let $\mathcal{P} = \mathcal{P}_1 \times \dots \times \mathcal{P}_Q$ be the joint distribution of (s_1, \dots, s_Q) . Then, instead of requiring Φ to be a permutation of \mathbf{D} (recalling that \mathbf{D} is the set of secrets (s_1, \dots, s_Q) that make \mathcal{A} win the wAOM-MISIS game), we require that the distribution of (s_1, \dots, s_Q) conditioned on $(s_1, \dots, s_Q) \in \mathbf{D}$ (denoted as $\mathcal{P}_{|\mathbf{D}}$) is identical to that of $\Phi(s_1, \dots, s_Q)$ conditioned on $(s_1, \dots, s_Q) \in \mathbf{D}$ (denoted as $\Phi(\mathcal{P}_{|\mathbf{D}})$). The original TZ approach corresponds to the case when \mathcal{P} is a uniform distribution over \mathbf{D} . Here, since arbitrary distributions over \mathbf{D} are allowed, we no longer require \mathbf{D} to be a finite set. The two other requirements for Φ remain the same, i.e., Φ has no fixed point over \mathbf{D} and $\text{View}_{\mathcal{A}}(s_1, \dots, s_Q) = \text{View}_{\mathcal{A}}(\Phi(s_1, \dots, s_Q))$.

Our key observation here is that the previous argument to lower bound the success probability of \mathcal{B} still applies. To see this, it helps us to extend Φ to the space of all potential secrets (not only those in \mathbf{D}) by defining $\Phi(s_1, \dots, s_Q) := (s_1, \dots, s_Q)$ for all $(s_1, \dots, s_Q) \in R^m \setminus \mathbf{D}$. Then, consider the following two adversaries \mathcal{B}' and \mathcal{B}'' .

- \mathcal{B}' is the same as \mathcal{B} except that \mathcal{B}' samples (s_1, \dots, s_Q) from $\Phi(\mathcal{P})$.
- \mathcal{B}'' is the same as \mathcal{B} except that 1. \mathcal{B}'' samples $(r_1, \dots, r_Q) \leftarrow \mathcal{P}$ and computes secrets as

³Since N is usually set to 512 or 1024, the leading term here is $q^{k/m}$.

$(s_1, \dots, s_Q) \leftarrow \Phi(r_1, \dots, r_Q)$; 2. after \mathcal{A} returns, \mathcal{B}'' outputs $\hat{s}_i - r_i$ if there exists $r_i \neq \hat{s}_i$.

We note that \mathcal{B}' and \mathcal{B}'' do not need to be efficient. They are only used to compute the winning probability of the (efficient) adversary \mathcal{B} .

Denote now by $\text{PWin}_{\mathcal{X}}$ the winning probability of $\mathcal{X} \in \{\mathcal{A}, \mathcal{B}, \mathcal{B}', \mathcal{B}''\}$. Then, the conclusion that $\text{PWin}_{\mathcal{B}} \geq 1/2 \text{PWin}_{\mathcal{A}}$ is a straightforward corollary of the following three facts.

Fact 1. $\text{PWin}_{\mathcal{B}} = \text{PWin}_{\mathcal{B}'}$;

Fact 2. $\text{PWin}_{\mathcal{B}} = \text{PWin}_{\mathcal{B}''}$;

Fact 3. $\text{PWin}_{\mathcal{B}'} + \text{PWin}_{\mathcal{B}''} \geq \text{PWin}_{\mathcal{A}}$.

Fact 1 is straightforward since \mathcal{P} is identical to $\Phi(\mathcal{P})$, which means \mathcal{B} and \mathcal{B}' behave the same. Fact 2 is also not hard to see. Since $\text{View}_{\mathcal{A}}(r_1, \dots, r_Q) = \text{View}_{\mathcal{A}}(\Phi(r_1, \dots, r_Q)) = \text{View}_{\mathcal{A}}(s_1, \dots, s_Q)$, even if \mathcal{B}'' sets the secrets to (r_1, \dots, r_Q) instead of s_1, \dots, s_Q , the output of \mathcal{B}'' remain the same. However, then, \mathcal{B}'' is identical to \mathcal{B} , which implies the second fact.

Finally, to prove the third fact, we can interpret the sampling process of \mathcal{B}' as first sampling (r_1, \dots, r_Q) from \mathcal{P} and then set $(s_1, \dots, s_Q) \leftarrow \Phi(r_1, \dots, r_Q)$. Then, the only difference between \mathcal{B}' and \mathcal{B}'' is that \mathcal{B}' check whether there exists $s_i \neq \hat{s}_i$, while \mathcal{B}'' check whether there exists $r_i \neq \hat{s}_i$. Since Φ has no fixed point over \mathbf{D} , we know $(r_1, \dots, r_Q) \neq (s_1, \dots, s_Q)$ if $(r_1, \dots, r_Q) \in \mathbf{D}$. Therefore, for each $(r_1, \dots, r_Q) \in \mathbf{D}$, at least one of \mathcal{B}' and \mathcal{B}'' wins, which implies Fact 3.

However, if we use Φ from Equation 11.1, this extended argument does not give us any benefit. Indeed, Φ as define is a bijection over \mathbf{D} , and therefore, the relaxed condition $\mathcal{P}_{|\mathbf{D}} = \Phi(\mathcal{P}_{|\mathbf{D}})$ still requires \mathcal{P} to be a uniform distribution over \mathbf{D} . Also, it is unclear whether there are other ways to construct such Φ .

FURTHER RELAXATION ON THE REQUIREMENTS OF Φ . The way out from the above situation is that that the two distributions do not need to be exactly identical. In fact, this condition is only needed by Fact 1. Concretely, we relax the condition that $\mathcal{P}_{|\mathbf{D}} = \Phi(\mathcal{P}_{|\mathbf{D}})$ to requiring that the Rényi divergence of $\mathcal{P}_{|\mathbf{D}}$ from $\Phi(\mathcal{P}_{|\mathbf{D}})$ is small, denoted as $R_{\alpha}(\Phi(\mathcal{P}_{|\mathbf{D}}) \parallel \mathcal{P}_{|\mathbf{D}})$, where $\alpha > 1$ is a parameter we can choose. Then, due to the property of Rényi divergence (formally see Lemma 11.2.5),

$$\text{PWin}_{\mathcal{B}'} \leq (R_{\alpha}(\Phi(\mathcal{P}_{|\mathbf{D}}) \parallel \mathcal{P}_{|\mathbf{D}}) \cdot \text{PWin}_{\mathcal{B}})^{(\alpha-1)/\alpha} .$$

Combining with Fact 2 and 3, it gives

$$\text{PWin}_{\mathcal{A}} \leq \text{PWin}_{\mathcal{B}} + (R_{\alpha}(\Phi(\mathcal{P}_{|\mathbf{D}})\|\mathcal{P}_{|\mathbf{D}})) \cdot \text{PWin}_{\mathcal{B}})^{(\alpha-1)/\alpha},$$

which upper bounds the winning probability of \mathcal{A} by the advantage of solving the MSIS problem. We note that it is possible to get analogous statements using other distance measures, but they fail to give equally good parameters in our application scenarios.

Also, we emphasize that although both our work and [41] use Rényi divergence, the latter work bounds Rényi divergence between the *views* of \mathcal{A} , when run with different secrets, whereas we bound the Rényi divergence between distributions of *secrets* before and after applying the Φ map. The context where we use Rényi divergence is also very different from that of other works in lattice-based cryptography.

11.1.2 Step 2: Upper bounding $R_{\alpha}(\Phi(\mathcal{P}_{|\mathbf{D}})\|\mathcal{P}_{|\mathbf{D}})$

It is left to show that $R_{\alpha}(\Phi(\mathcal{P}_{|\mathbf{D}})\|\mathcal{P}_{|\mathbf{D}})$ is bounded, for the specific construction of Φ given in Equation 11.1, and again assuming a suitable Δ exists. The intuition that the two distributions should be close is that Φ is local in the sense that it maps each element to a nearby one given $\|(u_1\Delta, \dots, u_Q\Delta)\|$ is small. Since each secret is sampled from a discrete Gaussian distribution, elements that are close in distance have similar probabilities of being sampled. However, turning this idea into a proof is still challenging, as both $\mathcal{P}_{|\mathbf{D}}$ and $\Phi(\mathcal{P}_{|\mathbf{D}})$ are not “well-structured” distributions, since both the set \mathbf{D} and \mathbf{u} can be arbitrarily determined by \mathcal{A} .

Our key insight here is that \mathbf{D} can be decomposed into several disjoint sets such that the distributions conditioned on each set is a one-dimensional discrete Gaussian. Then, we can upper bound the Rényi divergence of the distributions conditioned on each set, which implies an upper bound on $R_{\alpha}(\Phi(\mathcal{P}_{|\mathbf{D}})\|\mathcal{P}_{|\mathbf{D}})$.

More precisely, for each $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \in \mathbf{D}$, our definition of Φ in Equation 11.1 also ensures that $\text{View}_{\mathcal{A}}(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ is identical to $\text{View}_{\mathcal{A}}(\mathbf{s}_1 + ku_1 \cdot \Delta, \dots, \mathbf{s}_Q + ku_Q \cdot \Delta)$ for any $k \in \mathbb{Z}$, where \mathbf{u} is the vector output by \mathcal{A} given secrets being $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Denote $\mathcal{S}_{\mathbf{W}}[\mathbf{s}_1, \dots, \mathbf{s}_Q] := \{(\mathbf{s}_1 + ku_1 \cdot \Delta, \dots, \mathbf{s}_Q + ku_Q \cdot \Delta) | k \in \mathbb{Z}\}$. We can show that Φ is a bijection over $\mathcal{S}_{\mathbf{W}}[\mathbf{s}_1, \dots, \mathbf{s}_Q]$, and all the $\mathcal{S}_{\mathbf{W}}[\mathbf{s}_1, \dots, \mathbf{s}_Q]$ sets are either the same or disjoint with each other. Therefore, we just need to bound

$$R_{\alpha}(\Phi(\mathcal{P}_{|\mathcal{S}_{\mathbf{W}}[\mathbf{s}_1, \dots, \mathbf{s}_Q]})\|\mathcal{P}_{|\mathcal{S}_{\mathbf{W}}[\mathbf{s}_1, \dots, \mathbf{s}_Q]})$$

for each $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \in \mathbf{D}$.

Now the problem is much easier since $\mathcal{P}_{|\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]}$ is well structured: each \mathcal{P} is a discrete Gaussian distribution over R^{Qm} , and $\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]$ is a one-dimensional lattice over R^{Qm} . In fact, we can transform $\mathcal{P}_{|\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]}$ into a discrete Gaussian over \mathbb{Z} with standard deviation

$$\|(u_1/\sigma_1 \cdot \Delta, \dots, u_Q/\sigma_Q \cdot \Delta)\|^{-1},^4$$

while $\Phi(\mathcal{P}_{|\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]})$ is exactly the same distribution shifted by 1. Therefore, we can use the upper bound from [121] (see also Lemma 11.2.7) of the Rényi divergence.

We refer to Section 12.3 for the proof details.

Further optimizations. We explain in this subsection two further optimizations we do to the reduction, which give us better parameters.

OUTPUTTING ONLY ONE SOLUTION. We observe that the reduction \mathcal{B} uses only one of the solutions output by \mathcal{A} , which means most of the solutions are actually redundant. In fact, the reduction still works if \mathcal{A} outputs a single solution $(i, \hat{\mathbf{s}})$ for the i -th challenge satisfying $u_i \neq 0$ and \mathcal{B} just outputs $\hat{\mathbf{s}} - \mathbf{s}_i$ if $\hat{\mathbf{s}} \neq \mathbf{s}_i$. This is because for each $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \in \mathbf{D}$ and $(\mathbf{s}'_1, \dots, \mathbf{s}'_Q) = \Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q)$, we have $\mathbf{s}_i \neq \mathbf{s}_i + u_i \cdot \Delta = \mathbf{s}'_i$ if $u_i \neq 0$, and thus \mathcal{B} wins in at least one of the cases if \mathcal{A} outputs a solution for the i -th challenge. This leads us to define our AOM-MISIS problem, where \mathcal{A} is only required to output a single solution $\hat{\mathbf{s}}$ together with a relaxation vector $\hat{\mathbf{b}}$ such that $\sum_{i \in [Q]} \hat{b}_i \mathbf{t}_i = A\hat{\mathbf{s}}$. Similarly, the requirement on the special solution becomes $\sum_{i \in [Q]} u_i \hat{b}_i \neq 0$. The relaxation vector $\hat{\mathbf{b}}$ is needed when reducing the security of the threshold signatures to the hardness of AOM-MISIS.

REDUCING THE NORM OF Δ . One question we do not answer in the above discussion is how to guarantee the existence of a small non-zero vector $\Delta \in R^m$ such that $A\Delta = 0$. Such Δ must exist in \mathcal{B}_β^m if $|\mathcal{B}_{\beta/2}^m| > q^{kN}$, which implies $\beta = \Omega(q^{k/m})$. However, this results in a large $\|\Delta\|$ negatively impacting both the Rényi divergence bound and σ_i , which depends linearly on $\|\Delta\|$.

We can get a smaller $\|\Delta\|$ by relying on the MLWE assumption to embed a small Δ into A as a (very minimal) trapdoor. We sample A as $[\mathbf{d}|D|\mathbb{I}_k]$, where $D \leftarrow_{\mathfrak{s}} R_q^{(m-k-1) \times k}$ and $\mathbf{d} = D\mathbf{a} + \mathbf{e}$ with $\mathbf{a} \leftarrow_{\mathfrak{s}} \mathcal{B}_{\beta_{\text{lwe}}}^{m-k-1}$ and $\mathbf{e} \leftarrow_{\mathfrak{s}} \mathcal{B}_{\beta_{\text{lwe}}}^k$. Then, we can let $\Delta = (1, -\mathbf{a}, -\mathbf{e})$, which satisfies $A\Delta = 0$, and $\|\Delta\| \leq \sqrt{m}\beta_{\text{lwe}}$, which is much smaller than $q^{k/m}$. By the MLWE assumption, $[\mathbf{d}|D]$ is computationally indistinguishable from a matrix uniformly sampled from $R_q^{(m-k) \times k}$.

⁴This is the reason why we bound $\|u_1/\sigma_1, \dots, u_Q/\sigma_Q\|$ in the winning condition.

Two remarks are in order. The first is that this trick would not have worked to improve the parameters of [41] directly, as they need something stronger than the existence of a small Δ . Second, the way we use MLWE here is different than its use to improve parameters in prior works on lattice-based signatures and threshold signatures (e.g. [98, 34, 54]). It is really tailored at supporting our tighter analysis of the reduction via the embedding of a small enough Δ .

11.2 Preliminaries

11.2.1 Polynomial Rings

Let q be an odd prime and N be a power of 2. We denote the ring $R := \mathbb{Z}[X]/(X^N + 1)$, contained in the cyclotomic field $K := \mathbb{Q}[X]/(X^N + 1)$, and let $R_q := R/qR \cong \mathbb{Z}_q[X]/(X^N + 1)$. Denote $K_{\mathbb{R}} := \mathbb{R} \otimes K \cong \mathbb{R}[X]/(X^N + 1)$. For an element $v \in K_{\mathbb{R}}$, where $v = \sum_{i=0}^{N-1} v_i X^i$, we denote its conjugate as $v^* = \sum_{i=0}^{N-1} -v_i X^{N-i}$. We use ϕ to denote the coefficient embedding that embeds $K_{\mathbb{R}}$ in \mathbb{R}^N , and ϕ maps v to vector $(v_0, \dots, v_{N-1}) \in \mathbb{R}^N$. When applying ϕ to a vector $\mathbf{v} \in K_{\mathbb{R}}^m$, ϕ maps \mathbf{v} to a vector in \mathbb{R}^{mN} by applying ϕ to each entry of \mathbf{v} . The map ϕ is a bijection, and we denote its inverse by ϕ^{-1} . An ℓ_p -norm of $\mathbf{v} \in K_{\mathbb{R}}^m$ is given by

$$\|\mathbf{v}\|_p := \|\phi(\mathbf{v})\|_p = \left(\sum_{i=1}^m \sum_{j=0}^{N-1} |v_{i,j}|^p \right)^{\frac{1}{p}},$$

where $v_{i,j}$ denotes the coefficient of X^j of the i -th entry of \mathbf{v} . Additionally, the ℓ_{∞} -norm of \mathbf{v} is defined as $\|\mathbf{v}\|_{\infty} := \max_{i \in [m], j \in [0..N-1]} |v_{i,j}|$. For the ℓ_2 -norm, we omit the subscript and denote $\|\mathbf{v}\|$ as the ℓ_2 -norm of \mathbf{v} . Denote the conjugate transpose of $\mathbf{v} \in K_{\mathbb{R}}^m$ as $\mathbf{v}^{\dagger} := (\mathbf{v}^*)^T$. We define the inner product of two vectors $\mathbf{v}, \mathbf{v}' \in K_{\mathbb{R}}^m$ as $\langle \mathbf{v}, \mathbf{v}' \rangle := \phi(\mathbf{v})^T \phi(\mathbf{v}') = \langle \phi(\mathbf{v}), \phi(\mathbf{v}') \rangle$. We have $\|\mathbf{v}\| = \langle \mathbf{v}, \mathbf{v} \rangle$. We say \mathbf{v} is a unit vector if $\|\mathbf{v}\| = 1$.

For any integer $p > 0$ and any $x \in \mathbb{Z}_p$, denote $\bar{x} \in \mathbb{Z}$ to be the lift of x such that $\bar{x} \in [0..p-1]$ and $\bar{x} = x \pmod{p}$. We use $\lceil \cdot \rceil : \mathbb{R} \rightarrow \mathbb{Z}$ to denote the rounding operator that maps any $x \in \mathbb{R}$ to $\lceil x + 1/2 \rceil$. For any integers $v > 0$ and $q > 2^v$, denote $q_v = \lfloor q/2^v \rfloor$ and denote $\lfloor \cdot \rfloor_v : \mathbb{Z}_q \rightarrow \mathbb{Z}_{q_v}$ a function that maps $x \in \mathbb{Z}_q$ to $\lfloor \bar{x}/2^v \rfloor \in \mathbb{Z}_{q_v}$. These operations can be extended an element x in R or R_q by applying them to each coefficient of x .

Also, we define a map $\phi_{\mathbb{M}}$ that maps each element in $K_{\mathbb{R}}$ to a matrix in $\mathbb{R}^{N \times N}$ as follows. Let $M_X := \begin{pmatrix} \mathbf{0} & -1 \\ \mathbb{I}_{N-1} & \mathbf{0} \end{pmatrix} \in \mathbb{R}^N$, where \mathbb{I}_{N-1} is the identity matrix in \mathbb{R}^{N-1} . For each $v \in K_{\mathbb{R}}$,

$\phi_{\mathbb{M}}(v) := \sum_{i=0}^{N-1} v_i M_X^i$, which can be viewed as the matrix representation of v . In particular, for ϕ and $\phi_{\mathbb{M}}$, the following properties hold: for any $v, v' \in K_{\mathbb{R}}$, $\phi_{\mathbb{M}}(v^*) = \phi_{\mathbb{M}}(v)^T$, $\phi_{\mathbb{M}}(vv') = \phi_{\mathbb{M}}(v)\phi_{\mathbb{M}}(v')$ and $\phi_{\mathbb{M}}(v)\phi(v') = \phi(vv')$. We extend the above definitions to R_q by representing each $v \in R_q$ as $v = \sum_{i=0}^{N-1} v_i X^i$, where $v_i \in \{-(q-1)/2, \dots, (q-1)/2\}$.

For a matrix $M \in K_{\mathbb{R}}^{m \times m}$, we denote its conjugate transpose as $M^\dagger = (M^*)^T$, and we say M is *hermitian* if $M = M^\dagger$. We say M is *positive definite* if and only if M is hermitian and for all $\mathbf{x} \in K_{\mathbb{R}}^m \setminus \{\mathbf{0}\}$, $\langle \mathbf{x}, M\mathbf{x} \rangle > 0$, or equivalently, $\phi_{\mathbb{M}}(M)$ is positive definite. Also, denote $\sigma_{\min}(M) := \inf_{\mathbf{x} \in K_{\mathbb{R}}^m, \|\mathbf{x}\|=1} \langle \mathbf{x}, M\mathbf{x} \rangle$ as the smallest singular value of M and $\sigma_{\max}(M) := \sup_{\mathbf{x} \in K_{\mathbb{R}}^m, \|\mathbf{x}\|=1} \langle \mathbf{x}, M\mathbf{x} \rangle$ as the largest singular value of M .

We borrow the following lemma from [22], which establishes the property of the set of signed monomials $\mathcal{S}_b := \{\pm 1, \dots, \pm X^{N-1}\} \subseteq R_q$.

Lemma 11.2.1 (Lemma 3.1 of [22]). *Let $\mathcal{S}_b := \{\pm 1, \dots, \pm X^{N-1}\} \subseteq R_q$. For any $b, \bar{b} \in \mathcal{S}_b$ such that $b \neq \bar{b}$, there exists $\gamma \in R$ such that $(b - \bar{b})\gamma = 2 \pmod q$ and γ is a polynomial with coefficients only in $\{-1, 0, 1\}$.*

11.2.2 Lattices and Discrete Gaussian Distributions

In this subsection, we give definitions for lattices and discrete Gaussian distributions over \mathbb{R} and $K_{\mathbb{R}}$. An m -dimensional lattice Λ over \mathbb{Z} (resp. R) is a discrete additive subgroup of \mathbb{Z} (resp. R). Equivalently, $\Lambda = \mathcal{L}(\{\mathbf{b}_1, \dots, \mathbf{b}_k\}) := \{\sum_{i \in [k]} x_i \mathbf{b}_i : x_i \in \mathbb{Z}\}$ for a set of linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{Z}^m$ (resp. R^m), which is referred to as a basis of Λ . The size k is the *rank* of the lattice Λ . We say Λ is a *full rank* lattice if $k = m$ (resp. $k = mN$ for Λ over R). For any $a \in \mathbb{Z}^m$ (resp. R^m), $\Lambda + a$ is a *coset* of Λ . The *dual lattice* of Λ is denoted as $\Lambda^* = \{\mathbf{x} \in \text{Span}(\Lambda) : \forall \mathbf{y} \in \Lambda, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$. A Λ -subspace is the linear span of some subset of Λ , i.e., a subspace S such that $S = \text{Span}(S \cap \Lambda)$.

For a matrix $A \in R_q^{k \times m}$, we define the R -lattice $\Lambda_q^\perp(A) \subseteq R^m$ as

$$\Lambda_q^\perp(A) := \{\mathbf{x} \in R^m : A\mathbf{x} = 0 \pmod q\}.$$

We know $\Lambda_q^\perp(A)$ has full-rank since $qR^m \subseteq \Lambda_q^\perp(A)$.

For a positive definite matrix $\Sigma \in \mathbb{R}^{m \times m}$ (resp. an invertible positive definite matrix $\Sigma \in K_{\mathbb{R}}^{m \times m}$) and a vector $\mathbf{c} \in \mathbb{R}^m$ (resp. $K_{\mathbb{R}}^m$), we define the function $\rho_{\Sigma, \mathbf{c}}$ over \mathbb{R}^m (resp. $K_{\mathbb{R}}^m$) as

$$\rho_{\Sigma, \mathbf{c}}(\mathbf{x}) := \exp\left(-\pi \langle \mathbf{x} - \mathbf{c}, \Sigma^{-1}(\mathbf{x} - \mathbf{c}) \rangle\right).$$

Then, we denote $\mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^m$ as the discrete Gaussian distribution over a lattice coset $\Lambda + \mathbf{a} \subseteq \mathbb{Z}^m$ (resp. R^m) with covariance matrix Σ , centered at $\mathbf{c} \in \mathbb{R}^m$, where for $\mathbf{x} \in \Lambda + \mathbf{a}$, we define

$$\mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^m(\mathbf{x}) := \Pr[\mathbf{x} \leftarrow_s \mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^m] = \frac{\rho_{\Sigma,\mathbf{c}}(\mathbf{x})}{\rho_{\Sigma,\mathbf{c}}(\Lambda + \mathbf{a})}$$

where $\rho_{\Sigma,\mathbf{c}}(\Lambda + \mathbf{a}) = \sum_{\mathbf{x} \in \Lambda + \mathbf{a}} \rho_{\Sigma,\mathbf{c}}(\mathbf{x})$. For $\Lambda + \mathbf{a} \subseteq R^m$, we denote $\mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^{m,\text{mod } q}(\mathbf{x})$ as the distribution of $(\mathbf{x} \bmod q) \in R_q^m$ for \mathbf{x} sampled from $\mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^m$.

Also, we make some remarks about the notations we will use throughout the paper. When $\Sigma = \sigma^2 \mathbb{I}_m$ for $\sigma \in \mathbb{R}$, we will use $\rho_{\sigma,\mathbf{c}}$ and $\mathcal{D}_{\Lambda+\mathbf{a},\sigma,\mathbf{c}}^m$ as $\rho_{\Sigma,\mathbf{c}}$ and $\mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^m$, respectively. If the center $\mathbf{c} = 0$, then we omit the subscript \mathbf{c} from $\rho_{\Sigma,\mathbf{c}}$ and $\mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^m$. Moreover, when $\Lambda + \mathbf{a} = \mathbb{Z}^m$ (resp. $\Lambda + \mathbf{a} = R^m$), we omit $\Lambda + \mathbf{a}$ from the subscript of $\mathcal{D}_{\Lambda+\mathbf{a},\Sigma,\mathbf{c}}^m$.

The smoothing parameter of a lattice Λ with respect to $\varepsilon > 0$, denoted by $\eta_\varepsilon(\Lambda)$, is the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. Throughout the paper, we set $\varepsilon = 2^{-2\kappa}$.

We borrow the following lemma from [6] that bounds the ℓ_2 -norm of discrete Gaussian random variables and adapt it to lattices over $K_{\mathbb{R}}$.

Lemma 11.2.2 (Lemma 3 of [6] adapted to $K_{\mathbb{R}}$). *For any $\varepsilon \in (0, 1)$, a lattice $\Lambda \subseteq R^m$, $\mathbf{c} \in K_{\mathbb{R}}^m$, and $\sigma \geq \eta_\varepsilon(\Lambda)$, then*

$$\Pr[\|\mathbf{x} - \mathbf{c}\| \geq \sigma\sqrt{mN} : \mathbf{x} \leftarrow_s \mathcal{D}_{\Lambda,\sigma,\mathbf{c}}] \leq \frac{1 + \varepsilon}{1 - \varepsilon} \cdot 2^{-mN}.$$

We also borrow the following lemma from [76] that show upper bounds of smoothing parameters for general lattices.

Lemma 11.2.3 (Lemma 2.6 of [76]). *For any full-rank lattice Λ in \mathbb{R}^m and $\varepsilon > 0$, $\eta_\varepsilon(\Lambda) \leq \frac{\sqrt{\log(2m(1+1/\varepsilon))/\pi}}{\lambda_1^\infty(\Lambda^*)}$, where $\lambda_1^\infty(\Lambda^*)$ denotes the ℓ_2 norm of the shortest non-zero vector in the ℓ_∞ norm in the dual lattice Λ^* .*

11.2.3 Assumptions

We recall the module short integer solution (MSIS) problem and the module learning with error (MLWE) problem (defined in Figure 11.1). The advantage of \mathcal{A} for the MSIS problem is defined as

$$\text{Adv}_{q,N,k,m,\beta}^{\text{msis}}(\mathcal{A}) := \Pr[\text{MSIS}_{q,N,k,m,\beta}^{\mathcal{A}} = 1].$$

The advantage of \mathcal{A} for the MLWE problem is defined as

$$\text{Adv}_{q,N,k,m,\beta}^{\text{mlwe}}(\mathcal{A}) := \Pr[\text{MLWE}_{q,N,k,m,\beta}^{\mathcal{A}} = 1].$$

<p>Game MSIS$_{q,N,k,m,\beta}^A$:</p> <p>$\bar{A} \leftarrow_{\\$} R_q^{k \times (m-k)}$</p> <p>$A \leftarrow [\bar{A}]_{\mathbb{I}_k}$</p> <p>$\mathbf{x} \leftarrow \mathcal{A}(A)$</p> <p>Return ($\ \mathbf{x}\ \leq \beta \wedge A\mathbf{x} = 0$)</p>	<p>Game MLWE$_{q,N,k,m,\beta}^A$:</p> <p>$\bar{A} \leftarrow_{\\$} R_q^{k \times (m-k)} ; A \leftarrow [\bar{A}]_{\mathbb{I}_k}$</p> <p>$\mathbf{s} \leftarrow_{\\$} \mathcal{B}_{\beta}^m ; \mathbf{t}_0 \leftarrow A\mathbf{s} ; \mathbf{t}_1 \leftarrow_{\\$} R_q^k$</p> <p>$b \leftarrow_{\\$} \{0, 1\}$</p> <p>$\hat{b} \leftarrow \mathcal{A}(A, \mathbf{t}_b)$</p> <p>Return ($\hat{b} = b$)</p>
---	---

Figure 11.1: The module-SIS and module-LWE problems, where $R := \mathbb{Z}[X]/(X^N + 1)$, $R_q := R/qR$ and $\mathcal{B}_{\beta}^m := \{\mathbf{x} \in R^m \mid \|\mathbf{x}\|_{\infty} \leq \beta\}$.

11.2.4 Rényi Divergence

We define the notion of Rényi Divergence between two distributions P, Q which we will use in our analysis of the scheme.

Definition 11.2.4 (Rényi Divergence). *Let P, Q be two discrete probability distributions such that $\text{Supp}(P) \subseteq \text{Supp}(Q)$. We define the Rényi Divergence of order α , for $\alpha \in (1, \infty)$ as $R_{\alpha}(P\|Q) := \left(\sum_{x \in \text{Supp}(P)} \frac{P(x)^{\alpha}}{Q(x)^{\alpha-1}} \right)^{\frac{1}{\alpha-1}}$.*

The following lemma gives basic properties of the Rényi Divergence.

Lemma 11.2.5 (Lemma 4.1 of [93]). *Let $\alpha \in (1, \infty)$ and P, Q be discrete probability distributions with $\text{Supp}(P) \subseteq \text{Supp}(Q)$. Then, the following properties hold:*

- **Data Processing Inequality:** $R_{\alpha}(P^f\|Q^f) \leq R_{\alpha}(P\|Q)$ for any function f , where P^f (and Q^f) denotes the distribution which samples $x \leftarrow_{\$} P$ ($x \leftarrow_{\$} Q$) and outputs $f(x)$.
- **Probability Preservation:** Let $E \subseteq \text{Supp}(Q)$ be an arbitrary event. Then, for $\alpha \in (1, \infty)$,

$$\Pr_{x \leftarrow_{\$} P}[E] \leq (\Pr_{x \leftarrow_{\$} Q}[E] R_{\alpha}(P\|Q))^{(\alpha-1)/\alpha}.$$

Also, we show the following property.

Lemma 11.2.6. *Let P and Q denote two distributions over the union of a countable number of disjoint sets $\{S_i\}_{i \in U}$ such that $P(S_i) = Q(S_i)$ for any $i \in U$. Then, for any $\alpha > 1$, if there exists δ such that $R_\alpha(P|_{S_i} \| Q|_{S_i}) \leq \delta$ for any $i \in U$, then $R_\alpha(P \| Q) \leq \delta$.*

Proof.

$$\begin{aligned} R_\alpha(P \| Q)^{\alpha-1} &= \sum_{x \in \text{Supp}(P)} \frac{P(x)^\alpha}{Q(x)^{\alpha-1}} = \sum_{i \in U} \sum_{x \in S_i} \frac{P(x)^\alpha}{Q(x)^{\alpha-1}} \\ &= \sum_{i \in U} Q(S_i) \sum_{x \in S_i} \frac{(P(x)/P(S_i))^\alpha}{(Q(x)/Q(S_i))^{\alpha-1}} \\ &= \sum_{i \in U} Q(S_i) R_\alpha(P|_{S_i} \| Q|_{S_i})^{\alpha-1} \leq \sum_{i \in U} Q(S_i) \delta^{\alpha-1} = \delta^{\alpha-1}, \end{aligned}$$

which concludes the lemma. □ □

We borrow the following lemma from [121], which upperbounds the Rényi Divergence between two discrete Gaussian distributions with different centers.

Lemma 11.2.7 (Lemma 5 of [121]). *For any m -dimensional lattice $\Lambda \subseteq \mathbb{R}^m$, $\sigma > 0$, and two vectors $\mathbf{c}, \mathbf{c}' \in \mathbb{R}^m$, let $P = \mathcal{D}_{\Lambda, \sigma, \mathbf{c}}^m$ and $Q = \mathcal{D}_{\Lambda, \sigma, \mathbf{c}'}^m$. If $\mathbf{c}, \mathbf{c}' \in \Lambda$, set $\varepsilon = 0$. Otherwise, fix $\varepsilon \in (0, 1)$ and assume $\sigma \geq \eta_\varepsilon(\Lambda)$. Then,*

$$R_\alpha(P \| Q) \leq \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^{\frac{\alpha}{\alpha-1}} \exp \left(\alpha \pi \frac{\|\mathbf{c} - \mathbf{c}'\|^2}{\sigma^2} \right).$$

Chapter 12

THE ALGEBRAIC ONE-MORE MISIS PROBLEM

In this chapter, we first formally define the algebraic one-more MISIS (AOM-MISIS) problem, then provide a comparison with the AOM-MLWE problem, and finally present a formal reduction from standard lattice problems to AOM-MISIS.

12.1 Definition of AOM-MISIS

The algebraic one-more MISIS game is defined in Figure 12.1. The game is defined implicitly over the cyclotomic ring $R = \mathbb{Z}[X]/(X^N + 1)$, where N is a power of two, as well as the associated ring $R_q = R/qR \cong \mathbb{Z}_q[X]/(X^N + 1)$ for an odd prime q . The adversary \mathcal{A} is given Q MISIS challenges $\mathbf{t}_1, \dots, \mathbf{t}_Q$ and the goal is to output a short solution $\hat{\mathbf{s}}$ for a linear combination of the challenges $\sum_{i \in [Q]} \hat{b}_i \mathbf{t}_i$, where the norm of $\hat{\mathbf{b}}$ is suitably bounded. The adversary can also issue queries to the oracle PI which reveal linear combinations of the secrets. To win, the adversary also needs to output an additional vector \mathbf{u} which is orthogonal to all queries to PI, but not to the vector $\hat{\mathbf{b}}$. We can also view the existence of such \mathbf{u} as a constraint on the PI queries. We refer the reader to the introduction for some additional intuition. Here, for $\text{par} = (q, N, k, m, Q, (\sigma_i)_{i \in [Q]}, \beta_s, \beta_b, \beta_u)$, we define the *advantage* of an adversary \mathcal{A} as

$$\text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{A}, \kappa) = \Pr \left[\text{AOM-MISIS}_{\text{par}}^{\mathcal{A}}(\kappa) = 1 \right]. \quad (12.1)$$

We note that we slightly abuse notation in the asymptotic definition of the game, since it is understood that all parameters grow with κ , including Q , and thus the notation $(\sigma_i)_{i \in [Q]}$ is not entirely well-defined. This will not be an issue in actual use cases, and we dispense with a more rigorous definition.

Remark 12.1. We note that for each $i \in [Q]$, it is equivalent to sample \mathbf{s}_i from $\mathcal{D}_{\sigma_i}^{m, \text{mod } q}$ instead of $\mathcal{D}_{\sigma_i}^m$, since the view of \mathcal{A} remains unchanged. With this insight, it is natural to allow σ_i to be ∞ , in which case \mathbf{s}_i is sampled uniformly from R_q^m , $1/\sigma_i = 0$, and we require $\hat{b}_i = 0$ in the winning condition (since, o.w., $\|\hat{b}_i \cdot \sigma_i\|_1$ is ∞). This extension is used, in particular, in our security

<p>Game AOM-MISIS$_{q,N,k,m,Q,(\sigma_i)_{i \in [Q]},\beta_s,\beta_b,\beta_u}^A(\kappa)$:</p> <hr style="border: 0.5px solid black;"/> <p>$B \leftarrow \emptyset$ $\bar{A} \leftarrow_{\\$} R_q^{k \times (m-k)} ; A \leftarrow [\bar{A} \mathbb{I}_k]$ For $i \in [Q]$, $\mathbf{s}_i \leftarrow \mathcal{D}_{\sigma_i}^m ; \mathbf{t}_i \leftarrow A \mathbf{s}_i \bmod q$ $(\hat{\mathbf{s}}, \hat{\mathbf{b}}, \mathbf{u}) \leftarrow \mathcal{A}^{\text{PI}}(A, \{\mathbf{t}_i\}_{i \in [Q]}) \ // \ \hat{\mathbf{s}} \in R^m, \hat{\mathbf{b}}, \mathbf{u} \in R^Q$ Return $(\forall \mathbf{b} \in B : \mathbf{b}^T \mathbf{u} = 0 \bmod q) \wedge \hat{\mathbf{b}}^T \mathbf{u} \neq 0$ $\wedge \ (u_1/\sigma_1, \dots, u_Q/\sigma_Q)\ _2 \leq \beta_u$ $\wedge \ \hat{\mathbf{s}}\ _2 \leq \beta_s \wedge \ (\hat{b}_1 \cdot \sigma_1, \dots, \hat{b}_Q \cdot \sigma_Q)\ _1 \leq \beta_b$ $\wedge \sum_{i \in [Q]} \hat{b}_i \mathbf{t}_i = A \hat{\mathbf{s}} \bmod q$</p> <p>Oracle PI$(\mathbf{b} \in R^Q)$:</p> <hr style="border: 0.5px solid black;"/> <p>$B \leftarrow B \cup \{\mathbf{b}\}$ Return $\sum_{i \in [Q]} b_i \mathbf{s}_i \bmod q \ (\in R_q^m)$</p>

Figure 12.1: The AOM-MISIS game, where $R := \mathbb{Z}[X]/(X^N + 1)$ and $R_q := R/qR$.

reduction of the CATZ construction (see Section 13.1.1), where the random values for generating the key shares are sampled uniformly from R_q^m .

12.2 Comparison with prior work

We present in Figure 12.2 the AOM-MLWE problem proposed by Espitau et al. [60]. In our formulation, the PI oracle corresponds to the $\mathcal{O}_{\text{solve}}$ oracle in their notation, and the constraints on PI queries are explicitly stated within the game. More precisely, the constraints are:

- The number of PI queries is exactly $Q - 1$.
- Denote $\begin{pmatrix} \mathbf{v}^T \\ \underline{D} \end{pmatrix} = [\mathbf{d}_1 | \dots | \mathbf{d}_{Q-1}]$, where \mathbf{d}_i denotes the i -th PI query. \underline{D} is an invertible matrix.
- Let $\mathbf{w} \leftarrow (\mathbf{v}^T \underline{D}^{-1})^T$. The ℓ_2 -norm of each entry of \mathbf{w} is bounded by β_d .




Figure 12.2: The AOM-MLWE game.

We also define, for $\text{par} = (q, N, k, m, Q, (\sigma_i)_{i \in [Q]}, \beta_s, \beta_e, \beta_b, \beta_d)$, the corresponding advantage

$$\text{Adv}_{\text{par}}^{\text{aom-mlwe}}(\mathcal{A}, \kappa) = \Pr [\text{AOM-MLWE}_{\text{par}}^{\mathcal{A}}(\kappa) = 1] .$$

Compared with the AOM-MLWE problem, our problem differs in the following key aspects:

1. Our problem only asks the adversary to output one special solution $\hat{\mathbf{s}}$ for a linear combination of the challenges, while the AOM-MLWE problem demands one solution for each challenge. This makes our problem inherently easier to solve, as noted at the beginning of Section 11.1. Moreover, this relaxation simplifies the security reduction, as it only needs to extract one solution rather than multiple, making it easier to reduce from our assumption. Additionally, it enables better parameter selections, since the norm bound applies to a single solution instead of n solutions.

2. The constraints on the PI queries differ between AOM-MLWE and our problem. In particular, their constraints are a special case of ours: given the PI queries satisfying their constraints, we can verify that these queries also satisfy our constraints by setting $\mathbf{u} = (1, -\mathbf{w})$, since $\mathbf{u} \begin{pmatrix} \mathbf{v}^T \\ \underline{D} \end{pmatrix} = 0$. We will show this formally later in this section.
3. The norm bounding approach is different. In AOM-MLWE, the norm of each entry of $\hat{\mathbf{b}}$ and $\hat{\mathbf{w}}$ is bounded individually. In contrast, we bound the norm of the entire vectors $\hat{\mathbf{b}}$ and $\hat{\mathbf{u}}$, with each entry weighted by the standard deviation of the corresponding challenge. This adjustment leads to better parameter selections.

In short, our problem is easier to reduce from and enables better parameter selections. We formally show in Lemma 12.2.1 that the hardness of our problem implies the hardness of theirs, and we discuss the improvement in the parameter selections in Section 13.2.1, Remark 13.1.

Remark 12.2. *We note that Espitau et al. [60] propose an alternative way of defining constraints, which, like ours, requires the existence of a nonzero vector \mathbf{u} that is orthogonal to all PI queries. The key difference is that they impose a bound on the norm of each entry of \mathbf{u} . We can show that the hardness of our problem also implies the hardness of theirs under these constraints. However, we omit a formal analysis here, as the proof idea is very similar, and this alternative version is not used to establish the security of their threshold signature scheme.*

Lemma 12.2.1. *For any $\text{par} = (q, N, k, m, Q, (\sigma_i)_{i \in [Q]}, \beta_s, \beta_e, \beta_b, \beta_d)$ and any adversary \mathcal{A} playing the game $\text{AOM-MLWE}_{\text{par}}^{\mathcal{A}}$, there exists an AOM-MISIS adversary \mathcal{B} running in time roughly the same as \mathcal{A} such that*

$$\text{Adv}_{\text{par}}^{\text{aom-mlwe}}(\mathcal{A}, \kappa) \leq \text{Adv}_{\text{par}'}^{\text{aom-misis}}(\mathcal{B}, \kappa),$$

where $\text{par}' = (q, N, k, m, Q, (\sigma_i)_{i \in [Q]}, \beta'_s = \beta_s + \beta_e, \beta'_b = \sqrt{N}\beta_b\sigma_1, \beta_u)$ and $\beta_u = 1/\sigma_1 + \beta_d\sqrt{Q}/(\min_{i \in [2..Q]} \sigma_i)$.

Roughly, the proof idea is that for an adversary \mathcal{A} that wins the game AOM-MLWE, \mathcal{A} can win the game AOM-MISIS by outputting only the solution to the first challenge (i.e., outputting \hat{s}_1 and $(\hat{b}_1, 0, \dots, 0)$) and outputting $\mathbf{u} \leftarrow (1, -\mathbf{v}^T \underline{D}^{-1})$. It is not hard to check \mathbf{u} satisfies the constraints of the game AOM-MISIS. We do this explicitly in the following proof.

Proof of Lemma 12.2.1. For any adversary \mathcal{A} described in the lemma, we construct \mathcal{B} as follows. To start with, \mathcal{B} runs \mathcal{A} on its input $A, \{\mathbf{t}_i\}_{i \in [Q]}$ by forwarding all PI queries from \mathcal{A} to its own PI oracle. After receiving \mathcal{A} 's output $(\hat{\mathbf{b}}, \hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_Q)$, if \mathcal{A} wins the AOM-MLWE game simulated by \mathcal{B} , \mathcal{B} returns $((\hat{b}_1, 0, \dots, 0), \hat{\mathbf{s}}_1, (1, -\mathbf{w}^T))$, where \mathbf{w} is defined in the AOM-MLWE game. Otherwise \mathcal{B} aborts.

We now show \mathcal{B} wins the AOM-MISIS game given \mathcal{A} wins by checking all the winning conditions of \mathcal{B} .

- For the first condition, since $(1, -\mathbf{w}^T) \cdot \begin{pmatrix} \mathbf{v}^T \\ \underline{D} \end{pmatrix} = \mathbf{v}^T - \mathbf{v}^T \underline{D}^{-1} \underline{D} = 0$, for any $i \in [Q - 1]$, $(1, -\mathbf{w}^T) \cdot \mathbf{d}_i = 0$, so the first condition is satisfied.
- For the second condition, $(\hat{b}_1, 0, \dots, 0) \begin{pmatrix} 1 \\ -\mathbf{w} \end{pmatrix} = \hat{b}_1 \neq 0$.
- For the third condition, $\|(1/\sigma_1, w_1/\sigma_2, \dots, w_{Q-1}/\sigma_Q)\| \leq \frac{1}{\sigma_1} + \sum_{i \in [Q-1]} \frac{\|w_i\|}{\sigma_i} \leq \frac{1}{\sigma_1} + \frac{\beta_d}{\min_{i \in [2..Q]} \sigma_i} = \beta_u$.
- For the fourth condition, $\|\hat{\mathbf{s}}_1\| \leq \|\hat{\mathbf{s}}_{1,[m-k]}\| + \|\hat{\mathbf{s}}_{1,[(m-k+1)..k]}\| \leq \beta_s + \beta_e = \beta'_s$.
- For the fifth condition, $\|(\hat{b}_1 \cdot \sigma_1, 0, \dots, 0)\|_1 \leq \|\hat{b}_1 \cdot \sigma_1\|_1 \leq \sqrt{N} \sigma_1 \|\hat{b}_1\| = \beta'_b$.
- For the final condition, $\hat{b}_1 \mathbf{t}_1 = A \hat{\mathbf{s}} \pmod{q}$.

□

12.3 Reduction from MSIS and MLWE

This section shows our main result establishing hardness of AOM-MISIS from the hardness of MLWE and MSIS. (We remind the reader that we also provided a detailed overview of this proof in Section 11.1 above.)

Theorem 12.3.1. *For any $\varepsilon \in (0, 1)$, $\alpha > 1$, any $\text{par} = (q, N, k, m, Q, (\sigma_i)_{i \in [Q]}, \beta_s, \beta_b, \beta_u)$ such that $mN \geq 2\kappa$ and $\sigma_i \geq \sqrt{\log(6mN)/\pi}$, and any AOM-MISIS adversary \mathcal{A} , there exist a MSIS*

adversary \mathcal{B} and two MLWE adversaries \mathcal{C} and \mathcal{D} , such that

$$\begin{aligned} \text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{A}, \kappa) &\leq 2\delta_\alpha \left(\text{Adv}_{q,k,m,\beta_{\text{sis}}}^{\text{msis}}(\mathcal{B}, \kappa) + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{D}, \kappa) + Q \cdot 2^{-2\kappa+2} \right)^{\frac{\alpha}{\alpha-1}} \\ &\quad + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{C}, \kappa), \end{aligned}$$

where $\delta_\alpha = \frac{1+\varepsilon}{1-\varepsilon} \cdot \exp((\alpha-1)\pi^2/\log(2(1+1/\varepsilon)))$, $\beta_{\text{sis}} = \beta_s + \sqrt{mN}\beta_b$, and $\beta_{\text{lwe}} = 1/(\beta_u\sqrt{mN}\sqrt{\log(2(1+1/\varepsilon))/\pi})$. The three adversaries have, roughly the same running time as that of \mathcal{A} .

Proof. We prove the theorem via the following hybrid.

\mathbf{G}_0 : Same as the AOM-MISIS game.

\mathbf{G}_1 : Same as \mathbf{G}_0 except that A is sampled with a short solution embedded, i.e., $A \leftarrow [D\mathbf{a} + \mathbf{e}|D|\mathbb{I}_k]$ for $(\mathbf{a}, \mathbf{e}) \leftarrow_{\$} \mathcal{B}_{\beta_{\text{lwe}}}^{m-1}$, where $\mathcal{B}_{\beta_{\text{lwe}}}^{m-1} := \{\mathbf{x} \in R^{m-1} \mid \|\mathbf{x}\|_\infty \leq \beta_{\text{lwe}}\}$, and $D \leftarrow_{\$} R_q^{k \times (m-k-1)}$. Since the only difference between \mathbf{G}_0 and \mathbf{G}_1 is that \bar{A} is replaced with a MLWE challenge, there exists a MLWE adversary \mathcal{C} such that

$$\text{Adv}^{\mathbf{G}_0}(\mathcal{A}, \kappa) \leq \text{Adv}^{\mathbf{G}_1}(\mathcal{A}, \kappa) + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{C}, \kappa). \quad (12.2)$$

Consider a variant of the MSIS game, where A is sampled with a short solution embedded (same as the above \mathbf{G}_1), i.e., $A \leftarrow [D\mathbf{a} + \mathbf{e}|D|\mathbb{I}_k]$ for $(\mathbf{a}, \mathbf{e}) \leftarrow_{\$} \mathcal{B}_{\beta_{\text{lwe}}}^{m-1}$, where $\mathcal{B}_{\beta_{\text{lwe}}}^{m-1} := \{\mathbf{x} \in R^{m-1} \mid \|\mathbf{x}\|_\infty \leq \beta_{\text{lwe}}\}$, and $D \leftarrow_{\$} R_q^{k \times (m-k-1)}$. We refer to the game as td-MSIS. We construct \mathcal{B} playing the td-MSIS game as follows. Given the td-MSIS challenge $A' \in R_q^{k \times (m-k)}$, \mathcal{B} runs \mathcal{A} by simulating the game \mathbf{G}_1 with \mathcal{A} faithfully except that \mathcal{B} sets $A \leftarrow A'$ instead of sampling A by itself. After receiving the output $(\hat{\mathbf{s}}, \hat{\mathbf{b}}, \mathbf{u})$ from \mathcal{A} , if \mathcal{A} wins the game \mathbf{G}_1 simulated by \mathcal{B} , \mathcal{B} outputs $\mathbf{x} \leftarrow \hat{\mathbf{s}} - \sum_{i \in [Q]} \hat{b}_i \mathbf{s}_i$. Otherwise, \mathcal{B} aborts.

ANALYSIS OF \mathcal{B} . Similar to the hybrid between \mathbf{G}_0 and \mathbf{G}_1 , the td-MSIS game is computationally indistinguishable from the MSIS game under the MLWE assumption. Therefore, there exists a MLWE adversary \mathcal{D} such that

$$\text{Adv}_{q,k,m-1,\beta_{\text{sis}}}^{\text{td-msis}}(\mathcal{B}, \kappa) \leq \text{Adv}_{q,k,m-1,\beta_{\text{sis}}}^{\text{msis}}(\mathcal{B}, \kappa) + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{D}, \kappa). \quad (12.3)$$

Claim 12.3.2. $\text{Adv}^{\mathbf{G}_1}(\mathcal{A}, \kappa) \leq 2\delta_\alpha \left(\text{Adv}_{q,k,m-1,\beta_{\text{sis}}}^{\text{td-msis}}(\mathcal{B}, \kappa) + Q \cdot 2^{-2\kappa+2} \right)^{(\alpha-1)/\alpha}$.

We can conclude the proof since

$$\begin{aligned}
\text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{A}, \kappa) &\leq \text{Adv}^{\mathbf{G}_1}(\mathcal{A}, \kappa) + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{C}, \kappa) \\
&\leq 2\delta_\alpha \left(\text{Adv}_{q,k,m-1,\beta_{\text{sis}}}^{\text{td-msis}}(\mathcal{B}, \kappa) + Q \cdot 2^{-2\kappa+2} \right)^{(\alpha-1)/\alpha} + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{C}, \kappa) \\
&\leq 2\delta_\alpha \left(\text{Adv}_{q,k,m,\beta_{\text{sis}}}^{\text{msis}}(\mathcal{B}, \kappa) + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{D}, \kappa) + Q \cdot 2^{-2\kappa+2} \right)^{\frac{\alpha}{\alpha-1}} \\
&\quad + \text{Adv}_{q,k,m-1,\beta_{\text{lwe}}}^{\text{mlwe}}(\mathcal{C}, \kappa),
\end{aligned}$$

where the first inequality follows from Equation (12.2), the second inequality follows from the claim, and the third inequality follows from Equation (12.3). \square

Proof of Claim 12.3.2. We first show that, except for negligible probability, \mathcal{B} wins if its output vector \mathbf{x} is non-zero. It is clear that if \mathcal{A} wins the game \mathbf{G}_1 simulated by \mathcal{B} , by the winning conditions of the game, it holds that $A\mathbf{x} = A\hat{\mathbf{s}} - \sum_{i \in [Q]} \hat{b}_i \mathbf{t}_i = 0$.

Also, $\|\mathbf{x}\|$ is small except for negligible probability. Denote BNonce as the event that there exists $i \in [Q]$ such that $\|\mathbf{s}_i\| > \sigma_i \sqrt{mN}$ for $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ sampled as in the game \mathbf{G}_1 . If BNonce does not occur, we know

$$\begin{aligned}
\|\mathbf{x}\| &\leq \|\hat{\mathbf{s}}\| + \sum_{i \in [Q]} \|\hat{b}_i\|_1 \cdot \|\mathbf{s}_i\| \leq \beta_s + \sum_{i \in [Q]} \|\hat{b}_i\|_1 \cdot (\sqrt{mN} \sigma_i) \\
&= \beta_s + \sqrt{mN} \left\| (\hat{b}_1 \cdot \sigma_1, \dots, \hat{b}_Q \cdot \sigma_Q) \right\|_1 = \beta_s + \sqrt{mN} \beta_b \leq \beta_{\text{sis}}.
\end{aligned}$$

Therefore, \mathcal{B} wins if BNonce does not occur, which means $\text{Adv}_{q,k,m-1,\beta_{\text{sis}}}^{\text{td-msis}}(\mathcal{B}, \kappa) \geq \Pr[\mathcal{B} \text{ outputs } \mathbf{x} \neq 0] - \Pr[\text{BNonce}]$. To bound $\Pr[\text{BNonce}]$, since \mathbf{s}_i is sampled from $\mathcal{D}_{\sigma_i}^m$ and $\sigma_i \geq \sqrt{\log(6mN)}/\pi = \eta_{\varepsilon'}(R^m)$ with $\varepsilon' = 1/2$, by Lemma 11.2.2, $\Pr[\|\mathbf{s}_i\| > \sigma_i \sqrt{mN}] \leq 3 \cdot 2^{-mN} \leq 2^{-2\kappa+2}$. By the Union Bound, $\Pr[\text{BNonce}] \leq \sum_{i \in [Q]} \Pr[\|\mathbf{s}_i\| > \sigma_i \sqrt{mN}] \leq Q \cdot 2^{-2\kappa+2}$. Therefore,

$$\text{Adv}_{q,k,m-1,\beta_{\text{sis}}}^{\text{td-msis}}(\mathcal{B}, \kappa) \geq \Pr[\mathcal{B} \text{ outputs } \mathbf{x} \neq 0] - Q \cdot 2^{-2\kappa+2}. \quad (12.4)$$

To bound the probability that the output of \mathcal{B} is non-zero, consider a fixed randomness $(\mathbf{a}, \mathbf{e}, D)$ of td-MSIS. Also, w.l.o.g. assume \mathcal{A} is deterministic. Then, the execution of \mathcal{B} is determined by $\mathbf{s}_1, \dots, \mathbf{s}_Q$. We define a map $\Phi_{\mathbf{a}, \mathbf{e}, D, \mathcal{A}} : R^{Qm} \rightarrow R^{Qm}$ as follows such that the view of \mathcal{A} given $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ is exactly the same as that given $\Phi_{\mathbf{a}, \mathbf{e}, D, \mathcal{A}}(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Consider the execution given $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. To simplify notation, we omit the subscript of Φ for the rest of the proof. If \mathcal{A} does not win the \mathbf{G}_1 simulated by \mathcal{B} , we set $\Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q) = (\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Otherwise, we set

$\Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q) = (\mathbf{s}_1 + u_1\Delta, \dots, \mathbf{s}_Q + u_Q\Delta)$, where $\mathbf{u} \in R^Q$ is output by \mathcal{A} and $\Delta = (1, -\mathbf{a}, -\mathbf{e}) \in R^m$.

It is not hard to see that the view of \mathcal{A} given $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ is identical to that given $\Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. In particular, in the case that \mathcal{A} wins the game \mathbf{G}_1 , $\Phi(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ leads to the same view since $\sum_{i \in [Q]} b_i(\mathbf{s}_i + u_i\Delta) = \sum_{i \in [Q]} b_i\mathbf{s}_i + \Delta \sum_{i \in [Q]} b_i u_i = \sum_{i \in [Q]} b_i\mathbf{s}_i$ for any $(b_1, \dots, b_Q) \in B$ and $A(\mathbf{s}_i + u_i\Delta) = A\mathbf{s}_i$ for any $i \in [Q]$ due to the fact that $A \cdot \Delta = [D\mathbf{a} + \mathbf{e}|D|\mathbb{I}_k] \cdot (1, -\mathbf{a}, -\mathbf{e})^T = \mathbf{0}$.

Also, it is not hard to see that Φ is a bijection. For $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \in R^{mQ}$, suppose $\Phi(\mathbf{s}'_1, \dots, \mathbf{s}'_Q) = (\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Since $(\mathbf{s}'_1, \dots, \mathbf{s}'_Q)$ leads to the same view of \mathcal{A} as $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$, we have either $(\mathbf{s}'_1, \dots, \mathbf{s}'_Q) = (\mathbf{s}_1, \dots, \mathbf{s}_Q)$ in case that \mathcal{A} does not win the game \mathbf{G}_1 , or $(\mathbf{s}'_1, \dots, \mathbf{s}'_Q) = (\mathbf{s}_1 - u_1\Delta, \dots, \mathbf{s}_Q - u_Q\Delta)$ in case that \mathcal{A} wins the game \mathbf{G}_1 by outputting \mathbf{u} given $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Therefore, such $(\mathbf{s}'_1, \dots, \mathbf{s}'_Q)$ must exist uniquely, which means that Φ is a bijection.

In the game \mathbf{G}_1 , the distribution of $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ is \mathcal{D}_Σ^{mQ} , where $\Sigma = \mathbb{I}_m \otimes \text{diag}(\sigma_1^2, \dots, \sigma_Q^2)$. Denote $P = \Phi(\mathcal{D}_\Sigma^{mQ})$. Following the idea described in Section 11.1.1, consider the following two adversaries \mathcal{B}' and \mathcal{B}'' .

- \mathcal{B}' is the same as \mathcal{B} except that \mathcal{B}' samples $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ from P .
- \mathcal{B}'' is the same as \mathcal{B} except that 1. \mathcal{B}'' samples $(\mathbf{r}_1, \dots, \mathbf{r}_Q) \leftarrow_{\$} \mathcal{D}_\Sigma^{mQ}$ and computes secrets as $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \leftarrow \Phi(\mathbf{r}_1, \dots, \mathbf{r}_Q)$; 2. after \mathcal{A} returns, \mathcal{B}'' outputs $\mathbf{x} \leftarrow \hat{\mathbf{s}} - \sum_{i \in [Q]} \hat{b}_i \mathbf{r}_i$.

We note that \mathcal{B}' and \mathcal{B}'' do not need to be efficient. They are only used to compute the winning probability of the (efficient) adversary \mathcal{B} .

Denote now by $\text{PVal}_{\mathcal{X}}$ the probability that $\mathcal{X} \in \{\mathcal{B}, \mathcal{B}', \mathcal{B}''\}$ outputs a non-zero vector. Then, we will show the following three facts.

Fact 1. $\text{PVal}_{\mathcal{B}'} \leq \delta_\alpha \text{PVal}_{\mathcal{B}}^{(\alpha-1)/\alpha}$;

Fact 2. $\text{PVal}_{\mathcal{B}} = \text{PVal}_{\mathcal{B}''}$;

Fact 3. $\text{PVal}_{\mathcal{B}'} + \text{PVal}_{\mathcal{B}''} \geq \text{Adv}^{\mathbf{G}_1}(\mathcal{A}, \kappa)$.

We can conclude the proof from the facts since

$$\begin{aligned} \text{Adv}^{\mathbf{G}_1}(\mathcal{A}, \kappa) &\leq \text{PVal}_{\mathcal{B}'} + \text{PVal}_{\mathcal{B}''} \leq \delta_\alpha \text{PVal}_{\mathcal{B}}^{(\alpha-1)/\alpha} + \text{PVal}_{\mathcal{B}} \\ &\leq 2\delta_\alpha \text{PVal}_{\mathcal{B}}^{(\alpha-1)/\alpha} = 2\delta_\alpha \left(\text{Adv}_{q,k,m-1,\beta_{\text{sis}}}^{\text{td-msis}}(\mathcal{B}, \kappa) + Q \cdot 2^{-2\kappa+2} \right)^{(\alpha-1)/\alpha}, \end{aligned}$$

where the last inequality is due to the fact that $\delta_\alpha \geq 1$ and $(\alpha - 1)/\alpha < 1$ and the last equation is due to Equation (12.4).

We now show the above three facts. For Fact 1, since the only difference between \mathcal{B} and \mathcal{B}' is the distribution of $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$, by Lemma 11.2.5,

$$\text{PVal}_{\mathcal{B}'} \leq \left(\text{PVal}_{\mathcal{B}} \cdot R_\alpha \left(P \parallel \mathcal{D}_\Sigma^{mQ} \right) \right)^{(\alpha-1)/\alpha}.$$

Therefore, Fact 1 follows from the following lemma, which we will prove below.

Lemma 12.3.3. *For the discrete Gaussian distribution \mathcal{D}_Σ^{mQ} (defined in Section 11.2.2), where $\Sigma = \mathbb{I}_m \otimes \text{diag}(\sigma_1^2, \dots, \sigma_Q^2)$, and the distribution $P = \Phi(\mathcal{D}_\Sigma^{mQ})$,*

$$R_\alpha \left(P \parallel \mathcal{D}_\Sigma^{mQ} \right) \leq \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^{\frac{\alpha}{\alpha-1}} \cdot \exp \left(\alpha \pi^2 / \log(2(1 + 1/\varepsilon)) \right).$$

The arguments for Facts 2 and 3 follow exactly as in Section 11.1.1. We repeat here for completeness. For Fact 2, since $\text{View}_A(\mathbf{r}_1, \dots, \mathbf{r}_Q) = \text{View}_A(\Phi(\mathbf{r}_1, \dots, \mathbf{r}_Q)) = \text{View}_A(\mathbf{s}_1, \dots, \mathbf{s}_Q)$, even if \mathcal{B}'' sets the secrets to $(\mathbf{r}_1, \dots, \mathbf{r}_Q)$ instead of $\mathbf{s}_1, \dots, \mathbf{s}_Q$, the output of \mathcal{B}'' remains the same. However, then, \mathcal{B}'' is identical to \mathcal{B} , which implies the second fact.

Finally, to prove the third fact, we can interpret the sampling process of \mathcal{B}' as first sampling $(\mathbf{r}_1, \dots, \mathbf{r}_Q)$ from \mathcal{P} and then setting $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \leftarrow \Phi(\mathbf{r}_1, \dots, \mathbf{r}_Q)$. Then, the only difference between \mathcal{B}' and \mathcal{B}'' is that \mathcal{B}' returns $\hat{\mathbf{s}} - \sum_{i \in [Q]} \hat{b}_i \mathbf{s}_i$, while \mathcal{B}'' returns $\hat{\mathbf{s}} - \sum_{i \in [Q]} \hat{b}_i \mathbf{r}_i$. If \mathcal{A} wins the game \mathbf{G}_1 simulated by \mathcal{B}' or \mathcal{B}'' , we know $\sum_{i \in [Q]} \hat{b}_i \mathbf{r}_i = \sum_{i \in [Q]} \hat{b}_i (\mathbf{s}_i - u_i \Delta) = \sum_{i \in [Q]} \hat{b}_i \mathbf{s}_i - \Delta \cdot \hat{\mathbf{b}}^T \mathbf{u} \neq \sum_{i \in [Q]} \hat{b}_i \mathbf{s}_i$ by the definition of Φ , and thus, at least one of \mathcal{B}' and \mathcal{B}'' outputs a non-zero vector, which implies Fact 3. \square

Proof of Lemma 12.3.3. We first partition the support R^{mQ} into disjoint sets, then show that the Rényi divergence conditioning on each set is small, and finally use Lemma 11.2.6 to conclude the lemma. Denote \mathcal{S}_L as the set of $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ such that the event A_{Win} does not occur. Since Φ is the identity function over \mathcal{S}_L , $\mathcal{D}_\Sigma^{mQ}(\mathcal{S}_L) = P(\mathcal{S}_L)$. For each $(\mathbf{s}_1, \dots, \mathbf{s}_Q) \notin \mathcal{S}_L$, denote $\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q] := \{(\mathbf{s}_1 + ku_1 \Delta, \dots, \mathbf{s}_Q + ku_Q \Delta)\}_{k \in \mathbb{Z}}$, where \mathbf{u} is output by \mathcal{A} given $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. By a similar argument as above, we know any $(\mathbf{s}'_1, \dots, \mathbf{s}'_Q) \in \mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]$ leads to the same view of \mathcal{A} . Therefore, $\Phi(\mathbf{s}_1 + ku_1 \Delta, \dots, \mathbf{s}_Q + ku_Q \Delta) = (\mathbf{s}_1 + (k+1)u_1 \Delta, \dots, \mathbf{s}_Q + (k+1)u_Q \Delta)$, and Φ is a bijection over $\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]$, which implies $\mathcal{D}_\Sigma^{mQ}(\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]) = P(\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q])$. Also, for any $(\mathbf{s}_1, \dots, \mathbf{s}_Q), (\mathbf{s}'_1, \dots, \mathbf{s}'_Q) \notin \mathcal{S}_L$, $\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]$ and $\mathcal{S}_W[\mathbf{s}'_1, \dots, \mathbf{s}'_Q]$ are either equal

to disjoint. Therefore, R^{mQ} can be partitioned into disjoint sets $\{\mathcal{S}_L\} \cup \{\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]\}_{(\mathbf{s}_1, \dots, \mathbf{s}_Q) \in R^{mQ} \setminus \mathcal{S}_L}$, and for each set, the probability that $(\mathbf{s}_1, \dots, \mathbf{s}_Q)$ falls in the set is equal under both distributions \mathcal{D}_Σ^{mQ} and P .

Therefore, by Lemma 11.2.6, we just need to show the Rényi divergence conditioning on each set is small. For \mathcal{S}_L , since Φ is the identity function over \mathcal{S}_L , $R_\alpha \left(P_{|\mathcal{S}_L} \| \mathcal{D}_\Sigma^{mQ} |_{\mathcal{S}_L} \right) = 1$.

For any $\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]$, we show in the following that the conditioned distribution is identical to a one-dimensional discrete Gaussian distribution under a linear transformation. Let $\sqrt{\Sigma}^{-1} = \mathbb{I}_m \otimes \text{diag}(1/\sigma_1, \dots, 1/\sigma_Q)$. Denote $\mathbf{X} := \sqrt{\Sigma}^{-1}(u_1\Delta, \dots, u_Q\Delta)$ and $\mathbf{S} := \sqrt{\Sigma}^{-1}(\mathbf{s}_1, \dots, \mathbf{s}_Q)$. Denote $\mathbf{S}_\perp := \mathbf{S} - s_0\mathbf{X}$, where $s_0 := \langle \mathbf{S}, \mathbf{X} \rangle / \langle \mathbf{X}, \mathbf{X} \rangle \in \mathbb{R}$, and we have $\langle \mathbf{S}_\perp, \mathbf{X} \rangle = \langle \mathbf{S}, \mathbf{X} \rangle - s_0 \langle \mathbf{X}, \mathbf{X} \rangle = 0$. Then, for any $k \in \mathbb{Z}$,

$$\begin{aligned} \rho_\sigma(\mathbf{s}_1 + ku_1\Delta, \dots, \mathbf{s}_Q + ku_Q\Delta) &= \exp \left(-\pi \left\| \sqrt{\Sigma}^{-1}(\mathbf{s}_1 + ku_1\Delta, \dots, \mathbf{s}_Q + ku_Q\Delta) \right\|^2 \right) \\ &= \exp \left(-\pi \|\mathbf{S} + k\mathbf{X}\|^2 \right) \\ &= \exp \left(-\pi (\|\mathbf{S}_\perp\|^2 + (s_0 + k)^2 \|\mathbf{X}\|^2) \right) \\ &\propto \exp \left(-\pi (-s_0 - k)^2 \|\mathbf{X}\|^2 \right) . \end{aligned}$$

Therefore, $\mathcal{D}_\Sigma^{mQ} |_{\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]} = T \left(\mathcal{D}_{\|\mathbf{X}\|^{-1}, -s_0} \right)$, where $T : \mathbb{Z} \rightarrow R^{mQ}$ maps k to $(\mathbf{s}_1 + ku_1\Delta, \dots, \mathbf{s}_Q + ku_Q\Delta)$. Since $\Phi(\mathbf{s}_1 + (k-1)u_1\Delta, \dots, \mathbf{s}_Q + (k-1)u_Q\Delta) = (\mathbf{s}_1 + ku_1\Delta, \dots, \mathbf{s}_Q + ku_Q\Delta)$, we know $P(\mathbf{s}_1 + ku_1\Delta, \dots, \mathbf{s}_Q + ku_Q\Delta) = \mathcal{D}_\Sigma^{mQ}(\mathbf{s}_1 + (k-1)u_1\Delta, \dots, \mathbf{s}_Q + (k-1)u_Q\Delta)$. By a similar argument as above, $P_{|\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]} = T \left(\mathcal{D}_{\|\mathbf{X}\|^{-1}, -s_0+1} \right)$. Since $\|\mathbf{X}\|^2 = \sum_{i \in [Q]} (u_i/\sigma_i)^2 \|\Delta\|^2 = \|(u_1/\sigma_1, \dots, u_Q/\sigma_Q)\|^2 \|\Delta\|^2 \leq \sqrt{mN} \beta_u \beta_{\text{lwe}}$, we have

$$\|\mathbf{X}\|^{-1} \geq 1/(\sqrt{mN} \beta_u \beta_{\text{lwe}}) = \sqrt{\log(2(1+1/\varepsilon))/\pi} \geq \eta_\varepsilon(\mathbb{Z}) .$$

Therefore, by Lemma 11.2.7,

$$\begin{aligned} R \left(\mathcal{D}_\Sigma^{mQ} |_{\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]} \| P_{|\mathcal{S}_W[\mathbf{s}_1, \dots, \mathbf{s}_Q]} \right) &\leq \left(\frac{1+\varepsilon}{1-\varepsilon} \right)^{\frac{\alpha}{\alpha-1}} \cdot \exp(\alpha\pi \|\mathbf{X}\|^2) \\ &\leq \left(\frac{1+\varepsilon}{1-\varepsilon} \right)^{\frac{\alpha}{\alpha-1}} \cdot \exp(\alpha\pi mN \beta_u^2 \beta_{\text{lwe}}^2) \\ &= \left(\frac{1+\varepsilon}{1-\varepsilon} \right)^{\frac{\alpha}{\alpha-1}} \cdot \exp(\alpha\pi^2 / \log(2(1+1/\varepsilon))) . \end{aligned}$$

Therefore, we can conclude the lemma by Lemma 11.2.6. \square

Chapter 13

TWO-ROUND LATTICE-BASED THRESHOLD SIGNATURES

13.1 Analysis of the CATZ Construction

As our first application to threshold signatures, this section applies AOM-MISIS to the analysis of the CATZ construction [41].

We adopt the following definition from [41], simplifying the small coefficient property by bounding only the term required in the security proof (Section 13.1.3).

Definition 13.1.1 (Linear Threshold Secret Sharing with Small Coefficients). *Let $1 < t \leq n$ and B_{ss} be positive integers and \mathbb{G} be an abelian group. A t -out-of- n linear threshold secret sharing scheme $\text{SecSha}_{t,n,B_{ss}}$ for \mathbb{G} consists of two algorithms (Share, Recon) with the following syntax:*

- $\text{Share}(s \in \mathbb{G}; \boldsymbol{\rho} \in \mathbb{G}^K) \Rightarrow (\text{ss}_j)_{j \in [L]} \in \mathbb{G}^L$: *takes as input a secret $s \in \mathbb{G}$ and a randomness vector $\boldsymbol{\rho} \in \mathbb{G}^K$ (sampled uniformly from \mathbb{G}^K), and returns the secret shares $(\text{ss}_j)_{j \in [L]}$. We note that each party $i \in [n]$ has a subset of indices $T_i \subseteq [L]$ such that the share of party i is $(\text{ss}_j)_{j \in T_i}$. We say that the individual share size of party i is $|T_i|$, the **total share size** is L , and the **randomness size** is K .*
- $\text{Recon}(U, (\text{ss}_j)_{j \in \bigcup_{i \in U} T_i}) \Rightarrow s \in \mathbb{G}$: *takes as input a set $U \subseteq [n]$ with $|U| \geq t$ and the secret shares corresponding to each party in U , and returns the reconstructed secret s .*

We require that $\text{SecSha}_{t,n,B_{ss}}$ satisfies the following properties:

- **Linearity**: *The sharing algorithm Share can be written as an integer matrix $M \in \mathbb{Z}^{L \times (K+1)}$ mapping a vector $\mathbf{v} = (s, \rho_1, \dots, \rho_K)^T \in \mathbb{G}^{K+1}$ to $M\mathbf{v} \in \mathbb{G}^L$. We refer to M as the **sharing matrix** of $\text{SecSha}_{t,n,B_{ss}}$. Moreover, for any $U \subseteq [n]$ denote M_U as the matrix M restricted to the rows indexed with $\bigcup_{i \in U} T_i$, the following is also true:*
 - *For any $U \subseteq [n], |U| \geq t$, there exists a **reconstruction coefficient** vector $\boldsymbol{\lambda}^U \in \mathbb{Z}^L$ such that $\lambda_j^U = 0$ for $j \notin \bigcup_{i \in U} T_i$ and $(\boldsymbol{\lambda}^U)^T M = (1, 0, \dots, 0)$. Then, the output of $\text{Recon}(U, \cdot)$ on input $(\text{ss}_j)_{j \in \bigcup_{i \in U} T_i}$ can be written as $\sum_{i \in U} \sum_{j \in T_i} \lambda_j^U \text{ss}_j$. Hence, for*

- $(ss_j)_{j \in [L]} \leftarrow \text{Share}(s; \rho)$ for any $s \in \mathbb{G}$ and $\rho \in \mathbb{G}^K$, we have that $\sum_{i \in U} \sum_{j \in T_i} \lambda_j^U ss_j = s$.
- For any $CS \subseteq [n]$ with $|CS| < t$, there exists a vector $\mathbf{w}^{CS} \in \mathbb{Z}^{K+1}$ such that $w_1 = 1$ and $M_{CS} \mathbf{w}^{CS} = \mathbf{0}$. We call such \mathbf{w}^{CS} the **sweeping vector** of M_{CS} .
 - **Small Coefficients:** For any $U \subseteq [n]$ with $|U| \geq t$, any $i \in U$ and, any $CS \subset [n]$ with $|CS| < t$, it holds that $\sum_{(i, \hat{j}) \in T_i \times [K+1]} \lambda_i^U M_{i, \hat{j}} w_{\hat{j}}^{CS} \leq B_{ss}$.

Also, [41] shows the existence of such secret sharing scheme from the generic construction by Benaloh and Leichter [21], which can be stated as the following lemma.

Lemma 13.1.2 ([41]). *For any $1 < t \leq n$, there exists a t -out-of- n linear threshold secret sharing with small coefficients with total share size $L = O(t^{4.3} n \log n)$ making the individual share size $|T_i| \leq O(t^{4.3} n \log n)$ for $t' = \min(t, n - t)$ and the small coefficient bound $B_{ss} = O(t^{4.3} n (\log n)^2)$.*

13.1.1 Construction and main security theorem

We present the scheme CATZ[SecSha] in Figure 13.1, where $\text{SecSha} = \text{SecSha}_{t, n, B_{ss}}$ is a linear threshold secret sharing scheme with small coefficients, which is defined in Definition 13.1.1. In Figure 13.2, we give the description of the parameters used in the protocol. One small change here is that the signing key is sampled from a discrete Gaussian distribution with standard deviation σ_{sk} instead of a uniform distribution over the set of vectors with ℓ_∞ -norm bounded by σ_{sk} . We note that this does not affect the correctness of the scheme, as the ℓ_2 -norm of \hat{sk} remains bounded by $\sqrt{mN} \sigma_{sk}$ except for a negligible probability, which is the exact property needed in the correctness proof.

For unforgeability, we show that TS-UF-0 security of CATZ is implied by the hardness of the AOM-MISIS problem in the random oracle model, which is formally stated in the following theorem. In particular, we consider an extension of AOM-MISIS (See Remark 12.1 for details) in which some σ_i can be ∞ . The full proof of Theorem 13.1.3 is given in Section 13.1.3.

Theorem 13.1.3 (TS-UF-0 of CATZ). *For any integers $q = q(\kappa), k = k(\kappa), m = m(\kappa)$, any linear threshold secret sharing scheme $\text{SecSha} = \text{SecSha}_{t, n, B_{ss}}$ with small coefficients, (see Definition 13.1.1) and any TS-UF-0 adversary \mathcal{A} making at most $q_s = q_s(\kappa)$ queries to PPO and*

<p><u>Algorithm Setup(1^κ) :</u> $\bar{A} \leftarrow \\$ R_q^{k \times (m-k)} ; A \leftarrow [\bar{A} \mathbb{I}_k]$ Return A</p> <p><u>Algorithm KeyGen() :</u> $\hat{\mathbf{sk}} \leftarrow \\$ \mathcal{D}_{\sigma_{\mathbf{sk}}}^m ; \mathbf{pk} \leftarrow A \hat{\mathbf{sk}} \bmod q$ $\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_K \leftarrow \\$ R_q^m$ $(\mathbf{ss}_1, \dots, \mathbf{ss}_L)^T \leftarrow M \cdot (\hat{\mathbf{sk}}, \boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_K)^T$ For $i \in [n]$ do $\mathbf{sk}_i \leftarrow (\mathbf{ss}_j)_{j \in T_i}$ Return $(\mathbf{pk}, (\mathbf{sk}_i)_{i \in [n]})$</p> <p><u>Algorithm SPP(\mathbf{st}_i) :</u> For $j \in [0..\ell]$ do $\mathbf{r}_j \leftarrow \\$ \mathcal{D}_{\sigma_r}^m$ For $j \in [0..\ell]$ do $\mathbf{R}_j \leftarrow A \mathbf{r}_j \bmod q$ $pp \leftarrow (\mathbf{R}_j)_{j \in [0..\ell]}$ $\mathbf{st}_i.\text{mapPP}(pp) \leftarrow (\mathbf{r}_j)_{j \in [0..\ell]}$ Return (pp, \mathbf{st}_i)</p> <p><u>Algorithm CompPar(\mathbf{pk}, lr) :</u> $\mu \leftarrow lr.\text{msg}$ For $i \in lr.\text{SS}$ do $(b_j)_{j \in [\ell]} \leftarrow H_1(\mathbf{pk}, lr)$ $(\mathbf{R}_{i,j})_{j \in [0..\ell]} \leftarrow lr.\text{PP}(i)$ $\mathbf{R} \leftarrow \sum_{i \in lr.\text{SS}} (\mathbf{R}_{i,0} + \sum_{j \in [\ell]} b_j \mathbf{R}_{i,j})$ $c \leftarrow H_2(\mathbf{pk}, \mu, \mathbf{R})$ Return $(\mathbf{R}, c, (b_j)_{j \in [\ell]})$</p>	<p><u>Algorithm PS(lr, i, \mathbf{st}_i) :</u> $pp_i \leftarrow lr.\text{PP}(i)$ If $\mathbf{st}_i.\text{mapPP}(pp_i) = \perp$ then return (\perp, \mathbf{st}_i) $(\mathbf{r}_j)_{j \in [0..\ell]} \leftarrow \mathbf{st}_i.\text{mapPP}(pp_i)$ $\mathbf{st}_i.\text{mapPP}(pp_i) \leftarrow \perp$ $(\mathbf{R}, c, (b_j)_{j \in [\ell]}) \leftarrow \text{CompPar}(\mathbf{st}_i.\mathbf{pk}, lr)$ $(\mathbf{ss}_j)_{j \in T_i} \leftarrow \mathbf{st}_i.\mathbf{sk}$ $\mathbf{z} \leftarrow \mathbf{r}_0 + \sum_{j \in [\ell]} b_j \cdot \mathbf{r}_j$ $\quad + 2c \cdot \sum_{j \in T_i} \lambda_j^{lr.\text{SS}} \mathbf{ss}_j \bmod q$ Return $((\mathbf{R}, \mathbf{z}), \mathbf{st}_i)$</p> <p><u>Algorithm Agg(PS, \mathbf{st}_0) :</u> $\mathbf{R} \leftarrow \perp ; \mathbf{z} \leftarrow 0$ For $(\mathbf{R}', \mathbf{z}') \in \text{PS}$ do If $\mathbf{R} = \perp$ then $\mathbf{R} \leftarrow \mathbf{R}'$ If $\mathbf{R} \neq \mathbf{R}'$ then return (\perp, \mathbf{st}_0) $\mathbf{z} \leftarrow \mathbf{z} + \mathbf{z}'$ Return $((\mathbf{R}, \mathbf{z}), \mathbf{st}_0)$</p> <p><u>Algorithm Vf($\mathbf{pk}, \mu, sig$) :</u> $(\mathbf{R}, \mathbf{z}) \leftarrow sig$ If $\ \mathbf{z}\ > \beta_z$ then return 0 $c \leftarrow H_2(\mathbf{pk}, \mu, \mathbf{R})$ Return $(A\mathbf{z} = \mathbf{R} + 2c \cdot \mathbf{pk} \bmod q)$</p>
--	---

Figure 13.1: Lattice-based t -out-of- n threshold signatures CATZ[SecSha], where SecSha = SecSha $_{t,n,B_{\text{ss}}}$ is t -out-of- n a linear secret sharing scheme with small coefficients (see Definition 13.1.1). In particular, $\lambda_j^{lr.\text{SS}}$ denotes the reconstruction coefficient. Also, $H_1 : \{0, 1\}^* \rightarrow \mathcal{S}_b^\ell$ and $H_2 : \{0, 1\}^* \rightarrow \mathcal{S}_c$. The algorithms LPP and LR are defined the same as in Figure 10.3.

Parameters	Description
κ	Security parameter
n	Number of signers
t	Threshold for signing
N	$N \geq 2\kappa$, power of two defining the ring R
q	Prime modulus
k	Number of rows of A
m	$m > k$, number of columns of A ,
$\ell + 1$	$\ell = 2\kappa/\log(2N)$, number of nonces for each signer
\mathcal{S}_b	$\mathcal{S}_b = \{\pm 1, \pm X, \dots, \pm X^{N-1}\}$, set for the aggregating coefficients b_j
β_c	Satisfying $2^{\beta_c} \binom{N}{\beta_c} \geq 2^{2\kappa}$, the ℓ_1 -norm of the challenge c ,
\mathcal{S}_c	$\mathcal{S}_c = \{c \in R : \ c\ _\infty = 1, \ c\ _1 = \beta_c\}$, set of the challenges c
σ_{sk}	Standard deviation of the signing key sk
σ_r	Standard deviation of the nonces r_i
ν_{pk}	Only for EKT: satisfying $\lfloor q/2^{\nu_{pk}} \rfloor = \lfloor q/2^{\nu_{pk}} \rfloor$, number of bits saved on pk
ν_r	Only for EKT: satisfying $\lfloor q/2^{\nu_r} \rfloor = \lfloor q/2^{\nu_r} \rfloor$, number of bits saved on h
$q_{\nu_{pk}}, q_{\nu_r}$	Only for EKT: $(q_{\nu_{pk}}, q_{\nu_r}) = (\lfloor q/2^{\nu_{pk}} \rfloor, \lfloor q/2^{\nu_r} \rfloor)$ the rounded moduli
β_z	ℓ_2 -norm bound of a valid signature vector z (or $(z, 2^{\nu_{pk}}h)$ for EKT)

Figure 13.2: Parameters for CATZ and EKT. Some parameters only apply to EKT.

$q_h = q_h(\kappa)$ queries to RO, there exists an AOM-MISIS adversary \mathcal{B} running in time roughly two times that of \mathcal{A} such that

$$\text{Adv}_{\text{CATZ}}^{\text{ts-uf-0}}(\mathcal{A}, \kappa) \leq \sqrt{\mathbf{q} \text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{B}, \kappa) + 8\mathbf{q}^3 2^{-2\kappa}}.$$

where $\mathbf{q} = q_h + q_s + 1$ and $\text{par} = (q, k, m, Q = 1 + K + q_s(1 + \ell), (\sigma_i)_{i \in [Q]}, \beta_s = 2\beta_z, \beta_b = 4\sigma_{sk}\beta_c, \beta_u = 1/\sigma_{sk} + \beta_c B_{ss} \sqrt{8Nq_s}/\sigma_r)$ with $\sigma_1 = \sigma_{sk}$, $\sigma_{1+i} = \infty$ for $i \in [K]$, $\sigma_{1+K+i} = \sigma_r$ for $i \in [q_s(\ell + 1)]$.

For completeness, we recall the correctness theorem from [41], which is needed for parameter selections later.

Theorem 13.1.4 (Correctness of CATZ [41]). *For any integers $1 < t \leq n$, any linear threshold secret sharing scheme $\text{SecSha} = \text{SecSha}_{t,n,B_{ss}}$ with small coefficients, (see Definition 13.1.1) given $\sigma_r \geq 2\sqrt{6mN \log(2mN)\kappa/\pi}$ and $\beta_z \geq \sqrt{mN}(2\beta_c\sigma_{sk} + \sigma_r\sqrt{n(1+\ell)})$, the threshold signature scheme $\text{CATZ}[\text{SecSha}]$ is correct with correctness error $\varepsilon_{\text{cor}} \leq (2 + 4n(\ell + 1)) \cdot 2^{-2\kappa}$.*

13.1.2 Parameter selection

In this section, we first discuss the asymptotic parameters selection derived from the security theorems and the hardness of AOM-MISIS, then compare these parameters with those proposed in [41], and finally estimate the concrete efficiency based on the parameter selections.

ASYMPTOTIC PARAMETER SELECTIONS. Denote β_{lwe} as the norm of the underlying MLWE assumption. Initially, we select $N, m, k, \beta_{\text{lwe}}$ such that N is a power of $N \geq 2\kappa$, $m, k = \text{poly}(\kappa)$, and $\beta_{\text{lwe}} \geq m \log(N)$.¹ (We note that when estimating the concrete efficiency, we will enumerate through plausible $(N, m, k, \beta_{\text{lwe}})$ tuples and pick the one that yields the best efficiency.) Then, we set other parameters as follows.

- Set β_c as the smallest integer such that $2^{\beta_c} \binom{N}{\beta_c} \geq 2^{2\kappa}$.
- $\sigma_{sk} = \max\{2\beta_{\text{lwe}}\sqrt{mN}, \sqrt{\log(6mN)/\pi}\}$. The first term is usually the leading term.
- $\sigma_r = \max\{\sigma_{sk}\beta_c B_{ss}\sqrt{8Nq_s}, 2\sqrt{6mN \log(2mN)\kappa/\pi}\}$. The first term is usually the leading term.
- $\beta_z = \sqrt{mN}(2\beta_c\sigma_{sk} + \sigma_r\sqrt{n(1+\ell)})$
- Denote $\beta_{\text{sis}} = 2\beta_z + 4\sigma_{sk}\beta_c\sqrt{mN}$.
- Select q such that the problem $\text{MSIS}_{q,k,m,\beta_{\text{sis}}}$ and the problem $\text{MLWE}_{q,N,k,m,\beta_{\text{lwe}}}$ are assumed to be exponentially hard in κ .

By Theorem 13.1.3 and Theorem 12.3.1 with $\varepsilon = 1/2$ and $\alpha = 2$ (we can further optimize the concrete bound by adjusting α), TS-UF-0 of CATZ is implied by the hardness of $\text{MSIS}_{q,k,m,\beta_{\text{sis}}}$ and $\text{MLWE}_{q,N,k,m,\beta_{\text{lwe}}}$.

¹This is for guaranteeing the underlying MLWE is hard.

	$\log_2(q)$	k	m	σ_{sk}	σ_r	β_z	$ pk $	$ sig $	Comm.
[41]	119	8	50	2^{18}	$2^{105.06}$	$2^{117.25}$	60.73KB	440.32KB	2.02MB
This thesis	94.1	7	17	$2^{9.1}$	$2^{80.5}$	$2^{91.96}$	42.14KB	144.47KB	1.24MB

Figure 13.3: The concrete parameters and estimated efficiency for $\kappa = 128$ and $n = 32$ in [41] and this work. We set $(N, \ell, \beta_c) = (512, 26, 64)$. The last second column denotes the communication complexity per signer.

COMPARED WITH [41] The differences of the parameter selections between [41] and ours shows in the selections of σ_{sk} and σ_r . In particular, the prior work requires

- $\sigma_{sk} = q^{k/m} 2^{2\kappa/(Nm)}$.
- $\sigma_r = \max\{N\beta_c B_{ss} \sigma_{sk} \sqrt{32\pi q_s m N}, \frac{16N\sqrt{3m}}{\sqrt{\pi}} q^{\frac{k}{m}} \sqrt{N(\log(2mN) + 2\kappa)}\}$. The first term is usually the leading term.

For σ_{sk} , $q^{k/m}$ is significantly larger than $2\beta_{lwe}\sqrt{mN}$, which influences the choice of k and m . In particular, one needs to set m to be several times larger than k to ensure σ_{sk} does not cause the parameters to grow excessively. For σ_r , we can see a factor of $N\sqrt{m}$ improvement by comparing the first term inside the maximization.

CONCRETE EFFICIENCY. We show a set of concrete parameters and estimated efficiency for $\kappa = 128$ and $n = 32$, and compare them with those from [41] in Figure 13.3, where we can see improvements in both signature sizes and communication complexity. We derive the parameters following our parameter selections mentioned above, where SecSha are instantiated using Lemma 13.1.2. To estimate the concrete hardness of MSIS and MLWE, we use the state-of-art tool, lattice estimator,² which is also used in other prior works [54, 60, 33].

13.1.3 Security reduction of CATZ

The proof structure generally follows from the security proof of FROST [20], which reduces the security of FROST to the algebraic one-more discrete logarithm assumption, as the CATZ con-

²<https://github.com/malb/lattice-estimator>

struction can be seen as a lattice analog of FROST. The key differences lie in how the reduction computes \mathbf{u} and how we bound the norm of the output solution $(\hat{\mathbf{s}}, \hat{\mathbf{b}})$ and \mathbf{u} .

To prove Theorem 13.1.3, we slightly generalize the forking lemma (Lemma 9.2.3) such that the hash values (h_1, \dots, h_1) are independently sampled by an algorithm HG instead of uniformly sampled from a set H .

Proof of Theorem 13.1.3. Let \mathcal{A} be a TS-UF-0 adversary as described in the theorem. W.l.o.g. we assume that \mathcal{A} is deterministic and corrupts exactly $t - 1$ signers. Also, we assume if \mathcal{A} returns $(\mu^*, (\mathbf{R}^*, \mathbf{z}^*))$, the RO query $H_2(\text{pk}, \mu^*, \mathbf{R}^*)$ was made by \mathcal{A} , which adds at most one RO query. Also, since the game makes at most one RO query to H_1 and H_2 respectively for each signing query, the total number of RO queries to each of H_1 and H_2 is bounded $q = q_h + q_s + 1$. We first construct an algorithm \mathcal{C} compatible with the syntax in Lemma 9.2.3 and construct \mathcal{B} from $\text{Fork}^{\mathcal{C}}$.

CONSTRUCTION OF \mathcal{C} . The input of \mathcal{C} consists of $A, (\mathbf{t}_i)_{i \in [K+1+q_s(\ell+1)]}$, and a list of hash values h_1, \dots, h_{2q} , where $(A, (\mathbf{t}_i)_{i \in [K+1+q_s(\ell+1)]})$ are sampled following the AOM-MISIS game, and for each $i \in [q]$, h_{2i-1} is sampled uniformly from \mathcal{S}_b and h_{2i} is sampled uniformly from \mathcal{S}_c . To start with, \mathcal{C} sets $\text{par} \leftarrow A$, initializes $\text{st}_i.\text{mapPP} \leftarrow ()$ for $i \in [n]$, and in addition, initializes a counter $\text{ctr}_h \leftarrow 0$ for counting the number of random oracle queries. Then, \mathcal{C} runs \mathcal{A} on input par with access to oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}$ and $\widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$: \mathcal{C} sets $\text{pk} \leftarrow \mathbf{t}_1$ and views $\mathbf{t}_{1+\hat{j}}$ as $A\rho_{\hat{j}}$ for $\hat{j} \in [K]$. Then, for each $i \in CS$ and $j \in T_i$, \mathcal{C} computes $\text{ss}_j \leftarrow \text{PI}(\mathbf{d})$ with

$$d_{\hat{j}} = \begin{cases} M_{j,\hat{j}}, & \hat{j} \in [K + 1], \\ 0, & \text{o.w.} \end{cases} \quad (13.1)$$

Finally, \mathcal{C} returns $(\text{pk}, (\text{sk}_i = (\text{ss}_j)_{j \in T_i})_{i \in CS})$.

$\widetilde{\text{PPO}}(i)$: For the j -th query, \mathcal{C} sets $\mathbf{R}_{\hat{j}} \leftarrow \mathbf{t}_{K+1+(j-1)(\ell+1)+\hat{j}+1}$ for $\hat{j} \in [0..\ell]$. Since \mathcal{C} does not sample $\{\mathbf{r}_{\hat{j}}\}_{\hat{j} \in [0..\ell]}$, \mathcal{C} uses $\text{st}_i.\text{mapPP}$ to store the index j instead, i.e., \mathcal{C} sets $\text{st}_i.\text{mapPP}(\{\mathbf{R}_{\hat{j}}\}_{\hat{j} \in [0..\ell]}) \leftarrow j$.

$\widetilde{\text{PSIGNO}}(i, lr)$: The same as $\text{PSIGNO}(i, lr)$ except that \mathcal{C} computes \mathbf{z} using the PI oracle as fol-

lows. Let $j \leftarrow \text{st}_i.\text{mapPP}(lr.\text{PP}(i))$ and \mathbf{d} be a vector in R^Q such that

$$d_{\hat{j}} = \begin{cases} 2c \cdot \sum_{\hat{i} \in T_i} \lambda_i^{lr,SS} M_{\hat{i},\hat{j}}, & \hat{j} \in [K+1], \\ 1, & \hat{j} = K+1 + (j-1)(\ell+1) + 1, \\ b_{j'}, & \hat{j} = K+1 + (j-1)(\ell+1) + 1 + j', j' \in [\ell], \\ 0, & o.w. \end{cases} \quad (13.2)$$

\mathcal{C} computes $\mathbf{z} \leftarrow \text{PI}(\mathbf{d})$.

$\widetilde{\text{RO}}$ query $H_1(x)$: If $H_1(x) \neq \perp$, \mathcal{C} returns $H_1(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, lr)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_1(x) \leftarrow_s \mathcal{S}_b^\ell$ and returns $H_1(x)$. Otherwise, \mathcal{C} increases ctr_h by 1, sets $H_1(x) \leftarrow h_{2\text{ctr}_h-1}$. Also, \mathcal{C} computes $\mathbf{R} \leftarrow \sum_{i \in lr.SS} (\mathbf{R}_{i,0} + \sum_{j \in [\ell]} b_j \cdot \mathbf{R}_{i,j})$, where $(\mathbf{R}_{i,j})_{j \in [0..\ell]} \leftarrow lr.\text{PP}(i)$ and $\{b_j\}_{j \in [\ell]} \leftarrow h_{2\text{ctr}_h-1}$. If $H_2(\text{pk}, lr.\text{msg}, \mathbf{R}) = \perp$, \mathcal{C} sets $H_2(\text{pk}, lr.\text{msg}, \mathbf{R}) \leftarrow h_{2\text{ctr}_h}$. Finally, \mathcal{C} returns $H_1(x)$.

$\widetilde{\text{RO}}$ query $H_2(x)$: If $H_2(x) \neq \perp$, \mathcal{C} returns $H_2(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, \mu, \mathbf{R})$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_2(x) \leftarrow_s \mathcal{S}_c$. Otherwise, \mathcal{C} increases ctr_h by 1 and sets $H_2(x) \leftarrow h_{2\text{ctr}_h}$. Finally, \mathcal{C} returns $H_2(x)$.

After receiving the output $(\mu^*, (\mathbf{R}^*, \mathbf{z}^*))$ from \mathcal{A} , \mathcal{C} aborts if \mathcal{A} does not win the TS-UF-0 game. Otherwise, \mathcal{C} finds the index I such that $H_2(\text{pk}, \mu^*, \mathbf{R}^*)$ is set to h_I during the simulation. By our assumption of \mathcal{A} , we know such I must exist. Then, \mathcal{C} returns $(I, \text{Out} = (\mu^*, \mathbf{R}^*, \mathbf{z}^*))$.

ANALYSIS OF \mathcal{C} . To use Lemma 9.2.3, we define $S := \{2i\}_{i \in [q]}$, IG as the algorithm that samples $(A, (\mathbf{t}_i)_{i \in [K+1+q_s(\ell+1)]})$ following the AOM-MISIS game, and HG as the algorithm that samples h_{2i-1} uniformly from \mathcal{S}_b and samples h_{2i} uniformly from \mathcal{S}_c for each $i \in [q]$. From the simulation, we know that the output index I of \mathcal{C} is always in S . Also, it is not hard to check that \mathcal{C} simulates the game TS-SUF-0 perfectly, which implies $\text{acc}(\mathcal{C}) \geq \text{Adv}_{\text{CATZ}}^{\text{ts-uf-0}}(\mathcal{A}, \kappa)$. By Lemma 9.2.3, we have that

$$\text{acc}(\text{Fork}^{\mathcal{C}}) \geq \text{Adv}_{\text{CATZ}}^{\text{ts-uf-0}}(\mathcal{A}, \kappa)^2 / q.$$

CONSTRUCT \mathcal{B} FROM $\text{Fork}^{\mathcal{C}}$. We now construct the AOM-MISIS adversary \mathcal{B} using $\text{Fork}^{\mathcal{C}}$. To start with, \mathcal{B} receives $(A, \{\mathbf{t}_i\}_{i \in [Q]})$ from the AOM-MISIS game with $Q = K+1+q_s(\ell+1)$ and runs $\text{Fork}^{\mathcal{C}}(A, \{\mathbf{t}_i\}_{i \in [Q]})$ with access to the PI oracle from the AOM-MISIS game. If $\text{Fork}^{\mathcal{C}}$

outputs $(I, \text{Out} = (\mu^*, \mathbf{R}^*, \mathbf{z}^*), \overline{\text{Out}} = (\bar{\mu}^*, \bar{\mathbf{R}}^*, \bar{\mathbf{z}}^*))$, we know $A\mathbf{z}^* = \mathbf{R}^* + 2h_I \text{pk} \bmod q$ and $A\bar{\mathbf{z}}^* = \bar{\mathbf{R}}^* + 2\bar{h}_I \text{pk} \bmod q$, which implies $A(\mathbf{z}^* - \bar{\mathbf{z}}^*) = 2(h_I - \bar{h}_I) \text{pk} \bmod q$. Therefore, \mathcal{B} sets $\hat{\mathbf{s}} \leftarrow \mathbf{z}^* - \bar{\mathbf{z}}^*$ and $\hat{\mathbf{b}} \leftarrow (2(h_I - \bar{h}_I), 0, \dots, 0)$, and it holds that $\sum_{i \in [Q]} \hat{b}_i \mathbf{t}_i = 2(h_I - \bar{h}_I) \mathbf{t}_0 = 2(h_I - \bar{h}_I) \text{pk} = A(\mathbf{z}^* - \bar{\mathbf{z}}^*) = \hat{\mathbf{s}} \bmod q$.

We now show how \mathcal{B} sets \mathbf{u} such that $\hat{\mathbf{b}}^T \mathbf{u} \neq 0$ and $\mathbf{d}^T \mathbf{u} = 0$ for any oracle query $\text{PI}(\mathbf{d})$.

By the linearity of SecSha, there exists a sweeping vector $\mathbf{w} \in \mathbb{Z}^{K+1}$ such that $M_{CS} \mathbf{w} = \mathbf{0}$ and $w_1 = 1$, and we set $u_{[K+1]} = \mathbf{w}$. Therefore, $\hat{\mathbf{b}}^T \mathbf{u} = \hat{b}_1 u_1 = \hat{b}_1 \neq 0$. Also, for each PI query made during the execution of $\widetilde{\text{INIT}}$, the query is of the form $\mathbf{d} = (M_{j,1}, \dots, M_{j,K+1}, 0, \dots, 0)$, where $j \in \bigcup_{i \in CS} T_i$, and thus $\mathbf{d}^T \mathbf{u} = (M_{j,1}, \dots, M_{j,K+1}) \cdot \mathbf{w} = 0$.

For $j \in [q_s]$, \mathcal{B} sets $u_{K+1+(j-1)(\ell+1)+[\ell+1]}$ as follows. To simplify notation in the following analysis, we use \mathbf{v} to denote the vector $u_{K+1+(j-1)(\ell+1)+[\ell+1]} \in R^{\ell+1}$. We say a $\widetilde{\text{PSIGNO}}$ query (i, lr) corresponds to the j -th token if and only if it is the valid query with $\text{st}_i.\text{mapPP}(lr.\text{PP}(i)) = j$, where a valid query means the $\widetilde{\text{PSIGNO}}$ oracle does not return \perp . From the simulation, there is at most one $\widetilde{\text{PSIGNO}}$ query corresponding to the j -th token during each execution of \mathcal{A} . Therefore, there are the following cases:

Case 1: No query corresponds to the j -th token during both executions. In this case, \mathcal{B} set $\mathbf{v} \leftarrow \mathbf{0}$.

Case 2: Only one query corresponds to the j -th token during the two executions. Denote \mathbf{d} as the PI query made during the execution of the $\widetilde{\text{PSIGNO}}$ query corresponding to the j -th token, where \mathbf{d} follows the form given in Equation (13.2). \mathcal{B} sets $\mathbf{v} \leftarrow (-\sum_{\hat{j} \in [K+1]} d_{\hat{j}} u_{\hat{j}}, 0, \dots, 0)$, which implies $\mathbf{d}^T \mathbf{u} = \sum_{\hat{j} \in [K+1]} d_{\hat{j}} u_{\hat{j}} + v_1 = 0$.

Case 3: There is one query corresponding to the j -th token during each of the two executions. Denote \mathbf{d} (resp. $\bar{\mathbf{d}}$) as the PI query made during the execution of the $\widetilde{\text{PSIGNO}}$ query corresponding to the j -th token before (resp. after) rewinding. If $\mathbf{d} = \bar{\mathbf{d}}$, then \mathcal{C} sets \mathbf{v} in the same way as Case 2. Otherwise, let $\hat{k} \in [\ell]$ be the index such that $d_{K+1+(j-1)(\ell+1)+1+\hat{k}} \neq \bar{d}_{K+1+(j-1)(\ell+1)+1+\hat{k}}$. (If such \hat{k} does not exist, \mathcal{B} aborts.) Denote $b := d_{K+1+(j-1)(\ell+1)+1+\hat{k}}$ and $\bar{b} := \bar{d}_{K+1+(j-1)(\ell+1)+1+\hat{k}}$. By Equation (13.2), we know that 2 divides $d_{\hat{j}}$ and $\bar{d}_{\hat{j}}$ for $\hat{j} \in [K+1]$. Therefore, denote $\Delta := \sum_{\hat{j} \in [K+1]} (d_{\hat{j}}/2) u_{\hat{j}}$ and $\bar{\Delta} := \sum_{\hat{j} \in [K+1]} (\bar{d}_{\hat{j}}/2) u_{\hat{j}}$. Since

$b, \bar{b} \in \mathcal{S}_b$, by Lemma 11.2.1, there exists $\gamma \in R$ such that $\gamma(b - \bar{b}) = 2 \pmod q$. \mathcal{C} sets

$$v_{\hat{j}} \leftarrow \begin{cases} -2\Delta + b\gamma(\Delta - \bar{\Delta}), & \hat{j} = 1, \\ -\gamma(\Delta - \bar{\Delta}), & \hat{j} = \hat{k}, \\ 0, & \text{o.w.} \end{cases}$$

Then, it holds that $\mathbf{d}^T \mathbf{u} = \sum_{\hat{j} \in [K+1]} d_{\hat{j}} u_{\hat{j}} + v_1 + b v_{\hat{k}} = 2\Delta - 2\Delta + b\gamma(\Delta - \bar{\Delta}) - b\gamma(\Delta - \bar{\Delta}) = 0$ and $\bar{\mathbf{d}}^T \mathbf{u} = \sum_{\hat{j} \in [K+1]} \bar{d}_{\hat{j}} u_{\hat{j}} + v_1 + \bar{b} v_{\hat{k}} = 2\bar{\Delta} - 2\Delta + b\gamma(\Delta - \bar{\Delta}) - \bar{b}\gamma(\Delta - \bar{\Delta}) = 2\bar{\Delta} - 2\Delta + (b - \bar{b})\gamma(\Delta - \bar{\Delta}) = 2\bar{\Delta} - 2\Delta + 2(\Delta - \bar{\Delta}) = 0$.

ANALYSIS OF \mathcal{B} . Denote BadHash as the event that $h_1, \bar{h}_1, \dots, h_{2q}, \bar{h}_{2q}$ are *not* all distinct. We now show that \mathcal{B} wins the AOM-MISIS game if Fork^C returns and BadHash does not occur.

We first show that if Fork^C returns and BadHash does not occur, \mathcal{B} does not abort. Suppose Fork^C returns and BadHash does not occur. Then, the only step where \mathcal{B} might abort is in Case 3 above. In Case 3, suppose $\mathbf{d} \neq \bar{\mathbf{d}}$ and $d_{K+1+(j-1)(\ell+1)+1+[\ell]} = \bar{d}_{K+1+(j-1)(\ell+1)+1+[\ell]}$ (in which case \mathcal{B} aborts). Let (i, lr) be the $\widetilde{\text{PSIGNO}}$ query that corresponds to the j -th token *before* rewinding. Then, there exists $J \in [q]$ such that $h_{2J-1} = H_1(\text{pk}, lr) = d_{K+1+(j-1)(\ell+1)+1+[\ell]}$. Since BadHash does not occur, the only possibility is that the $2J - 1 < I$ and (i, lr) is also the $\widetilde{\text{PSIGNO}}$ query that corresponds to the j -th token *after* rewinding. Let \mathbf{R} be the aggregated nonce computed from CompPar(pk, lr). Denote $\hat{J} \in [q]$ be the index such that $h_{2\hat{J}} = H_2(\text{pk}, lr.\text{msg}, \mathbf{R})$ before rewinding. From the simulation of the random oracles, it holds that $\hat{J} \leq J$ and thus $2\hat{J} \leq 2J \leq I$. Also, since $lr.\text{msg} \neq \mu^*$ (o.w., \mathcal{C} would not win the game), we have $2\hat{J} \neq I$ and thus $2\hat{J} < I$. Therefore, $H_2(\text{pk}, lr.\text{msg}, \mathbf{R})$ in the second execution of \mathcal{C} is also $h_{2\hat{J}}$. This implies $\mathbf{d} = \bar{\mathbf{d}}$, which contradicts our assumption.

If \mathcal{B} does not abort, we have $\|\hat{\mathbf{s}}\| \leq \|\mathbf{z}^*\| + \|\bar{\mathbf{z}}^*\| \leq 2\beta_z = \beta_s$ and

$$\left\| (\hat{b}_1 \cdot \sigma_1, \dots, \hat{b}_Q \cdot \sigma_Q) \right\|_1 = \|2(h_I - \bar{h}_I)\sigma_{\text{sk}}\|_1 \leq 4\sigma_{\text{sk}}\beta_c = \beta_b.$$

It is left to bound $\|(u_1/\sigma_1, \dots, u_Q/\sigma_Q)\|$. For $j \in [q_s]$, denote $\mathbf{v} := u_{K+1+(j-1)(\ell+1)+[\ell+1]}$. There are three cases as mentioned in the construction of \mathcal{B} . For the first case, $\|\mathbf{v}\| = 0$. For the second case, $\|\mathbf{v}\| = \left\| 2c \cdot \sum_{(\hat{i}, \hat{j}) \in T_i \times [K+1]} \lambda_i^{lr, \text{SS}} M_{i, \hat{j}} w_{\hat{j}} \right\| \leq 2B_{\text{ss}} \|c\| \leq 2B_{\text{ss}} \sqrt{\beta_c}$. For the third case, $\|\Delta\|_1 = \left\| c \cdot \sum_{(\hat{i}, \hat{j}) \in T_i \times [K+1]} \lambda_i^{lr, \text{SS}} M_{i, \hat{j}} w_{\hat{j}} \right\|_1 \leq B_{\text{ss}} \beta_c$ and similarly $\|\bar{\Delta}\|_1 \leq B_{\text{ss}} \beta_c$. Since $\|\gamma\| \leq \sqrt{N}$

(by Lemma 11.2.1),

$$\begin{aligned} \|\mathbf{v}\|^2 &= \|\gamma(\bar{b}\Delta - b\bar{\Delta}), 0, \dots, 0, -\gamma(\Delta - \bar{\Delta}), 0, \dots, 0\|^2 \\ &\leq 2(\|\Delta\|_1 + \|\bar{\Delta}\|_1)^2 \|\gamma\|^2 \leq 8NB_{\text{ss}}^2\beta_c^2. \end{aligned}$$

Therefore, $\|(u_1/\sigma_1, \dots, u_Q/\sigma_Q)\| \leq 1/\sigma_{\text{sk}} + \sqrt{\sum_{j \in [q_s]} \|\mathbf{v}/\sigma_r\|^2} \leq 1/\beta_s + \beta_c B_{\text{ss}} \sqrt{8Nq_s}/\sigma_r \leq \beta_u$.

Since each of $h_1, h_3, \dots, h_{2q-1}$ and h_2, h_4, \dots, h_{2q} are sampled uniformly from \mathcal{S}_b^ℓ and \mathcal{S}_c respectively, $\Pr[\text{BadHash}] \leq (2q)^2/|\mathcal{S}_b|^\ell + (2q)^2/|\mathcal{S}_c| \leq 8q^2 2^{-2\kappa}$. Therefore,

$$\begin{aligned} \text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{B}, \kappa) &\geq \text{acc}(\text{Fork}^C) - \Pr[\text{BadHash}] \\ &\geq \text{Adv}_{\text{CATZ}}^{\text{ts-uf-0}}(\mathcal{A}, \kappa)^2/q - 8q^2 2^{-2\kappa}, \end{aligned}$$

which concludes the theorem. \square

13.2 Analysis of the EKT Construction

We apply AOM-MISIS to the analysis of the EKT construction and discuss the parameter selections in this section.

13.2.1 Construction and main security theorem

We recall the EKT construction from [60] in Figure 13.4. A summary of the parameters used by the scheme is provided in Figure 13.2. The scheme also depends on a pseudorandom function PRF with suitable domain and range. The Lagrange coefficient $\lambda_{U,j} \in \mathbb{Z}_q$ for any subset $U \subseteq [n]$ and $j \in U$ is defined as $\lambda_{U,j} := \prod_{i \in U \setminus \{j\}} \frac{-i}{j-i}$.

The following theorem establishes TS-UF-4 security for EKT from the hardness of AOM-MISIS. In the theorem statement, $\text{Adv}_{\text{PRF}}^{\text{prf}}(\mathcal{C}, \kappa)$ refers to the standard PRF advantage of an adversary \mathcal{C} . The full proof of Theorem 13.2.1 is given in Section 13.2.3.

Theorem 13.2.1 (TS-UF-4 of EKT). *For any integers $q = q(\kappa), k = k(\kappa), m = m(\kappa)$, any pseudorandom function PRF, and any TS-UF-4 adversary \mathcal{A} making at most $q_s = q_s(\kappa)$ queries to PPO and $q_h = q_h(\kappa)$ queries to RO, given $\sigma_r > 2N \cdot q^{\frac{1}{m} + \frac{1}{N(m-k)}}$, there exist an AOM-MISIS adversary \mathcal{B} and a PRF adversary \mathcal{C} running in time roughly two times that of \mathcal{A} such that*

$$\text{Adv}_{\text{EKT}[\text{PRF}]}^{\text{ts-uf-4}}(\mathcal{A}, \kappa) \leq \sqrt{q \text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{B}, \kappa) + 8q^3 2^{-2\kappa} + n^2 \cdot \text{Adv}_{\text{PRF}}^{\text{prf}}(\mathcal{C}, \kappa) + q_s^2 \cdot 2^{-2\kappa+1}}.$$

<p><u>Algorithm Setup(1^κ) :</u> $\hat{A} \leftarrow_s R_q^{k \times (m-k)} ; A \leftarrow [\hat{A}]_{\mathbb{I}_k}$ Return A</p> <p><u>Algorithm KeyGen() :</u> $\hat{\mathbf{s}}k \leftarrow_s \mathcal{D}_{\sigma_{\mathbf{s}k}}^m ; \mathbf{pk} \leftarrow \left[2A \cdot \hat{\mathbf{s}}k \right]_{\nu_{\mathbf{pk}}}$ For $(i, j) \in [n] \times [n]$ do $\text{seed}_{i,j} \leftarrow \{0, 1\}^\kappa$ $\mathbf{a}_1, \dots, \mathbf{a}_{t-1} \leftarrow_s R_q^{m-k}$ For $i \in [n]$ do $\text{ssk}_i \leftarrow \hat{\mathbf{s}}k_{[m-k]} + \sum_{j=1}^{t-1} \mathbf{a}_j i^j$ $\text{sk}_i \leftarrow (\text{ssk}_i, (\text{seed}_{i,j}, \text{seed}_{j,i})_{j \in [n]})$ Return $(\mathbf{pk}, (\text{sk}_i)_{i \in [n]})$</p> <p><u>Algorithm SPP(st_i) :</u> For $j \in [0..\ell]$ do $\mathbf{r}_j \leftarrow_s \mathcal{D}_{\sigma_r}^m$ For $j \in [0..\ell]$ do $\mathbf{R}_j \leftarrow A \cdot \mathbf{r}_j$ $\text{st}_i.\text{mapPP}((\mathbf{R}_j)_{j \in [0..\ell]}) \leftarrow (\mathbf{r}_j)_{j \in [0..\ell]}$ Return $((\mathbf{R}_j)_{j \in [0..\ell]}, \text{st}_i)$</p> <p><u>Algorithm CompPar(\mathbf{pk}, lr) :</u> For $i \in lr.SS$ do $(b_j)_{j \in [\ell]} \leftarrow H_1(\mathbf{pk}, lr)$ $(\mathbf{R}_{i,j})_{j \in [0..\ell]} \leftarrow lr.PP(i)$ $\mathbf{R} \leftarrow \left[\sum_{i \in lr.SS} \left(\mathbf{R}_{i,0} + \sum_{j \in [\ell]} b_j \mathbf{R}_{i,j} \right) \right]_{\nu_r}$ $c \leftarrow H_2(\mathbf{pk}, lr.\text{msg}, \mathbf{R})$ Return $(\mathbf{R}, c, (b_j)_{j \in [\ell]})$</p>	<p><u>Algorithm PS(lr, i, st_i) :</u> $pp_i \leftarrow lr.PP(i)$ If $\text{st}_i.\text{mapPP}(pp_i) = \perp$ then return (\perp, st_i) $(\mathbf{r}_j)_{j \in [0..\ell]} \leftarrow \text{st}_i.\text{mapPP}(pp_i)$ $\text{st}_i.\text{mapPP}(pp_i) \leftarrow \perp$ $(\mathbf{R}, c, (b_j)_{j \in [\ell]}) \leftarrow \text{CompPar}(\text{st}_i.\mathbf{pk}, lr)$ $(\text{ssk}_i, (\text{seed}_{i,j}, \text{seed}_{j,i})_{j \in [n]}) \leftarrow \text{st}_i.\text{sk}$ $\text{mask} \leftarrow \sum_{j \in lr.SS} \text{PRF}(\text{seed}_{i,j}, (\mathbf{pk}, lr))$ $\text{mask}' \leftarrow \sum_{j \in lr.SS} \text{PRF}(\text{seed}_{j,i}, (\mathbf{pk}, lr))$ $\mathbf{z} \leftarrow \mathbf{r}_{0,[m-k]} + \sum_{j \in [\ell]} b_j \cdot \mathbf{r}_{j,[m-k]}$ $\quad + 2c \cdot \lambda_{lr.SS,j} \cdot \text{ssk}_j + \text{mask} - \text{mask}' \bmod q$ Return $((\mathbf{R}, \mathbf{z}), \text{st}_i)$</p> <p><u>Algorithm Agg(PS, st_0) :</u> $\mathbf{R} \leftarrow \perp ; \mathbf{z} \leftarrow 0$ For $(\mathbf{R}', \mathbf{z}') \in \text{PS}$ do $\mathbf{R} \leftarrow \mathbf{R}' ; \mathbf{z} \leftarrow \mathbf{z} + \mathbf{z}'$ $c \leftarrow H_2(\mathbf{pk}, \mu, \mathbf{R})$ $\mathbf{h} \leftarrow \mathbf{R} - \left[\hat{A}\mathbf{z} - 2^{\nu_{\mathbf{pk}}} \cdot c \cdot \overline{\mathbf{pk}} \right]_{\nu_r} \quad // \mathbf{h} \in R_{q\nu_r}^k$ Return $((c, \mathbf{z}, \mathbf{h}), \text{st}_0)$</p> <p><u>Algorithm Vf($\mathbf{pk}, \mu, sig$) :</u> $(c, \mathbf{z}, \mathbf{h}) \leftarrow sig$ If $\ (\mathbf{z}, 2^{\nu_r} \overline{\mathbf{h}} \bmod q)\ _2 > \beta_z$ then return 0 $c' \leftarrow H_2(\mathbf{pk}, \mu, \left[\hat{A}\mathbf{z} - 2^{\nu_{\mathbf{pk}}} \cdot c \cdot \overline{\mathbf{pk}} \right]_{\nu_r} + \mathbf{h})$ Return $(c' = c)$</p>
--	--

Figure 13.4: Lattice-based t -out-of- n threshold signatures EKT[PRF], where PRF is a pseudo-random function. Here, $H_1 : \{0, 1\}^* \rightarrow \mathcal{S}_b^\ell$ and $H_2 : \{0, 1\}^* \rightarrow \mathcal{S}_c$. Also, $\lambda_{lr.SS,j}$ denotes the Lagrange coefficient, and $\overline{\mathbf{pk}}, \overline{\mathbf{h}} \in R^k$ denote the lift (see Section 11.2.1 for more details) of \mathbf{pk} and \mathbf{h} respectively. The algorithms LPP and LR are defined the same as in Figure 10.3.

where $\mathbf{q} = \mathbf{q}_h + \mathbf{q}_s + 1$ and $\text{par} = (q, k, m, Q = 1 + \mathbf{q}_s(1 + \ell), (\sigma_i)_{i \in [Q]}, \beta_s = (2^{\nu_r+2} + \beta_c \cdot 2^{\nu_{pk}+1}) \cdot \sqrt{Nk} + 4\beta_z, \beta_b = 4\sigma_{sk}\beta_c, \beta_u = 1/\sigma_{sk} + \beta_c\sqrt{8N\mathbf{q}_s}/\sigma_r)$ with $\sigma_1 = \sigma_{sk}, \sigma_{1+i} = \sigma_r$ for $i \in [\mathbf{q}_s(\ell + 1)]$.

Remark 13.1. We note that we can directly establish TS-UF-0 security of EKT from AOM-MISIS by applying the unforgeability theorem from [60] and Lemma 12.2.1. However, the resulting parameters are worse than those in Theorem 13.2.1. Specifically, β_u becomes $1/\sigma_{sk} + 2\beta_c\sqrt{\ell N\mathbf{q}_s}/\sigma_r$, and β_b becomes $4\sigma_{sk}\sqrt{\beta_c N}$.

Also, we recall the correctness theorem from [60], which is needed for parameter selection later.

Theorem 13.2.2 (Correctness of EKT [60]). *For any integers $1 < t \leq n$, any pseudorandom function PRF, given $\sigma_r \geq \sqrt{(\log(2Nm) + \kappa)/\pi}$ and $\beta_z \geq (\beta_c 2^{\nu_{pk}} + 2^{\nu_r})\sqrt{mN} + e^{1/4}(2\beta_c\sigma_{sk} + \sigma_r\sqrt{n(1 + \ell)})\sqrt{N}(\sqrt{k} + \sqrt{m - k})$, the threshold signature scheme EKT[PRF] is correct with correctness error negligible in κ .*

13.2.2 Parameter selection

In this section, we first discuss the asymptotic parameters selection derived from the security theorems and the hardness of AOM-MISIS, then compare these parameters with those proposed in [60], and finally estimate the concrete efficiency based on the parameter selections. We also discuss how our parameters are compared to the parameters of Ringtail [33].

ASYMPTOTIC PARAMETER SELECTIONS. Denote β_{lwe} as the norm of the underlying MLWE assumption. Initially, we select $N, m, k, \beta_{\text{lwe}}$ such that N is a power of $N \geq 2\kappa$, $m, k = \text{poly}(\kappa)$, and $\beta_{\text{lwe}} \geq m \log(N)$.³ (We note that when estimating the concrete efficiency, we will enumerate through plausible $(N, m, k, \beta_{\text{lwe}})$ tuples and pick the one that yields the best efficiency.) Then, we set other parameters as follows.

- Set β_c as the smallest integer such that $2^{\beta_c} \binom{N}{\beta_c} \geq 2^{2\kappa}$.
- $\sigma_{sk} = \max\{2\beta_{\text{lwe}}\sqrt{mN}, \sqrt{\log(6mN)/\pi}\}$. The first term is usually the leading term.
- $\sigma_r = \max\{\sigma_{sk}\beta_c\sqrt{8N\mathbf{q}_s}, \sqrt{(\log(2Nm) + \kappa)/\pi}, 2N \cdot q^{\frac{1}{m} + \frac{1}{N(m-k)}}\}$. The first term is usually the leading term.

³This is for guaranteeing the underlying MLWE is hard.

- $\nu_{pk} = \log_2(\sigma_r/\beta_c)$ and $\nu_r = \log_2(\sigma_r)$.
- β_z is set as shown in Theorem 13.2.2
- Denote $\beta_{\text{sis}} = 2\beta_z + 4\sigma_{\text{sk}}\beta_c\sqrt{mN}$.
- Select q such that the problem $\text{MSIS}_{q,k,m,\beta_{\text{sis}}}$ and the problem $\text{MLWE}_{q,N,k,m,\beta_{\text{lwe}}}$ are assumed to be exponentially hard in κ .

By Theorem 13.2.1 and Theorem 12.3.1 with $\varepsilon = 1/2$ and $\alpha = 2$ (we can further optimize the concrete bound by adjusting α), TS-UF-4 of EKT is implied by the hardness of $\text{MSIS}_{q,k,m,\beta_{\text{sis}}}$ and $\text{MLWE}_{q,N,k,m,\beta_{\text{lwe}}}$.

COMPARISON WITH [60]. Although the prior work does not give a security reduction to standard lattice assumptions, they provide candidate asymptotic parameters based on the heuristic assumption that the selective version of AOM-MLWE is as hard as the adaptive version. Still, our asymptotic parameters are slightly better than their candidate asymptotic parameters provided in [60]. The key difference lies in the choice of σ_r , which significantly impacts efficiency. In particular, the prior work requires $\sigma_r = \Omega(\beta_c\beta_{\text{lwe}}N\sqrt{q_sNk})$, while we require $\sigma_r = \Omega(\beta_c\beta_{\text{lwe}}N\sqrt{q_s m})$. Therefore, there is roughly a factor of \sqrt{N} improvement.

COMPARISON WITH RINGTAIL [33]. Ringtail is very close to EKT. The main differences are that the output space of hash function H_1 changes and the nonces \mathbf{r}_j are sampled from unbalanced discrete Gaussian distributions. In particular, in Ringtail, the first $m - k$ entries of each nonce $\mathbf{r}_{j,[m-k]}$ is sampled from $\mathcal{D}_{\sigma_r}^{m-k}$, while the rest $\mathbf{r}_{j,[(m-k+1)..m]}$ is sampled from $\mathcal{D}_{\sigma'_r}^k$ with $\sigma'_r \neq \sigma_r$. Therefore, it is possible to compare the parameter selections directly.

The key difference still lies in the choice of σ_r . In particular, Ringtail requires $\sigma_r = \Omega(\beta_c\sqrt{q_h})$, where q_h denotes the number of random oracle queries. Here we offer different trade-offs. A key drawback of their parameters is that σ_r depends on q_h , which is typically assumed to be much larger than q_s . This is because q_s refers to the number of online signing queries and is a system parameter that can be enforced, while q_h scales with the offline computational power of the adversary. However, if we set $q_h = q_s$, their σ_r is smaller than ours roughly by a factor of $\beta_{\text{lwe}}N\sqrt{m}$.

CONCRETE EFFICIENCY. We show a set of concrete parameters and estimated efficiency for $\kappa \in \{128, 192, 256\}$ and $n = 1024$ in Figure 13.5. We derive the parameters following our parameter selections mentioned above, and similar to the CTZ scheme, we estimate the concrete hardness of

κ	$\log_2(q)$	k	m	σ_{sk}	σ_r	β_z	$ pk $	$ sig $	Comm.
128	69.3	6	11	$2^{10.6}$	$2^{52.6}$	$2^{66.2}$	8.72KB	30.88KB	766.83KB
192	69.6	8	15	$2^{10.6}$	$2^{52.6}$	$2^{66.5}$	11.74KB	42.90KB	1.01MB
256	70.2	10	18	$2^{11.17}$	$2^{53.2}$	$2^{67.1}$	14.76KB	50.72KB	1.29MB

Figure 13.5: The concrete parameters and estimated efficiency of the EKT scheme for $\kappa = 128, 192, 256$ and $n = 1024$. We set $(N, \ell, \beta_c) = (512, 26, 64)$. The last second column denotes the communication complexity per signer.

MSIS and MLWE using the lattice estimator. We note that our concrete parameters are worse than those given in [60], although in a similar ballpark. However, worse parameters are to be expected. This is because their parameter selection is based on their direct cryptanalysis of AOM-MLWE, whereas we rely on a reduction from two standard lattice assumptions. In a similar spirit, our parameters are worse than those claimed for Ringtail in [33], however this is to be expected, too, as they heuristically assume $q_h = q_s$ (also see the above discussion) to set parameters. In practice, however, we expect q_h to be best approximated conservatively by the running time of the adversary, and this can be as high as 2^{256} , whereas q_s could typically be 2^{60} . The authors of Ringtail were aware of this fact, and their choice was motivated by their conjecture that a better dependency would be possible. We confirm their conjecture for the case of EKT.

13.2.3 Security reduction of EKT

Unlike other security analyses [54, 60] of lattice-based threshold signatures that use masking techniques, our reduction follows a two-step approach, which we believe has independent interest.

In the first step, we reduce the TS-UF-4 game of EKT[PRF] to an *ideal* unforgeability game Ideal-TUF for threshold signatures (defined in Figure 13.6). In this game, no secret sharing of signing key or masking is involved. Moreover, the adversary directly obtains an *aggregation* of all partial signatures from honest parties in $lr.SS$ via a second-round signing oracle SIGNO, provided that the tokens in lr for honest signers are all valid. Intuitively, the ideal game captures the information hidden by the masks. In particular, all the secret key shares and the partial signatures

<p><u>Game Ideal-TUF^A(κ) :</u> $\hat{A} \leftarrow_{\\$} R_q^{k \times (m-k)} ; A \leftarrow [\hat{A}]_{\mathbb{I}_k}$ $H \leftarrow_{\\$} \text{TS.HF} ; S \leftarrow \emptyset$ For $i \in [n]$ do $\text{st}_i.\text{mapPP} \leftarrow ()$ $(\mu, \text{sig}) \leftarrow \mathcal{A}^{\text{INIT,PPO,PSIGNO,RO}}(A)$ Return $(\mu \notin S \wedge \text{Vf}(\text{pk}, \mu, \text{sig}) = 1)$</p> <p><u>Oracle INIT($CS$) :</u> Require: $CS \subseteq [n]$ and $CS < t$ $HS \leftarrow [n] \setminus CS$ $\text{sk} \leftarrow_{\\$} \mathcal{D}_{\sigma_{\text{sk}}}^m$ $\text{pk} \leftarrow [2A \cdot \text{sk}]_{\nu_{\text{pk}}}$ Return pk</p> <p><u>Oracle RO(x) :</u> Return $H(x)$</p>	<p><u>Oracle PPO(i) :</u> Require: $i \in HS$ $\mathbf{r}_j \leftarrow_{\\$} \mathcal{D}_{\sigma_r}^m$ For $j \in [0..\ell]$ do $\mathbf{R}_j \leftarrow A \cdot \mathbf{r}_j$ $pp \leftarrow (\mathbf{R}_j)_{j \in [0..\ell]}$ $\text{st}_i.\text{mapPP}(pp) \leftarrow (\mathbf{r}_j)_{j \in [0..\ell]}$ Return pp</p> <p><u>Oracle SIGNO(lr) :</u> Require: $lr.SS \subseteq [n]$ and $lr.SS \geq t$ $\text{hon} \leftarrow lr.SS \cap HS$ If $\exists i \in \text{hon} : \text{st}_i.\text{mapPP}(lr.PP(i)) = \perp$ then Return \perp $S \leftarrow S \cup \{lr.\text{msg}\}$ $(c, b_1, \dots, b_\ell) \leftarrow \text{CompPar}(\text{pk}, lr)$ $\mathbf{z} \leftarrow 2c \cdot \text{sk}_{[m-k]}$ For $i \in \text{hon}$ do $(\mathbf{r}_j)_{j \in [0..\ell]} \leftarrow lr.PP(i)$ $\text{st}_i.\text{mapPP}(lr.PP(i)) \leftarrow \perp$ $\mathbf{z} \leftarrow \mathbf{z} + \mathbf{r}_{0,[m-k]} + \sum_{j \in [\ell]} b_j \cdot \mathbf{r}_{j,[m-k]}$ Return \mathbf{z}</p>
--	--

Figure 13.6: The Ideal-TUF game, where the algorithms CompPar and Vf are defined in Figure 13.4.

(except their aggregation) are entirely hidden by the masks.

In the second step, we establish the hardness of the ideal unforgeability game based on AOM-MISIS. The approach is cleaner than prior proofs, as it clearly separates the effects of masks from the main security reduction to the AOM-MISIS problem. In particular, the first step relies only on the security of PRF, while the second step does not involve masking at all.

Concretely, Theorem 13.2.1 is a corollary from the following two lemmas.

Lemma 13.2.3. *For any integers $q = q(\kappa), k = k(\kappa), m = m(\kappa)$, any PRF scheme PRF, and any TS-UF-4 adversary \mathcal{A} making at most $q_s = q_s(\kappa)$ queries to PPO and $q_h = q_h(\kappa)$ queries to RO, given $\sigma_r > 2N \cdot q^{\frac{1}{m} + \frac{1}{N(m-k)}}$, there exists an Ideal-TUF adversary \mathcal{B} making at most q_s queries to PPO and q_h queries to RO and an PRF adversary \mathcal{C} running in time roughly the same as \mathcal{A} such that*

$$\text{Adv}_{\text{EKT}[\text{PRF}]}^{\text{ts-uf-4}}(\mathcal{A}, \kappa) \leq \text{Adv}^{\text{ideal-tuf}}(\mathcal{B}, \kappa) + n^2 \cdot \text{Adv}_{\text{PRF}}^{\text{prf}}(\mathcal{C}, \kappa) + q_s^2 \cdot 2^{-2\kappa+1}.$$

where n denotes the number of signers.

Lemma 13.2.4. *For any integers $q = q(\kappa), k = k(\kappa), m = m(\kappa)$ and any Ideal-TUF adversary \mathcal{A} making at most q_s queries to PPO and q_h queries, there exists an AOM-MISIS adversary \mathcal{B} running in time roughly two times that of \mathcal{A} such that*

$$\text{Adv}^{\text{ideal-tuf}}(\mathcal{A}, \kappa) \leq \sqrt{q \text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{B}, \kappa) + 8q^3 2^{-2\kappa}}.$$

where $q = q_h + q_s + 1$ and $\text{par} = (q, k, m, Q = 1 + q_s(1 + \ell), (\sigma_i)_{i \in [Q]}, \beta_s = (2^{\nu_r+2} + \beta_c \cdot 2^{\nu_{pk}+1}) \cdot \sqrt{Nk} + 4\beta_z, \beta_b = 4\sigma_{sk}\beta_c, \beta_u = 1/\sigma_{sk} + \beta_c\sqrt{8Nq_s}/\sigma_r)$ with $\sigma_1 = \sigma_{sk}, \sigma_{1+i} = \sigma_r$ for $i \in [q_s(\ell + 1)]$.

The rough idea behind the first reduction (Lemma 13.2.3) is as follows. Due to masking, the reduction can simulate the PSIGNO oracle by responding with a uniformly random vector, unless the adversary has made PSIGNO queries to all honest party in lr .SS. In this case, the reduction queries its own SIGNO oracle to obtain an aggregated signature and derives the requested partial signature from it. Thus, the reduction only queries SIGNO for message lr .msg when all honest signers in lr .SS were queried. This implies that the set S of messages considered signed in TS-UF-4 is exactly the same as the set S defined in Ideal-TUF. The second reduction (Lemma 13.2.4) is the similar to the proof of Theorem 13.1.3, and its proof is provided in Section 13.2.4.

Proof of Lemma 13.2.3. Let \mathcal{A} be a TS-UF-4 adversary described in the theorem. We show the lemma via the following series of games.

G₀: This is the same as TS-UF-4. The game is formally defined in Figure 13.7

G₁: The same as G₀ except that in the oracle PSIGNO, the response z is computed in a different way, and the game aborts if there are two valid PSIGNO queries for the same input (i, lr) (denoted as BNonce). The game is formally defined in Figure 13.8. We first

<p>Game $G_0^A(\kappa)$:</p> $\hat{A} \leftarrow_{\$} R_q^{k \times (m-k)} ; A \leftarrow [\hat{A}]_{\mathbb{I}_k}$ $H \leftarrow_{\$} \text{TS.HF} ; S \leftarrow \emptyset ; \text{curSSL} \leftarrow ()$ $(\mu^*, \text{sig}^*) \leftarrow \mathcal{A}^{\text{INIT,PPO,PSIGNO,OPEN,RO}}(A)$ <p>Return $(\mu^* \notin S \wedge \text{Vf}(\text{pk}, \mu^*, \text{sig}^*) = 1)$</p> <p>Oracle INIT($CS$) :</p> $HS \leftarrow [n] \setminus CS$ $\hat{\text{sk}} \leftarrow_{\$} \mathcal{D}_{\sigma_{\text{sk}}}^m ; \text{pk} \leftarrow \left[2A \cdot \hat{\text{sk}} \right]_{\nu_{\text{pk}}}$ <p>For $(i, j) \in [n] \times [n]$ do</p> $\text{seed}_{i,j} \leftarrow \{0, 1\}^\kappa$ $\mathbf{a}_1, \dots, \mathbf{a}_{t-1} \leftarrow_{\$} R_q^{m-k}$ <p>For $i \in [n]$ do</p> $\text{ssk}_i \leftarrow \hat{\text{sk}}_{[m-k]} + \sum_{j=1}^{t-1} \mathbf{a}_j i^j$ $\text{sk}_i \leftarrow (\text{ssk}_i, (\text{seed}_{i,j}, \text{seed}_{j,i})_{j \in [n]})$ <p>Return $(\text{sk}_i)_{i \in CS}$</p> <p>Oracle PPO(i) :</p> <p>Require: $i \in HS$</p> <p>For $j \in [0..\ell]$ do $\mathbf{r}_j \leftarrow_{\\$} \mathcal{D}_{\sigma_r}^m$</p> <p>For $j \in [0..\ell]$ do $\mathbf{R}_j \leftarrow A \cdot \mathbf{r}_j$</p> $pp \leftarrow (\mathbf{R}_j)_{j \in [0..\ell]}$ $\text{st}_i.\text{mapPP}(pp) \leftarrow (\mathbf{r}_j)_{j \in [0..\ell]}$ <p>Return pp</p>	<p>Oracle PSIGNO(i, lr) :</p> <p>Require: $lr.SS \subseteq [n]$ and $i \in HS \cap lr.SS$</p> $pp_i \leftarrow lr.PP(i)$ <p>If $\text{st}_i.\text{mapPP}(pp_i) = \perp$ then return \perp</p> <p>If $\text{curSSL}(lr) = \perp$ then</p> $\text{curSSL}(lr) \leftarrow \{i\}$ <p>Else $\text{curSSL}(lr) \leftarrow \text{curSSL}(lr) \cup \{i\}$</p> <p>If $\text{curSSL} = lr.SS \cap HS$ then</p> $S \leftarrow S \cup \{lr.\text{msg}\}$ $(\mathbf{r}_j)_{j \in [0..\ell]} \leftarrow \text{st}_i.\text{mapPP}(pp_i)$ $\text{st}_i.\text{mapPP}(pp_i) \leftarrow \perp$ $(\mathbf{R}, c, (b_j)_{j \in [\ell]}) \leftarrow \text{CompPar}(\text{pk}, lr)$ $(\text{ssk}_i, (\text{seed}_{i,j}, \text{seed}_{j,i})_{j \in [n]}) \leftarrow \text{sk}_i$ $\text{mask} \leftarrow \sum_{j \in lr.SS} \text{PRF}(\text{seed}_{i,j}, (\text{pk}, lr))$ $\text{mask}' \leftarrow \sum_{j \in lr.SS} \text{PRF}(\text{seed}_{j,i}, (\text{pk}, lr))$ $\mathbf{z} \leftarrow \mathbf{r}_{0,[m-k]} + \sum_{j \in [\ell]} b_j \cdot \mathbf{r}_{j,[m-k]} + 2c \cdot \lambda_{lr.SS,j} \cdot \text{ssk}_j + \text{mask} - \text{mask}' \bmod q$ <p>Return (\mathbf{R}, \mathbf{z})</p> <p>Oracle RO(x) :</p> <p>Return $H(x)$</p>
--	---

Figure 13.7: The G_0 game, where the algorithms CompPar and Vf are defined in Figure 13.4.

show that if BNonce does not occur, the game is identical to G_0 . From the description of G_1 , if $\text{curSSL}(lr) \neq \text{hon}$, we have $\mathbf{z} = \mathbf{v} + \text{mask}_c = \mathbf{r}_{0,[m-k]} + (\sum_{j \in [\ell]} b_j \cdot \mathbf{r}_{j,[m-k]}) + 2c \cdot \lambda_{lr.SS,i} \cdot \text{ssk}_i + \text{mask}_h + \text{mask}_c$, which is exactly the same as G_0 . Otherwise, $\mathbf{z} =$

<p>Game $\boxed{G_1^A(\kappa)}$, $G_2^A(\kappa)$, $\boxed{G_3^A(\kappa)}$:</p> <p>Oracle PSIGNO(i, lr) :</p> <p>Require: $lr.SS \subseteq [n]$ and $i \in HS \cap lr.SS$; $pp_i \leftarrow lr.PP(i)$</p> <p>If $st_i.mapPP(pp_i) = \perp$ then return \perp</p> <p>If $i \in curSSL(lr)$ then the game aborts // Bad event BNonce</p> <p>$hon \leftarrow lr.SS \setminus CS$; $cor \leftarrow lr.SS \cap CS$</p> <p>$curSSL(lr) \leftarrow curSSL(lr) \cup \{i\}$</p> <p>If $curSSL = hon$ then $S \leftarrow S \cup \{lr.msg\}$</p> <p>$(R, c, (b_j)_{j \in [\ell]}) \leftarrow CompPar(pk, lr)$</p> <p>$mask_c \leftarrow \sum_{j \in cor} PRF(seed_{i,j}, (pk, lr)) - \sum_{j \in cor} PRF(seed_{j,i}, (pk, lr))$</p> <p>$(r_j)_{j \in [0..\ell]} \leftarrow st_i.mapPP(pp_i)$; $st_i.mapPP(pp_i) \leftarrow \perp$</p> <p>$curSumR(lr) \leftarrow curSumR(lr) + r_{0,[m-k]} + \sum_{j \in [\ell]} b_j \cdot r_{j,[m-k]}$</p> <p>If $curSSL(lr) \neq hon$ then // i is not the last queried honest party in $lr.SS$</p> <p>$(ssk_i, (seed_{i,j}, seed_{j,i})_{j \in [n]}) \leftarrow sk_i$</p> <p>$mask_h \leftarrow \sum_{j \in hon} (PRF(seed_{i,j}, (pk, lr)) - PRF(seed_{j,i}, (pk, lr)))$</p> <p>$mask_h \leftarrow R_q^{m-k}$</p> <p>$v \leftarrow r_{0,[m-k]} + (\sum_{j \in [\ell]} b_j \cdot r_{j,[m-k]}) + 2c \cdot \lambda_{lr.SS,i} \cdot ssk_i + mask_h$</p> <p>$v \leftarrow R_q^{m-k}$</p> <p>$curSum(lr) \leftarrow curSum(lr) + v$</p> <p>Else // i is the last queried honest signer</p> <p>$v \leftarrow curSumR(lr) + 2c \cdot \hat{sk}_{[m-k]} - curSum(lr) - \sum_{j \in cor} 2c \cdot \lambda_{lr.SS,j} \cdot ssk_j$</p> <p>$z \leftarrow v + mask_c$</p> <p>Return (R, z)</p>

Figure 13.8: The PSIGNO oracle of the games G_1 , G_2 , and G_3 , where G_1 only contains dashed boxes, G_2 only contains highlighted boxes, and G_3 only contains solid boxes. In addition, each entry of the tables $curSum$ and $curSumR$ is initialized to 0. The rest of each game is identical to G_0 .

$\text{curSumR}(lr) + 2c \cdot \hat{\mathbf{s}}k_{[m-k]} - \text{curSum}(lr) - \sum_{j \in \text{cor}} 2c \cdot \lambda_{lr.SS,j} \cdot \text{ssk}_j + \text{mask}_c$. Here, $\text{curSumR}(lr) = \sum_{i' \in \text{hon}} \mathbf{r}_{0,[m-k]}^{(i')} + (\sum_{j \in [\ell]} b_j \cdot \mathbf{r}_{j,[m-k]}^{(i')})$, where $\{\mathbf{r}_j^{(i')}\}_{j \in [0..\ell]}$ denotes the nonces for signer $i' \in \text{hon}$, and $\text{curSum}(lr) = \sum_{i' \in \text{hon} \setminus \{i\}} \mathbf{v}^{(i')} = \sum_{i' \in \text{hon} \setminus \{i\}} (\mathbf{r}_{0,[m-k]}^{(i')} + (\sum_{j \in [\ell]} b_j \cdot \mathbf{r}_{j,[m-k]}^{(i')})) + 2c \cdot \lambda_{lr.SS,i'} \cdot \text{ssk}_{i'} + \text{mask}_h^{(i')}$, where $(\cdot)^{(i')}$ denotes the value computed during the query (i', lr) . Since $\hat{\mathbf{s}}k = \sum_{i' \in lr.SS} \lambda_{lr.SS,i'} \cdot \text{ssk}_{i'}$ and $\sum_{i' \in \text{hon}} \text{mask}_h^{(i')} = \sum_{i' \in \text{hon}} \sum_{j \in \text{hon}} (\text{PRF}(\text{seed}_{i',j}, (\text{pk}, lr)) - \text{PRF}(\text{seed}_{j,i'}, (\text{pk}, lr))) = 0$, we have $\mathbf{z} = \mathbf{r}_{0,[m-k]}^{(i)} + (\sum_{j \in [\ell]} b_j \cdot \mathbf{r}_{j,[m-k]}^{(i)}) + 2c \cdot \lambda_{lr.SS,i} \cdot \text{ssk}_i + \text{mask}_h^{(i)} + \text{mask}_c$, which is identical to \mathbf{G}_0 .

We now argue that BNonce occurs with a negligible probability. Since $\text{st}_i.\text{mapPP}(pp_i)$ is set to \perp after the query (i, lr) , BNonce occurs only if the PPO oracle generates a new token $(\mathbf{R}_0, \dots, \mathbf{R}_\ell)$ that is exactly the same as pp_i . Therefore, by the following lemma from [54], the probability that this occurs is at most $q_s^2 \cdot 2^{-N+1}$, and thus, since $N \geq 2\kappa$,

$$\text{Adv}^{\mathbf{G}_1}(\mathcal{A}, \kappa) \geq \text{Adv}^{\mathbf{G}_0}(\mathcal{A}, \kappa) - q_s^2 \cdot 2^{-2\kappa+1}. \quad (13.3)$$

Lemma 13.2.5 (Lemma 3.8, [54]). *For any integers $q, m, k > 0$, any real number $\sigma > 0$ and any matrix $A \in R_q^{k \times m}$, denote a distribution $\mathcal{D}(A) := \{[A|\mathbb{I}_k] \cdot \mathbf{s} \mid \mathbf{s} \leftarrow_{\$} \mathcal{D}_\sigma^{m+k}\}$. If $\sigma > 2N \cdot q^{\frac{1}{k+m} + \frac{1}{Nm}}$,*

$$\Pr_{A \leftarrow_{\$} R_q^{k \times m}} [H_\infty(\mathcal{D}(A)) \geq N - 1] \geq 1 - 2^{-N+1},$$

where $H_\infty(\mathcal{D}(A)) := -\log_2(\max_{\mathbf{x}' \in R_q^k} \Pr_{\mathbf{x} \leftarrow_{\$} \mathcal{D}(A)}[\mathbf{x} = \mathbf{x}'])$ denotes the min-entropy of $\mathcal{D}(A)$.⁴

G₂: The same as **G₁** except in the oracle PSIGNO, the aggregated mask mask_h computed from honest parties' seeds are replaced with a uniformly random value. The game is formally defined in Figure 13.8. We can show this game is computationally close to **G₁** by first replacing each $\text{PRF}(\text{seed}_{i,j}, \cdot)$ for honest parties i and j with a truly random function, which incurs at most a reduction loss of $n^2 \cdot \text{Adv}_{\text{PRF}}^{\text{prf}}(\mathcal{C}, \kappa)$. Then, the game is identical to **G₁**, since for each lr , denoting $\text{hon} = lr.SS \cap HS$ and $\text{mask}_h^{(i)} = \sum_{j \in \text{hon}} (\text{PRF}(\text{seed}_{i,j}, (\text{pk}, lr)) - \text{PRF}(\text{seed}_{j,i}, (\text{pk}, lr)))$, we have $\{\text{mask}_h^{(i)}\}_{i \in \text{hon} \setminus \{i'\}}$ is uniformly distributed over $R_q^{k(|\text{hon}|-1)}$ for any $i' \in \text{hon}$. Therefore,

$$\text{Adv}^{\mathbf{G}_2}(\mathcal{A}, \kappa) \geq \text{Adv}^{\mathbf{G}_1}(\mathcal{A}, \kappa) - n^2 \cdot \text{Adv}_{\text{PRF}}^{\text{prf}}(\mathcal{C}, \kappa). \quad (13.4)$$

⁴We omit the bit dropping from the original lemma, as it is not needed here and it only reduces the min-entropy.

G₃: The same as **G₂** except the value v is uniformly sampled from R_q^{m-k} if i is not the last queried honest party in $lr.ss$. The game is formally defined in Figure 13.8. Since mask_h is sampled uniformly from R_q^{m-k} and only used to mask v , the distribution of v is identical in both games. Therefore, we have

$$\text{Adv}^{\mathbf{G}_3}(\mathcal{A}, \kappa) = \text{Adv}^{\mathbf{G}_2}(\mathcal{A}, \kappa). \quad (13.5)$$

We construct an Ideal-TUF adversary \mathcal{B} as follows. To start with, after receiving A from the Ideal-TUF game, \mathcal{B} initializes $\text{st}_i.\text{mapPP}$ for $i \in [n]$, tables curSSL and curSum following the game **G₃**, then initializes a map curLR to an empty map, recording whether a lr request has made, and runs \mathcal{A} with input A and access to oracles $\widetilde{\text{INIT}}$, $\widetilde{\text{PPO}}$, $\widetilde{\text{PSIGNO}}$ and $\widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$: \mathcal{B} queries $\text{pk} \leftarrow \text{INIT}(CS)$. For each $i \in CS$, \mathcal{B} samples $\text{ssk}_i \leftarrow_{\$} R_q^{m-k}$ and $\text{seed}_{i,j} \leftarrow_{\$} \{0, 1\}^\kappa$ for each $j \in [n]$. Finally, \mathcal{B} returns $(\text{ssk}_i, (\text{seed}_{i,j}, \text{seed}_{j,i})_{j \in [n]})_{i \in CS}$.

$\widetilde{\text{PPO}}, \widetilde{\text{RO}}$: \mathcal{B} forwards queries directly to PPO and RO respectively.

$\widetilde{\text{PSIGNO}}(i, lr)$: The same as **PSIGNO**(i, lr) in **G₃** except that when $\text{curSSL}(lr) = \text{hon}$, \mathcal{B} queries $\hat{z} \leftarrow \text{SIGNO}(lr)$ and sets $z \leftarrow \hat{z} - \text{curSum}(lr) - \sum_{j \in \text{cor}} 2c \cdot \lambda_{lr.ss,j} \cdot \text{ssk}_j + \text{mask}_c$. Also, \mathcal{B} does not need to retrieve $\{r_{i,j}\}_{j \in [0..\ell]}$ and update curSumR , as the table curSumR is not used anymore.

After \mathcal{A} returns, \mathcal{B} outputs the output of \mathcal{A} .

We observe that \mathcal{B} wins the Ideal-TUF game if \mathcal{A} wins the game **G₃**, since the message $lr.\text{msg}$ is added to S in **G₃** if and only if $\text{curSSL}(lr) = \text{hon}$, which is exactly the scenario where \mathcal{B} makes a **SIGNO** query on lr . Also, since \mathcal{B} simulates the game **G₃** perfectly, it follows that $\text{Adv}^{\text{ideal-tuf}}(\mathcal{B}, \kappa) \geq \text{Adv}^{\mathbf{G}_3}(\mathcal{A}, \kappa)$. Therefore, we can conclude the lemma by Equations (13.3) to (13.5). \square \square

13.2.4 Proof of Lemma 13.2.4

Let \mathcal{A} be a Ideal-TUF adversary as described in the lemma. W.l.o.g. we assume that \mathcal{A} is deterministic. Also, we assume if \mathcal{A} returns $(\mu^*, (\mathbf{R}^*, z^*))$, the RO query $\text{H}_2(\text{pk}, \mu^*, \mathbf{R}^*)$ was made by \mathcal{A} , which adds at most one RO query. Also, since the game makes at most one RO query to H_1 and H_2 respectively for each signing query, the total number of RO queries to each of H_1 and

H_2 is bounded $q = q_h + q_s + 1$. We now construct an algorithm \mathcal{C} compatible with the syntax in Lemma 9.2.3 and construct \mathcal{B} from Fork ^{\mathcal{C}} .

CONSTRUCTION OF \mathcal{C} . The input of \mathcal{C} consists of $A, (\mathbf{t}_i)_{i \in [1+q_s(\ell+1)]}$, and a list of hash values h_1, \dots, h_{2q} , where $(A, (\mathbf{t}_i)_{i \in [1+q_s(\ell+1)]})$ are sampled following the AOM-MISIS game, and for each $i \in [q]$, h_{2i-1} is sampled uniformly from \mathcal{S}_b and h_{2i} is sampled uniformly from \mathcal{S}_c . To start with, \mathcal{C} initializes $\text{st}_i.\text{mapPP} \leftarrow ()$ for $i \in [n]$, and in addition, initializes a counter $\text{ctr}_h \leftarrow 0$ for counting the number of random oracle queries. Then, \mathcal{C} runs \mathcal{A} on input A with access to oracles $\widetilde{\text{INIT}}, \widetilde{\text{PPO}}, \widetilde{\text{PSIGNO}}$ and $\widetilde{\text{RO}}$, which are simulated as follows.

$\widetilde{\text{INIT}}(CS)$: The same as $\text{INIT}(CS)$ except \mathcal{C} sets $\text{pk} \leftarrow [2\mathbf{t}_0]_{\nu_{\text{pk}}}$.

$\widetilde{\text{PPO}}(i)$: For the j -th query, \mathcal{C} sets $\mathbf{R}_{\hat{j}} \leftarrow \mathbf{t}_{1+(j-1)(\ell+1)+\hat{j}+1}$ for $\hat{j} \in [0..\ell]$ and sets

$$\text{st}_i.\text{mapPP}((\mathbf{R}_{\hat{j}})_{\hat{j} \in [0..\ell]}) \leftarrow j .$$

Note that since \mathcal{C} does not sample $(\mathbf{r}_{\hat{j}})_{\hat{j} \in [0..\ell]}$, \mathcal{C} uses $\text{st}_i.\text{mapPP}$ to store the index j instead.

$\widetilde{\text{SIGNO}}(lr)$: The same as $\text{SIGNO}(lr)$ except that \mathcal{C} computes $\mathbf{z} \leftarrow \text{PI}(\mathbf{d})$, where

$$d_{\hat{j}} = \begin{cases} 2c, & \hat{j} = 1, \\ 1, & \hat{j} = 1 + (j-1)(\ell+1) + 1, j \in \text{honR} \\ b_{j'}, & \hat{j} = 1 + (j-1)(\ell+1) + 1 + j', j' \in [\ell], j \in \text{honR}, \\ 0, & \text{o.w.}, \end{cases} \quad (13.6)$$

and $\text{honR} := \{\text{st}_i.\text{mapPP}(lr.\text{PP}(i))\}_{i \in \text{hon}}$.

$\widetilde{\text{RO}}$ query $H_1(x)$: If $H_1(x) \neq \perp$, \mathcal{C} returns $H_1(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, lr)$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_1(x) \leftarrow_s \mathcal{S}_b^\ell$ and returns $H_1(x)$. Otherwise, \mathcal{C} increases ctr_h by 1, sets $H_1(x) \leftarrow h_{2\text{ctr}_h-1}$. Also, \mathcal{C} computes $\mathbf{R} \leftarrow \sum_{i \in lr.\text{SS}} (\mathbf{R}_{i,0} + \sum_{j \in [\ell]} b_j \cdot \mathbf{R}_{i,j})$, where $(\mathbf{R}_{i,j})_{j \in [0..\ell]} \leftarrow lr.\text{PP}(i)$ and $\{b_j\}_{j \in [\ell]} \leftarrow h_{2\text{ctr}_h-1}$. If $H_2(\text{pk}, lr.\text{msg}, \mathbf{R}) = \perp$, \mathcal{C} sets $H_2(\text{pk}, lr.\text{msg}, \mathbf{R}) \leftarrow h_{2\text{ctr}_h}$. Finally, \mathcal{C} returns $H_1(x)$.

$\widetilde{\text{RO}}$ query $H_2(x)$: If $H_2(x) \neq \perp$, \mathcal{C} returns $H_2(x)$. Otherwise, parse x as $(\tilde{\text{pk}}, \mu, \mathbf{R})$. If the parsing fails or $\tilde{\text{pk}} \neq \text{pk}$, \mathcal{C} sets $H_2(x) \leftarrow_s \mathcal{S}_c$. Otherwise, \mathcal{C} increases ctr_h by 1 and sets $H_2(x) \leftarrow h_{2\text{ctr}_h}$. Finally, \mathcal{C} returns $H_2(x)$.

After receiving the output $(\mu^*, (c^*, \mathbf{z}^*, \mathbf{h}^*))$ from \mathcal{A} , \mathcal{C} aborts if \mathcal{A} does not win the Ideal-TUF game. Otherwise \mathcal{C} computes $\mathbf{R}^* \leftarrow \left[\widehat{A} \mathbf{z}^* - 2^{\nu_{\text{pk}}} \cdot c^* \cdot \text{pk} \right]_{\nu_r} + \mathbf{h}^*$ and finds the index I such that $H_2(\text{pk}, \mu^*, \mathbf{R}^*)$ is set to h_I during the simulation. By our assumption of \mathcal{A} , we know such I must exist. Then, \mathcal{C} returns $(I, \text{Out} = (\mu^*, \mathbf{R}^*, c^*, \mathbf{z}^*, \mathbf{h}^*))$.

ANALYSIS OF \mathcal{C} . To use Lemma 9.2.3, we define $S := \{2j\}_{j \in [q]}$ and IG as the algorithm that samples $(A, (\mathbf{t}_i)_{i \in [1+q_s(\ell+1)]})$ following the AOM-MISIS game, and HG as the algorithm that samples h_{2i-1} uniformly from \mathcal{S}_b and samples h_{2i} uniformly from \mathcal{S}_c for each $i \in [q]$. From the simulation, we know that the output index I of \mathcal{C} is always in S . Also, it is not hard to check that \mathcal{C} simulates the game Ideal-TUF perfectly, which implies $\text{acc}(\mathcal{C}) \geq \text{Adv}^{\text{ideal-tuf}}(\mathcal{A}, \kappa)$. By Lemma 9.2.3, we have that

$$\text{acc}(\text{Fork}^{\mathcal{C}}) \geq \text{Adv}^{\text{ideal-tuf}}(\mathcal{A}, \kappa)^2/q.$$

CONSTRUCT \mathcal{B} FROM $\text{Fork}^{\mathcal{C}}$. We now construct the AOM-MISIS adversary \mathcal{B} using $\text{Fork}^{\mathcal{C}}$. To start with, \mathcal{B} receives $(A, \{\mathbf{t}_i\}_{i \in [Q]})$ from the AOM-MISIS game with $Q = K + 1 + q_s(\ell + 1)$ and runs $\text{Fork}^{\mathcal{C}}(A, \{\mathbf{t}_i\}_{i \in [Q]})$ with access to the PI oracle from the AOM-MISIS game. If $\text{Fork}^{\mathcal{C}}$ outputs $(I, \text{Out} = (\mu^*, \mathbf{R}^*, c^*, \mathbf{z}^*, \mathbf{h}^*), \widetilde{\text{Out}} = (\widetilde{\mu}^*, \widetilde{\mathbf{R}}^*, \widetilde{c}^*, \widetilde{\mathbf{z}}^*, \widetilde{\mathbf{h}}^*))$ and $c^* \neq \widetilde{c}^*$, \mathcal{B} sets $\hat{\mathbf{s}} \leftarrow (\mathbf{z}^* - \widetilde{\mathbf{z}}^*, 2(c^* - \widetilde{c}^*)\mathbf{t}_0 - \widehat{A}(\mathbf{z}^* - \widetilde{\mathbf{z}}^*))$ and $\hat{\mathbf{b}} \leftarrow (2(c^* - \widetilde{c}^*), 0, \dots, 0)$. Otherwise, \mathcal{B} aborts. It is clear that $A\hat{\mathbf{s}} = \widehat{A}(\mathbf{z}^* - \widetilde{\mathbf{z}}^*) + 2(c^* - \widetilde{c}^*)\mathbf{t}_0 - \widetilde{A}(\mathbf{z}^* - \widetilde{\mathbf{z}}^*) = 2(c^* - \widetilde{c}^*)\mathbf{t}_0 = \sum_{i \in [Q]} \hat{b}_i \mathbf{t}_i$.

We now show how \mathcal{B} sets \mathbf{u} such that $\hat{\mathbf{b}}^T \mathbf{u} \neq 0$ and $\mathbf{d}^T \mathbf{u} = 0 \pmod q$ for any oracle query $\text{PI}(\mathbf{d})$. Note that \mathcal{B} only makes PI queries while simulating oracle $\widetilde{\text{SIGNO}}$. We set $u_1 = 1$ and thus $\hat{\mathbf{b}}^T \mathbf{u} = 2(c^* - \widetilde{c}^*) \neq 0$.

Enumerating j from 1 to q_s , \mathcal{C} sets $u_{1+(j-1)(\ell+1)+[\ell+1]}$ such that

$$\sum_{i \in [1+j(\ell+1)]} u_i d_i = 0 \text{ for each PI query } \mathbf{d} \text{ with } d_{1+[j(\ell+1)]} \neq \mathbf{0}. \quad (13.7)$$

To simplify notation in the following analysis, we use \mathbf{v} to denote the vector $u_{1+(j-1)(\ell+1)+[\ell+1]} \in R^{\ell+1}$. Concretely, \mathcal{C} sets \mathbf{v} as follows. Suppose Equation (13.7) holds for $j-1$ (except when $j=1$, i.e., no condition is required for the case $j=1$). We say a $\widetilde{\text{SIGNO}}$ query lr corresponds to the j -th token if and only if it is the valid query with $\text{st}_i.\text{mapPP}(lr.\text{PP}(i)) = j$, where a valid query means the $\widetilde{\text{SIGNO}}$ oracle does not return \perp . From the simulation, there is at most one $\widetilde{\text{SIGNO}}$ query corresponding to the j -th token during each execution of \mathcal{A} . Therefore, there are the following cases:

Case 1: No query corresponds to the j -th token during both executions. In this case, \mathcal{B} set $\mathbf{v} \leftarrow \mathbf{0}$.

Case 2: Only one query corresponds to the j -th token during the two executions. Denote \mathbf{d} as the PI query made during the execution of the $\widetilde{\text{SIGNO}}$ query corresponding to the j -th token, where \mathbf{d} follows the form given in Equation (13.6). \mathcal{B} sets $\mathbf{v} \leftarrow (-\sum_{i \in [1+(j-1)(\ell+1)]} u_i d_i, 0, \dots, 0)$, and it follows $\sum_{i \in [1+j(\ell+1)]} u_i d_i = \sum_{i \in [1+(j-1)(\ell+1)]} u_i d_i + v_1 = 0$. Also, since \mathbf{d} is the PI query with $d_{1+(j-1)(\ell+1)+[\ell+1]} \neq \mathbf{0}$, it follows that Equation (13.7) holds.

Case 3: There is one query corresponding to the j -th token during each of the two executions. Denote \mathbf{d} (resp. $\tilde{\mathbf{d}}$) as the PI query made during the execution of the $\widetilde{\text{SIGNO}}$ query corresponding to the j -th token before (resp. after) rewinding. If $\mathbf{d} = \tilde{\mathbf{d}}$, then \mathcal{C} sets \mathbf{v} in the same way as Case 2.

Otherwise, let $\hat{k} \in [\ell]$ be the index such that $d_{1+(j-1)(\ell+1)+1+\hat{k}} \neq \tilde{d}_{1+(j-1)(\ell+1)+1+\hat{k}}$. (If such \hat{k} does not exist, \mathcal{B} aborts.) Denote $b := d_{1+(j-1)(\ell+1)+1+\hat{k}}$ and $\tilde{b} := \tilde{d}_{1+(j-1)(\ell+1)+1+\hat{k}}$. Since Equation (13.7) holds for $j-1$, we have either $\sum_{i \in [1+(j-1)(\ell+1)]} u_i d_i = 0$ or $\sum_{i \in [1+(j-1)(\ell+1)]} u_i d_i = d_1 \in 2\mathcal{S}_c$ by Equation (13.6), where $2\mathcal{S}_c := \{2c \mid c \in \mathcal{S}_c\}$. Therefore, denote $\Delta := \frac{1}{2} \sum_{i \in [1+(j-1)(\ell+1)]} u_i d_i$ and $\tilde{\Delta} := \frac{1}{2} \sum_{i \in [1+(j-1)(\ell+1)]} u_i \tilde{d}_i$. Since $b, \tilde{b} \in \mathcal{S}_b$, by Lemma 11.2.1, there exists $\gamma \in R$ such that $\gamma(b - \tilde{b}) = 2 \pmod q$. \mathcal{C} sets

$$v_{\hat{j}} \leftarrow \begin{cases} -2\Delta + b\gamma(\Delta - \tilde{\Delta}), & \hat{j} = 1, \\ -\gamma(\Delta - \tilde{\Delta}), & \hat{j} = \hat{k}, \\ 0, & \text{o.w.} \end{cases}$$

Then, it holds that $\sum_{i \in [1+j(\ell+1)]} u_i d_i = \sum_{i \in [1+(j-1)(\ell+1)]} u_i d_i + v_1 + b v_{\hat{k}} = 2\Delta - 2\Delta + b\gamma(\Delta - \tilde{\Delta}) - b\gamma(\Delta - \tilde{\Delta}) = 0$ and $\sum_{i \in [1+j(\ell+1)]} u_i \tilde{d}_i = \sum_{i \in [1+(j-1)(\ell+1)]} u_i \tilde{d}_i + v_1 + \tilde{b} v_{\hat{k}} = 2\tilde{\Delta} - 2\Delta + b\gamma(\Delta - \tilde{\Delta}) - \tilde{b}\gamma(\Delta - \tilde{\Delta}) = 2\tilde{\Delta} - 2\Delta + (b - \tilde{b})\gamma(\Delta - \tilde{\Delta}) = 2\tilde{\Delta} - 2\Delta + 2(\Delta - \tilde{\Delta}) = 0$. Finally, since \mathbf{d} and $\tilde{\mathbf{d}}$ are the PI queries with $d_{1+(j-1)(\ell+1)+[\ell+1]} \neq \mathbf{0}$, it follows that Equation (13.7) holds.

ANALYSIS OF \mathcal{B} . Denote BadHash as the event that $h_1, \tilde{h}_1, \dots, h_{2q}, \tilde{h}_{2q}$ are *not* all distinct. We now show that \mathcal{B} wins the AOM-MISIS game if Fork^C returns and BadHash does not occur.

We first show that if Fork^C returns and BadHash does not occur, \mathcal{B} does not abort. (The following argument is similar to the one provided in the security reduction of the CTZ protocol.) Suppose

Fork^C returns and BadHash does not occur. Then, the only step where \mathcal{B} might abort is in Case 3 above. In Case 3, suppose $\mathbf{d} \neq \tilde{\mathbf{d}}$ and $d_{1+(j-1)(\ell+1)+1+[\ell]} = \tilde{d}_{1+(j-1)(\ell+1)+1+[\ell]}$ (in which case \mathcal{B} aborts). Let lr be the $\widetilde{\text{SIGNO}}$ query that corresponds to the j -th token *before* rewinding. Then, there exists $J \in [q]$ such that $h_{2J-1} = H_1(\text{pk}, lr) = d_{1+(j-1)(\ell+1)+1+[\ell]}$. Since BadHash does not occur, the only possibility is that the $2J - 1 < I$ and lr is also the $\widetilde{\text{SIGNO}}$ query that corresponds to the j -th token *after* rewinding. Let \mathbf{R} be the aggregated nonce computed from $\text{CompPar}(\text{pk}, lr)$. Denote $\hat{J} \in [q]$ be the index such that $h_{2\hat{J}} = H_2(\text{pk}, lr.\text{msg}, \mathbf{R})$ before rewinding. From the simulation of the random oracles, it holds that $\hat{J} \leq J$ and thus $2\hat{J} \leq 2J \leq I$. Also, since $lr.\text{msg} \neq \mu^*$ (o.w., \mathcal{C} would not win the game), we have $2\hat{J} \neq I$ and thus $2\hat{J} < I$. Therefore, $H_2(\text{pk}, lr.\text{msg}, \mathbf{R})$ in the second execution of \mathcal{C} is also $h_{2\hat{J}}$. This implies $\mathbf{d} = \tilde{\mathbf{d}}$, which contradicts our assumption.

Suppose \mathcal{B} does not abort. From the way \mathcal{B} sets \mathbf{u} , Equation (13.7) holds for $j = q_s$. Then, since for each PI query \mathbf{d} , the component $d_{1+[q_s(\ell+1)]} \neq \mathbf{0}$, it follows that $\mathbf{d}^T \mathbf{u} = \sum_{i \in [1+q_s(\ell+1)]} u_i d_i = 0$. It is left to bound the norms of vectors $\hat{\mathbf{s}}$, $\hat{\mathbf{b}}$ and \mathbf{u} .

We first bound the norm of $\hat{\mathbf{s}}$. We have

$$\|\hat{\mathbf{s}}[m - k]\| = \|\mathbf{z}^* - \tilde{\mathbf{z}}^*\| \leq 2\beta_z \quad (13.8)$$

and, by Lemma 13.2.6 (proved in [60]),

$$\begin{aligned} \|\hat{\mathbf{s}}[(m - k + 1)..m]\| &= \left\| 2(c^* - \tilde{c}^*)\mathbf{t}_0 - \hat{A}(\mathbf{z}^* - \tilde{\mathbf{z}}^*) \right\| \\ &\leq \left\| (c^* - \tilde{c}^*)2^{\nu_{\text{pk}}}\overline{\text{pk}} - \hat{A}(\mathbf{z}^* - \tilde{\mathbf{z}}^*) \right\| + \left\| (c^* - \tilde{c}^*) \cdot (2\mathbf{t}_0 - 2^{\nu_{\text{pk}}}\overline{\text{pk}}) \right\| \\ &\leq \left\| (c^* - \tilde{c}^*)2^{\nu_{\text{pk}}}\text{pk} - \hat{A}(\mathbf{z}^* - \tilde{\mathbf{z}}^*) \right\| + \beta_c 2^{\nu_{\text{pk}}+1}\sqrt{Nk} \quad (\text{by Lemma 13.2.6}) \\ &\leq \left\| 2^{\nu_r} \left[\overline{\hat{A}\tilde{\mathbf{z}}^* - 2^{\nu_{\text{pk}}}\tilde{c}^*\text{pk}} \right]_{\nu_r} - 2^{\nu_r} \left[\overline{\hat{A}\mathbf{z}^* - 2^{\nu_{\text{pk}}}c^*\text{pk}} \right]_{\nu_r} \right\| \pmod{q} \\ &\quad + \left\| (\hat{A}\mathbf{z}^* - 2^{\nu_{\text{pk}}}c^*\text{pk}) - 2^{\nu_r} \left[\overline{\hat{A}\mathbf{z}^* - 2^{\nu_{\text{pk}}}c^*\text{pk}} \right]_{\nu_r} \right\| \\ &\quad + \left\| (\hat{A}\tilde{\mathbf{z}}^* - 2^{\nu_{\text{pk}}}\tilde{c}^*\text{pk}) - 2^{\nu_r} \left[\overline{\hat{A}\tilde{\mathbf{z}}^* - 2^{\nu_{\text{pk}}}\tilde{c}^*\text{pk}} \right]_{\nu_r} \right\| + \beta_c 2^{\nu_{\text{pk}}+1}\sqrt{Nk} \\ &\leq \left\| 2^{\nu_r} \left(\left[\overline{\hat{A}\tilde{\mathbf{z}}^* - 2^{\nu_{\text{pk}}}\tilde{c}^*\text{pk}} \right]_{\nu_r} - \left[\overline{\hat{A}\mathbf{z}^* - 2^{\nu_{\text{pk}}}c^*\text{pk}} \right]_{\nu_r} \right) \right\| \pmod{q} \\ &\quad + 2^{\nu_r+1}\sqrt{Nk} + \beta_c 2^{\nu_{\text{pk}}+1}\sqrt{Nk}. \quad (\text{by Lemma 13.2.6}) \end{aligned} \quad (13.9)$$

Lemma 13.2.6 (Lemma 3.14 [60]). *For any integers $v \geq 4$ and $q > 2^v$, let $q_v = \lfloor q/2^v \rfloor$. Moreover,*

assume q and v satisfies $q_v = \lfloor q/2^v \rfloor$. Then, for any $x \in \mathbb{Z}_q$, we have

$$\left| x - 2^v \cdot \overline{[x]_v} \right| \leq 2^v - 1.$$

Since $\left[\widehat{A}z^* - 2^{\nu_{pk}} c^* pk \right]_{\nu_r} + \mathbf{h}^* = \mathbf{R}^* = \widetilde{\mathbf{R}}^* = \left[\widehat{A}\widetilde{z}^* - 2^{\nu_{pk}} \widetilde{c}^* pk \right]_{\nu_r} + \widetilde{\mathbf{h}}^*$, there exists $\boldsymbol{\delta} \in R^k$ such that $\|\boldsymbol{\delta}\|_\infty \leq 2$ and $\left[\widehat{A}z^* - 2^{\nu_{pk}} c^* pk \right]_{\nu_r} + \overline{\mathbf{h}^*} = \left[\widehat{A}\widetilde{z}^* - 2^{\nu_{pk}} \widetilde{c}^* pk \right]_{\nu_r} + \overline{\widetilde{\mathbf{h}}^*} + q_{\nu_r} \cdot \boldsymbol{\delta}$. Therefore,

$$\begin{aligned} & \left\| 2^{\nu_r} \left(\left[\widehat{A}\widetilde{z}^* - 2^{\nu_{pk}} \widetilde{c}^* pk \right]_{\nu_r} - \left[\widehat{A}z^* - 2^{\nu_{pk}} c^* pk \right]_{\nu_r} \right) \bmod q \right\| \\ & \leq \left\| 2^{\nu_r} \overline{\mathbf{h}^*} \bmod q \right\| + \left\| 2^{\nu_r} \overline{\widetilde{\mathbf{h}}^*} \bmod q \right\| + \left\| 2^{\nu_r} q_{\nu_r} \cdot \boldsymbol{\delta} \bmod q \right\| \\ & \leq 2\beta_z + 2^{\nu_r} \|\boldsymbol{\delta}\| \leq 2\beta_z + 2^{\nu_r+1} \sqrt{Nk}. \end{aligned} \quad (13.10)$$

Therefore, by Equations (13.8) to (13.10), $\|\hat{\mathbf{s}}\| \leq (2^{\nu_r+2} + \beta_c \cdot 2^{\nu_{pk}+1}) \cdot \sqrt{Nk} + 4\beta_z \leq \beta_s$.

Also, $\left\| (\hat{b}_1 \cdot \sigma_1, \dots, \hat{b}_Q \cdot \sigma_Q) \right\|_1 = \|2(c^* - \widetilde{c}^*)\sigma_{sk}\|_1 \leq 4\sigma_{sk}\beta_c \leq \beta_b$.

It is left to bound $\|(u_1/\sigma_1, \dots, u_Q/\sigma_Q)\|$. For $j \in [q_s]$, denote $\mathbf{v} := u_{K+1+(j-1)(\ell+1)+[\ell+1]}$. There are three cases as mentioned in the construction of \mathcal{B} . For the first case, $\|\mathbf{v}\| = 0$. For the second case, since $v_1 \in 2\mathcal{S}_c \cup \{0\}$, $\|\mathbf{v}\| = \|v\| \leq 2\sqrt{\beta_c}$. For the third case, since $\Delta, \widetilde{\Delta} \in \mathcal{S}_c \cup \{0\}$, $\|\Delta\|_1 \leq \beta_c$ and $\|\widetilde{\Delta}\|_1 \leq \beta_c$. Since $\|\gamma\| = \sqrt{N}$ (by Lemma 11.2.1),

$$\begin{aligned} \|\mathbf{v}\|^2 &= \left\| \gamma(\widetilde{b}\Delta - b\widetilde{\Delta}), 0, \dots, 0, -\gamma(\Delta - \widetilde{\Delta}), 0, \dots, 0 \right\|^2 \\ &\leq 2(\|\Delta\|_1 + \|\widetilde{\Delta}\|_1)^2 \|\gamma\|^2 \leq 8N\beta_c^2. \end{aligned}$$

Therefore, $\|(u_1/\sigma_1, \dots, u_Q/\sigma_Q)\| \leq u_1/\sigma_{sk} + \sqrt{\sum_{j \in [q_s]} \|\mathbf{v}/\sigma_r\|} \leq 1/\sigma_{sk} + \beta_c \sqrt{8Nq_s}/\sigma_r \leq \beta_u$. The above shows that \mathcal{B} wins the AOM-MISIS game, given that Fork^C returns and BadHash does not occur.

Finally, since each of $h_1, h_3, \dots, h_{2q-1}$ and h_2, h_4, \dots, h_{2q} are sampled uniformly from \mathcal{S}_b^ℓ and \mathcal{S}_c respectively, $\Pr[\text{BadHash}] \leq (2q)^2/|\mathcal{S}_b|^\ell + (2q)^2/|\mathcal{S}_c| \leq 8q^2 2^{-2\kappa}$. Therefore,

$$\begin{aligned} \text{Adv}_{\text{par}}^{\text{aom-misis}}(\mathcal{B}, \kappa) &\geq \text{acc}(\text{Fork}^C) - \Pr[\text{BadHash}] \\ &\geq \text{Adv}^{\text{ideal-tuf}}(\mathcal{A}, \kappa)^2/q - 8q^2 2^{-2\kappa}, \end{aligned}$$

which concludes the lemma. \square

BIBLIOGRAPHY

- [1] PCM: Click fraud prevention and attribution sent to advertiser. <https://webkit.org/blog/11940/pcm-click-fraud-prevention-and-attribution-sent-to-advertiser/>. Accessed: 2021-09-30.
- [2] Trust tokens. <https://developer.chrome.com/docs/privacy-sandbox/trust-tokens/>. Accessed: 2022-01-11.
- [3] Masayuki Abe. A secure three-move blind signature scheme for polynomially many signatures. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 136–151. Springer, Berlin, Heidelberg, May 2001.
- [4] Masayuki Abe and Eiichiro Fujisaki. How to date blind signatures. In Kwangjo Kim and Tsutomu Matsumoto, editors, *ASIACRYPT'96*, volume 1163 of *LNCS*, pages 244–251. Springer, Berlin, Heidelberg, November 1996.
- [5] Masayuki Abe and Tatsuaki Okamoto. Provably secure partially blind signatures. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 271–286. Springer, Berlin, Heidelberg, August 2000.
- [6] Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete Gaussian left-over hash lemma over infinite domains. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 97–116. Springer, Berlin, Heidelberg, December 2013.
- [7] Shweta Agrawal, Elena Kirshanova, Damien Stehlé, and Anshu Yadav. Practical, round-optimal lattice-based blind signatures. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 39–53. ACM Press, November 2022.
- [8] Shweta Agrawal, Damien Stehlé, and Anshu Yadav. Round-optimal lattice-based threshold signatures, revisited. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *ICALP 2022*, volume 229 of *LIPICs*, pages 8:1–8:20. Schloss Dagstuhl, July 2022.
- [9] Renas Bacho, Julian Loss, Stefano Tessaro, Benedikt Wagner, and Chenzhi Zhu. Twinkle: Threshold signatures from DDH with full adaptive security. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part I*, volume 14651 of *LNCS*, pages 429–459. Springer, Cham, May 2024.

- [10] Matilda Backendal, Mihir Bellare, Jessica Sorrell, and Jiahao Sun. The fiat-shamir zoo: relating the security of different signature variants. In *Nordic Conference on Secure IT Systems*, pages 154–170. Springer, 2018.
- [11] Foteini Baldimtsi and Anna Lysyanskaya. Anonymous credentials light. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 1087–1098. ACM Press, November 2013.
- [12] Balthazar Bauer, Georg Fuchsbauer, and Antoine Plouviez. The one-more discrete logarithm assumption in the generic group model. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 587–617. Springer, Cham, December 2021.
- [13] Mira Belenkiy, Jan Camenisch, Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Hovav Shacham. Randomizable proofs and delegatable anonymous credentials. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 108–125. Springer, Berlin, Heidelberg, August 2009.
- [14] Mihir Bellare, Elizabeth C. Crites, Chelsea Komlo, Mary Maller, Stefano Tessaro, and Chenzhi Zhu. Better than advertised security for non-interactive threshold signatures. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 517–550. Springer, Cham, August 2022.
- [15] Mihir Bellare, Wei Dai, and Lucy Li. The local forking lemma and its application to deterministic encryption. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 607–636. Springer, Cham, December 2019.
- [16] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-RSA-inversion problems and the security of Chaum’s blind signature scheme. *Journal of Cryptology*, 16(3):185–215, June 2003.
- [17] Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 390–399. ACM Press, October / November 2006.
- [18] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [19] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Berlin, Heidelberg, May / June 2006.

- [20] Mihir Bellare, Stefano Tessaro, and Chenzhi Zhu. Stronger security for non-interactive threshold signatures: BLS and FROST. Cryptology ePrint Archive, Report 2022/833, 2022.
- [21] Josh Cohen Benaloh and Jerry Leichter. Generalized secret sharing and monotone functions. In Shafi Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 27–35. Springer, New York, August 1990.
- [22] Fabrice Benhamouda, Jan Camenisch, Stephan Krenn, Vadim Lyubashevsky, and Gregory Neven. Better zero-knowledge proofs for lattice encryption and their application to group signatures. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 551–572. Springer, Berlin, Heidelberg, December 2014.
- [23] Fabrice Benhamouda, Shai Halevi, Hugo Krawczyk, Yiping Ma, and Tal Rabin. SPRINT: High-throughput robust distributed Schnorr signatures. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part V*, volume 14655 of *LNCS*, pages 62–91. Springer, Cham, May 2024.
- [24] Fabrice Benhamouda, Tancrède Lepoint, Julian Loss, Michele Orrù, and Mariana Raykova. On the (in)security of ROS. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 33–53. Springer, Cham, October 2021.
- [25] Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. *Journal of Cryptographic Engineering*, 2(2):77–89, September 2012.
- [26] Olivier Blazy, Georg Fuchsbauer, David Pointcheval, and Damien Vergnaud. Short blind signatures. *J. Comput. Secur.*, 21(5):627–661, 2013.
- [27] Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Yvo Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 31–46. Springer, Berlin, Heidelberg, January 2003.
- [28] Dan Boneh, Rosario Gennaro, and Steven Goldfeder. Using level-1 homomorphic encryption to improve threshold DSA signatures for bitcoin wallet security. In Tanja Lange and Orr Dunkelman, editors, *LATINCRYPT 2017*, volume 11368 of *LNCS*, pages 352–377. Springer, Cham, September 2019.
- [29] Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 565–596. Springer, Cham, August 2018.

- [30] Dan Boneh, Rosario Gennaro, Steven Goldfeder, and Sam Kim. A lattice-based universal thresholdizer for cryptographic systems. Cryptology ePrint Archive, Report 2017/251, 2017.
- [31] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532. Springer, Berlin, Heidelberg, December 2001.
- [32] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. *Journal of Cryptology*, 17(4):297–319, September 2004.
- [33] Cecilia Boschini, Darya Kaviani, Russell W. F. Lai, Giulio Malavolta, Akira Takahashi, and Mehdi Tibouchi. Ringtail: Practical Two-Round Threshold Signatures from Learning with Errors . In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 149–164, Los Alamitos, CA, USA, May 2025. IEEE Computer Society.
- [34] Cecilia Boschini, Akira Takahashi, and Mehdi Tibouchi. MuSig-L: Lattice-based multi-signature with single-round online phase. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 276–305. Springer, Cham, August 2022.
- [35] Stefan Brands. Untraceable off-line cash in wallets with observers (extended abstract). In Douglas R. Stinson, editor, *CRYPTO’93*, volume 773 of *LNCS*, pages 302–318. Springer, Berlin, Heidelberg, August 1994.
- [36] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 93–118. Springer, Berlin, Heidelberg, May 2001.
- [37] Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 56–72. Springer, Berlin, Heidelberg, August 2004.
- [38] Ran Canetti, Rosario Gennaro, Steven Goldfeder, Nikolaos Makriyannis, and Udi Peled. UC non-interactive, proactive, threshold ECDSA with identifiable aborts. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1769–1787. ACM Press, November 2020.
- [39] Rutchathon Chairattana-Apirom, Lucjan Hanzlik, Julian Loss, Anna Lysyanskaya, and Benedikt Wagner. PI-cut-choo and friends: Compact blind signatures via parallel instance cut-and-choose and more. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part III*, volume 13509 of *LNCS*, pages 3–31. Springer, Cham, August 2022.

- [40] Rutchathon Chairattana-Apirom and Anna Lysyanskaya. Compact cut-and-choose: Boosting the security of blind signature schemes, compactly. *Cryptology ePrint Archive*, Report 2022/003, 2022.
- [41] Rutchathon Chairattana-Apirom, Stefano Tessaro, and Chenzhi Zhu. Partially non-interactive two-round lattice-based threshold signatures. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IV*, volume 15487 of *LNCS*, pages 268–302. Springer, Singapore, December 2024.
- [42] David Chaum. Verification by anonymous monitors. In Allen Gersho, editor, *CRYPTO'81*, volume ECE Report 82-04, pages 138–139. U.C. Santa Barbara, Dept. of Elec. and Computer Eng., 1981.
- [43] David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO'82*, pages 199–203. Plenum Press, New York, USA, 1982.
- [44] David Chaum. Blind signature system. In David Chaum, editor, *CRYPTO'83*, page 153. Plenum Press, New York, USA, 1983.
- [45] David Chaum. Elections with unconditionally-secret ballots and disruption equivalent to breaking RSA. In C. G. Günther, editor, *EUROCRYPT'88*, volume 330 of *LNCS*, pages 177–182. Springer, Berlin, Heidelberg, May 1988.
- [46] David Chaum, Amos Fiat, and Moni Naor. Untraceable electronic cash. In Shafi Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 319–327. Springer, New York, August 1990.
- [47] David Chaum and Torben P. Pedersen. Wallet databases with observers. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 89–105. Springer, Berlin, Heidelberg, August 1993.
- [48] Yanbo Chen. DualMS: Efficient lattice-based two-round multi-signature with trapdoor-free simulation. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 716–747. Springer, Cham, August 2023.
- [49] Deirdre Connolly, Chelsea Komlo, Ian Goldberg, and Christopher A. Wood. Two-Round Threshold Schnorr Signatures with FROST. Internet-Draft draft-irtf-cfrg-frost-10, Internet Engineering Task Force, September 2022. Work in Progress.
- [50] Elizabeth Crites, Chelsea Komlo, and Mary Maller. How to prove Schnorr assuming Schnorr: Security of multi- and threshold signatures. *Cryptology ePrint Archive*, Report 2021/1375, 2021.

- [51] Ivan Damgård, Claudio Orlandi, Akira Takahashi, and Mehdi Tibouchi. Two-round n-out-of-n and multi-signatures and trapdoor commitment from lattices. In Juan Garay, editor, *PKC 2021, Part I*, volume 12710 of *LNCS*, pages 99–130. Springer, Cham, May 2021.
- [52] Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. Privacy pass: Bypassing internet challenges anonymously. *PoPETs*, 2018(3):164–180, July 2018.
- [53] Alfredo De Santis, Yvo Desmedt, Yair Frankel, and Moti Yung. How to share a function securely. In *26th ACM STOC*, pages 522–533. ACM Press, May 1994.
- [54] Rafaël Del Pino, Shuichi Katsumata, Mary Maller, Fabrice Mouhartem, Thomas Prest, and Markku-Juhani O. Saarinen. Threshold raccoon: Practical threshold signatures from standard lattice assumptions. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 219–248. Springer, Cham, May 2024.
- [55] Rafaël del Pino, Shuichi Katsumata, Thomas Prest, and Mélissa Rossi. Raccoon: A masking-friendly signature proven in the probing model. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part I*, volume 14920 of *LNCS*, pages 409–444. Springer, Cham, August 2024.
- [56] Frank Denis, Frederic Jacobs, and Christopher A. Wood. RSA Blind Signatures. Internet-Draft draft-irtf-cfrg-rsa-blind-signatures-02, Internet Engineering Task Force, August 2021. Work in Progress.
- [57] Frank Denis, Frederic Jacobs, and Christopher A. Wood. RSA Blind Signatures. RFC 9474, October 2023.
- [58] Yvo Desmedt. Society and group oriented cryptography: A new concept. In Carl Pomerance, editor, *CRYPTO'87*, volume 293 of *LNCS*, pages 120–127. Springer, Berlin, Heidelberg, August 1988.
- [59] Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In Gilles Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 307–315. Springer, New York, August 1990.
- [60] Thomas Espitau, Shuichi Katsumata, and Kaoru Takemure. Two-round threshold signature from algebraic one-more learning with errors. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 387–424. Springer, Cham, August 2024.
- [61] Thomas Espitau, Guilhem Niot, and Thomas Prest. Flood and submerge: Distributed key generation and robust threshold signature from lattices. In Leonid Reyzin and Douglas

- Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 425–458. Springer, Cham, August 2024.
- [62] Georg Fuchsbauer, Christian Hanser, Chethan Kamath, and Daniel Slamanig. Practical round-optimal blind signatures in the standard model from weaker assumptions. In Vassilis Zikas and Roberto De Prisco, editors, *SCN 16*, volume 9841 of *LNCS*, pages 391–408. Springer, Cham, August / September 2016.
- [63] Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. Practical round-optimal blind signatures in the standard model. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 233–253. Springer, Berlin, Heidelberg, August 2015.
- [64] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Cham, August 2018.
- [65] Georg Fuchsbauer, Antoine Plouviez, and Yannick Seurin. Blind Schnorr signatures and signed ElGamal encryption in the algebraic group model. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 63–95. Springer, Cham, May 2020.
- [66] Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A practical secret voting scheme for large scale elections. In Jennifer Seberry and Yuliang Zheng, editors, *AUSCRYPT'92*, volume 718 of *LNCS*, pages 244–251. Springer, Berlin, Heidelberg, December 1993.
- [67] Sanjam Garg and Divya Gupta. Efficient round optimal blind signatures. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 477–495. Springer, Berlin, Heidelberg, May 2014.
- [68] Sanjam Garg, Vanishree Rao, Amit Sahai, Dominique Schröder, and Dominique Unruh. Round optimal blind signatures. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 630–648. Springer, Berlin, Heidelberg, August 2011.
- [69] Rosario Gennaro and Steven Goldfeder. Fast multiparty threshold ECDSA with fast trustless setup. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1179–1194. ACM Press, October 2018.
- [70] Rosario Gennaro, Steven Goldfeder, and Arvind Narayanan. Threshold-optimal DSA/ECDSA signatures and an application to bitcoin wallet security. In Mark Manulis, Ahmad-Reza Sadeghi, and Steve Schneider, editors, *ACNS 2016*, volume 9696 of *LNCS*, pages 156–174. Springer, Cham, June 2016.

- [71] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Robust threshold DSS signatures. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 354–371. Springer, Berlin, Heidelberg, May 1996.
- [72] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure applications of Pedersen’s distributed key generation protocol. In Marc Joye, editor, *CT-RSA 2003*, volume 2612 of *LNCS*, pages 373–390. Springer, Berlin, Heidelberg, April 2003.
- [73] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. *Journal of Cryptology*, 20(1):51–83, January 2007.
- [74] Rosario Gennaro, Tal Rabin, Stanislaw Jarecki, and Hugo Krawczyk. Robust and efficient sharing of RSA functions. *Journal of Cryptology*, 13(2):273–300, March 2000.
- [75] Rosario Gennaro, Tal Rabin, Stanislaw Jarecki, and Hugo Krawczyk. Erratum: Robust and efficient sharing of RSA functions. *Journal of Cryptology*, 20(3):393, July 2007.
- [76] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.
- [77] Essam Ghadafi. Efficient round-optimal blind signatures in the standard model. In Aggelos Kiayias, editor, *FC 2017*, volume 10322 of *LNCS*, pages 455–473. Springer, Cham, April 2017.
- [78] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.
- [79] Kamil Doruk Gür, Jonathan Katz, and Tjerand Silde. Two-round threshold lattice-based signatures from threshold homomorphic encryption. In Markku-Juhani Saarinen and Daniel Smith-Tone, editors, *Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Part II*, pages 266–300. Springer, Cham, June 2024.
- [80] Eduard Hauck, Eike Kiltz, and Julian Loss. A modular treatment of blind signatures from identification schemes. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 345–375. Springer, Cham, May 2019.
- [81] Eduard Hauck, Eike Kiltz, Julian Loss, and Ngoc Khanh Nguyen. Lattice-based blind signatures, revisited. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 500–529. Springer, Cham, August 2020.

- [82] Scott Hendrickson, Jana Iyengar, Tommy Pauly, Steven Valdez, and Christopher A. Wood. Private Access Tokens. Internet-Draft draft-private-access-tokens-01, Internet Engineering Task Force, October 2021. Work in Progress.
- [83] Nicholas Hopper. Proving security of tor’s hidden service identity blinding protocol. <https://www-users.cse.umn.edu/~hoppernj/basic-proof.pdf>, 2013.
- [84] Joseph Jaeger and Stefano Tessaro. Expected-time cryptography: Generic techniques and applications to concrete soundness. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 414–443. Springer, Cham, November 2020.
- [85] Julia Kastner, Julian Loss, and Jiayu Xu. On pairing-free blind signature schemes in the algebraic group model. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part II*, volume 13178 of *LNCS*, pages 468–497. Springer, Cham, March 2022.
- [86] Shuichi Katsumata, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. Round-optimal blind signatures in the plain model from classical and quantum standard assumptions. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 404–434. Springer, Cham, October 2021.
- [87] Shuichi Katsumata, Michael Reichle, and Kaoru Takemure. Adaptively secure 5 round threshold signatures from MLWE/MSIS and DL with rewinding. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 459–491. Springer, Cham, August 2024.
- [88] Jonathan Katz, Julian Loss, and Michael Rosenberg. Boosting the security of blind signature schemes. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 468–492. Springer, Cham, December 2021.
- [89] Neal Koblitz and Alfred Menezes. Another look at non-standard discrete log and diffie-hellman problems. *J. Math. Cryptol.*, 2(4):311–326, 2008.
- [90] Neal Koblitz and Alfred J. Menezes. Another look at “provable security”. *Journal of Cryptology*, 20(1):3–37, January 2007.
- [91] Chelsea Komlo and Ian Goldberg. FROST: Flexible round-optimized Schnorr threshold signatures. In Orr Dunkelman, Michael J. Jacobson, Jr., and Colin O’Flynn, editors, *SAC 2020*, volume 12804 of *LNCS*, pages 34–65. Springer, Cham, October 2020.
- [92] Chelsea Komlo, Ian Goldberg, and T Wilson-Brown. Two-Round Threshold Signatures with FROST. Internet-Draft draft-irtf-cfrg-frost-01, Internet Engineering Task Force, August 2021. Work in Progress.

- [93] Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 239–256. Springer, Berlin, Heidelberg, May 2014.
- [94] Benoît Libert, Marc Joye, and Moti Yung. Born and raised distributively: fully distributed non-interactive adaptively-secure threshold signatures with short shares. In Magnús M. Halldórsson and Shlomi Dolev, editors, *33rd ACM PODC*, pages 303–312. ACM, July 2014.
- [95] Yehuda Lindell. Simple three-round multiparty Schnorr signing with full simulatability. *CiC*, 1(1):25, 2024.
- [96] Yehuda Lindell, Ariel Nof, and Samuel Ranellucci. Fast secure multiparty ECDSA with practical distributed key generation and applications to cryptocurrency custody. Cryptology ePrint Archive, Report 2018/987, 2018.
- [97] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 598–616. Springer, Berlin, Heidelberg, December 2009.
- [98] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Berlin, Heidelberg, April 2012.
- [99] Vadim Lyubashevsky, Leo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehle, and Shi Bai. CRYSTALS – DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. <https://csrc.nist.gov/projects/post-quantum-cryptography/selected-algorithms>.
- [100] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [101] Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Berlin, Heidelberg, December 2005.
- [102] National Institute of Standards and Technology. Multi-Party Threshold Cryptography, 2018–Present. <https://csrc.nist.gov/Projects/threshold-cryptography>.

- [103] National Institute of Standards and Technology. Post-Quantum Cryptography: Additional Digital Signature Schemes, 2022–Present. <https://csrc.nist.gov/projects/pqc-dig-sig>.
- [104] Vassiliy Ilyich Nechaev. Complexity of a determinate algorithm for the discrete logarithm. *Mathematical Notes*, 55(2):165–172, February 1994.
- [105] Jonas Nick, Tim Ruffing, and Yannick Seurin. MuSig2: Simple two-round Schnorr multi-signatures. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 189–221, Virtual Event, August 2021. Springer, Cham.
- [106] Miyako Ohkubo and Masayuki Abe. Security of some three-move blind signature schemes reconsidered. In *The 2003 Symposium on Cryptography and Information Security*, 2003.
- [107] Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 31–53. Springer, Berlin, Heidelberg, August 1993.
- [108] Tatsuaki Okamoto and Kazuo Ohta. Universal electronic cash. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 324–337. Springer, Berlin, Heidelberg, August 1992.
- [109] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 129–140. Springer, Berlin, Heidelberg, August 1992.
- [110] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.
- [111] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon. Technical report, National Institute of Standards and Technology, 2022. <https://csrc.nist.gov/projects/post-quantum-cryptography/selected-algorithms/>.
- [112] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.

- [113] Tim Ruffing, Viktoria Ronge, Elliott Jin, Jonas Schneider-Bensch, and Dominique Schröder. ROAST: Robust asynchronous Schnorr threshold signatures. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 2551–2564. ACM Press, November 2022.
- [114] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 239–252. Springer, New York, August 1990.
- [115] Claus-Peter Schnorr. Efficient identification and signatures for smart cards (abstract) (rump session). In Jean-Jacques Quisquater and Joos Vandewalle, editors, *EUROCRYPT'89*, volume 434 of *LNCS*, pages 688–689. Springer, Berlin, Heidelberg, April 1990.
- [116] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, January 1991.
- [117] Claus-Peter Schnorr. Security of blind discrete log signatures against interactive attacks. In Sihan Qing, Tatsuaki Okamoto, and Jianying Zhou, editors, *ICICS 01*, volume 2229 of *LNCS*, pages 1–12. Springer, Berlin, Heidelberg, November 2001.
- [118] Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Berlin, Heidelberg, May 1997.
- [119] Victor Shoup. Practical threshold signatures. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 207–220. Springer, Berlin, Heidelberg, May 2000.
- [120] Douglas R. Stinson and Reto Strobl. Provably secure distributed Schnorr signatures and a (t, n) threshold scheme for implicit certificates. In Vijay Varadharajan and Yi Mu, editors, *ACISP 01*, volume 2119 of *LNCS*, pages 417–434. Springer, Berlin, Heidelberg, July 2001.
- [121] Katsuyuki Takashima and Atsushi Takayasu. Tighter security for efficient lattice cryptography via the Rényi divergence of optimized orders. In Man Ho Au and Atsuko Miyaji, editors, *ProvSec 2015*, volume 9451 of *LNCS*, pages 412–431. Springer, Cham, November 2015.
- [122] Stefano Tessaro and Chenzhi Zhu. Threshold and multi-signature schemes from linear hash functions. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 628–658. Springer, Cham, April 2023.
- [123] Benedikt Wagner, Lucjan Hanzlik, and Julian Loss. PI-cut-choo! Parallel instance cut and choose for practical blind signatures. Cryptology ePrint Archive, Report 2022/007, Version 20220107:165256, 2022.

- [124] Hoeteck Wee. Threshold and revocation cryptosystems via extractable hash proofs. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 589–609. Springer, Berlin, Heidelberg, May 2011.