

Constructing Provably Secure Identity-Based Signature Schemes.

A Thesis
Submitted for the Degree of
Master of Science (Engineering)
in the Faculty of Engineering

by

Chethan Kamath H



Computer Science and Automation
INDIAN INSTITUTE OF SCIENCE
BANGALORE – 560 012, INDIA

JULY 2015

To my Teachers.

Sir Bedevere: There are ways of telling whether she is a witch.

Peasant 1: Are there? Oh well, tell us.

Sir Bedevere: Tell me. What do you do with witches?

Peasant 1: Burn them.

Sir Bedevere: And what do you burn, apart from witches?

Peasant 1: More witches.

Peasant 2: Wood.

Sir Bedevere: Good. Now, why do witches burn?

Peasant 3: ...because they're made of... wood?

Sir Bedevere: Good. So how do you tell whether she is made of wood?

Peasant 1: Build a bridge out of her.

Sir Bedevere: But can you not also build bridges out of stone?

Peasant 1: Oh yeah.

Sir Bedevere: Does wood sink in water?

Peasant 1: No, no, it floats!... It floats! Throw her into the pond!

Sir Bedevere: No, no. What else floats in water?

Peasant 1: Bread.

Peasant 2: Apples.

Peasant 3: Very small rocks.

Peasant 1: Cider.

Peasant 2: Gravy.

Peasant 3: Cherries.

Peasant 1: Mud.

Peasant 2: Churches.

Peasant 3: Lead! Lead!

King Arthur: A Duck.

Sir Bedevere: ...Exactly. So, logically...

Peasant 1: If she weighed the same as a duck... she's made of wood.

Sir Bedevere: And therefore...

Peasant 2: ...A witch!

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my primary advisor, Dr. Sanjit Chatterjee. I owe him not only for helping develop my academic nous, but also for enlightening me on the subtleties of life, in general. Working under his guidance has been an absolute joy! I am also grateful to Prof. C. E. Veni Madhavan, my co-advisor, for his constant support. I am indebted to faculty of mathematics (in particular, Prof. Dilip Patil and Dr. Chandan Saha) who made me appreciate the elegance, the beauty of mathematics once again. Finally, I would like to thank Prof. C. Pandu Rangan for the kind reviews.

I am thankful to my labmates: Vikas, Nitish, Srikanth, Kumar, Srinivas Vivek, Sudarshan, Kabaleesh, the two Siddharthas, Sayantan, Indradeep, Manjunath, Yogesh, Aniket and Nithin, for their support and forbearance; in particular, I am indebted to Vikas—we were partners in the buddy system of academia. I would also like to thank Ashwin, Maria, Ninad, Subramanya (CSA, IISc), Somindu, Sanjay and Rishiraj (ISI, Kolkata) for some insightful discussions. I am also grateful to my batchmates and our football club for some pleasant, lasting memories.

Last but not the least, I would like to thank my family to whom I owe everything.

Notation and Abbreviations

We adopt the notations that is commonly used in the literature (see [AB09]). Sets are denoted by blackboard bold letters in upper case (*e.g.*, \mathbb{S}, \mathbb{Z}), algorithms using calligraphic letters in upper case (*e.g.*, \mathcal{A}, \mathcal{B}) and tables, schemes *etc.* using upper case letters in fraktur (*e.g.*, \mathfrak{C}). Events are denoted using typewriter font (*e.g.*, $\text{E}, \text{abort}_{1,1}$). We use \mathbb{Z}^+ to indicate the set of positive integers. For an event E , the complementary event is denoted by $\neg\text{E}$. In a group, the discrete-log to a generator g is denoted by $\log g$.

$s \stackrel{\text{U}}{\leftarrow} \mathbb{S}$ denotes picking an element s uniformly at random from the set \mathbb{S} . In general, $\{s_1, \dots, s_n\} \stackrel{\text{U}}{\leftarrow} \mathbb{S}$ denotes picking elements s_1, \dots, s_n independently and uniformly at random from the set \mathbb{S} . In a similar manner, $s \stackrel{\mathbb{S}}{\leftarrow} \mathbb{S}$ and $\{s_1, \dots, s_n\} \stackrel{\mathbb{S}}{\leftarrow} \mathbb{S}$ denote random sampling, but with some underlying probability distribution on \mathbb{S} .

$(y_1, \dots, y_n) \stackrel{\mathbb{S}}{\leftarrow} \mathcal{A}(x_1, \dots, x_m)$ denotes a probabilistic algorithm \mathcal{A} which takes as input (x_1, \dots, x_m) to produce output (y_1, \dots, y_n) . The internal coins ρ of an algorithm, when explicitly given as input, is distinguished from the normal input using a semi-colon, *e.g.*, $y \leftarrow \mathcal{A}(x; \rho)$. Oracle access is indicated using superscripts, *e.g.*, access to an oracle H taking one parameter as inputs is denoted by $y \stackrel{\mathbb{S}}{\leftarrow} \mathcal{A}^{\text{H}(\cdot)}(x)$. The number of dots indicate the number of parameters, *e.g.*, $\text{H}(\cdot, \cdot)$ is an oracle taking two parameters as input. Throughout the thesis, we denote the adversary by \mathcal{A} .

In reductionist security arguments, $\Pi \leq_{\gamma} \mathcal{M}[\mathfrak{P}]$ denotes that Π is γ -reducible to \mathfrak{P} : given an adversary against \mathfrak{P} in the security model \mathcal{M} , one can construct an algorithm that breaks Π with a degradation of γ . $\Pi \stackrel{\gamma}{\Leftarrow} \mathcal{M}[\mathfrak{P}]$ conveys the same message. A straightforward generalisation is $\mathcal{M}'[\mathfrak{P}'] \leq_{\gamma} \mathcal{M}[\mathfrak{P}]$: given an adversary against \mathfrak{P} in the security model \mathcal{M} , one can construct an algorithm that breaks \mathfrak{P}' in the security model \mathcal{M}' with a degradation of γ . (Note that we use long arrows to denote logical implication and short arrows to denote reductions.)

Next, we introduce some notations pertaining to random oracles. The symbol $<$ is used to order the random oracle calls; *e.g.*, $\text{H}(x) < \text{G}(y)$ indicates that the random oracle call $\text{H}(x)$ precedes the

random oracle call $G(y)$. More generally, $H < G$ indicates that the target H-oracle call precedes the target G-oracle call. The convention applies to hash functions as well. The symbol, on the other hand, \prec is used to indicate random oracle dependency; *e.g.* $H \prec G$ indicates that the random oracle G is dependent on the random oracle H. In the discussion involving the forking algorithms, Q_j^k denotes the j^{th} random oracle query in **round k** of simulation.

Finally, in schematics involving IBC, dashed lines indicate an open channel, whereas solid lines indicate a secure channel.

Publications based on this Thesis

1. Sanjit Chatterjee and Chethan Kamath. A closer look at multiple forking: Leveraging (in)dependence for a tighter bound. In *Algorithmica*, pages 1–42, 2015.
2. Sanjit Chatterjee, Chethan Kamath and Vikas Kumar. Galindo-Garcia Identity-Based Signature Revisited. In *Information Security and Cryptology–ICISC’12*, volume 7839 of *Lecture Notes in Computer Science*, pages 456-471, Springer, 2013.
3. Sanjit Chatterjee and Chethan Kamath. From Selective-ID to Full-ID IBS without Random Oracles. To appear in *Security, Privacy and Cryptography–SPACE’13*, volume 8204 of *Lecture Notes in Computer Science*, pages 172-190, Springer, 2013.

Abstract

An identity-based cryptosystem is a public-key system where the public key can be represented by any arbitrary string such as an e-mail address. This notion was introduced by Shamir with the primary goal of simplifying certificate management. In this thesis we study identity-based signatures (IBS)—the identity-based counterpart of digital signatures. In particular, we focus on two techniques—one concrete and the other, generic—used to construct IBS. The concrete scheme we look into is the one proposed by Galindo and Garcia in Africacrypt 2009. This constitutes the primary contribution of the thesis. The generic technique, on the other hand, involves constructing fully-secure IBS given an IBS secure in the selective-identity model.

The Galindo-Garcia IBS (GG-IBS) is derived through a simple and elegant concatenation of two Schnorr signatures. It works in the discrete-log setting but does not require pairing. The security is established through two algorithms (both of) which use the Multiple-Forking (MF) Algorithm to *reduce* the problem of computing the discrete-log to breaking the IBS. The primary contribution of the thesis pertains to the security argument: it turns out that the argument is ridden with flaws and ambiguities and, as a remedy, we provide a new, detailed, security argument. However, the resulting security bound is still quite *loose*, chiefly down to the usage of the MF Algorithm. We explore possible avenues for improving this bound and, to this end, introduce two notions pertaining to random oracles termed *dependency* and *independency*. Incorporating (in)dependency allows us to launch the nested replay attack more effectively than in the MF Algorithm leading to a cleaner, (significantly) tighter security argument for GG-IBS, completing the final piece of the GG-IBS jigsaw.

In the second part of the thesis we look into the notion of *selective-identity* (sID) in the context of IBS. Since its induction, the sID model for identity-based cryptosystems and its relationship with various other notions of security has been extensively studied. As a result, it is a general consensus that the sID model is much weaker than the full-identity (ID) model. The main focus

is on the problem of constructing an ID-secure IBS given an sID-secure IBS without using random oracles—the so-called *standard* model—and with reasonable security degradation. We accomplish this by devising a generic construction which uses as black-box: *i*) a *chameleon* hash function and *ii*) a *weakly*-secure public-key signature. We argue that the resulting IBS is ID-secure but with a tightness gap of $O(q_s)$, where q_s is the upper bound on the number of signature queries that the adversary is allowed to make.

Contents

Acknowledgements	ii
Notation and Abbreviations	iii
Publications based on this Thesis	v
Abstract	vi
Keywords	x
1 Introduction	1
1.1 Identity-Based Signatures	2
1.2 Provable Security	3
1.2.1 Reductionist Security Arguments	5
1.2.2 The Random-Oracle Methodology	6
1.3 Preliminaries	7
1.3.1 Public-Key Signatures	7
1.3.2 Identity-Based Signatures	9
1.3.3 Discrete-Logarithm Assumption	13
1.4 Organisation of the Thesis	14
2 Schnorr Signature and the Oracle Replay Attack	16
2.1 Introduction	16
2.2 Schnorr Signature	16
2.2.1 Construction	17
2.2.2 Security of Schnorr Signature: An Intuition	17
2.2.3 Basic Security	18
2.3 The Oracle Replay Attack	21
2.3.1 The Splitting Lemma	21
2.3.2 Launching the Oracle Replay Attack	24
2.4 General Forking	32
2.5 Nested Oracle Replay Attacks and Multiple Forking	35
2.5.1 Analysis	35
2.5.2 Multiple Forking	41
3 Galindo-Garcia IBS, Revisited	44
3.1 Introduction	44
3.2 Revisiting the Galindo-Garcia Security Argument	45

3.2.1	The Construction	45
3.2.2	The Security Argument and Problems with it	47
3.3	New Security Argument	54
3.3.1	Reduction R_1	56
3.3.2	Reduction R_2	63
3.3.3	Reduction R_3	66
3.3.4	A Comparison with the Original Reduction.	68
4	Galindo-Garcia IBS, Improved	69
4.1	Introduction	69
4.2	Degradation: A Closer Look	70
4.3	Galindo-Garcia IBS, Improved	75
4.3.1	Security Argument	76
4.3.2	Analysis	76
4.3.3	Taking Stock	83
5	From sID IBS to ID IBS without Random Oracles	84
5.1	Introduction	84
5.2	Chameleon Hash Function	85
5.3	The Generic Transformation	86
5.3.1	Security Argument	87
5.4	Transforming from the EU-wID-CMA model	95
6	Conclusions	99
	Appendices	109
A	Galindo-Garcia IBS	109
A.1	The Fixed Security Argument	109
A.1.1	Reduction B_1	109
A.1.2	Reduction B_2	110
A.2	A Security Argument without Wrappers	111
A.2.1	Reduction R_1	112
A.2.2	Reduction R_2	120
A.2.3	Reduction R_3	124
A.3	Reduction R'_1	126
A.3.1	Analysis	129

Keywords

Provable Security, Random Oracles, Schnorr Signature, Oracle Replay Attack, Forking, Identity-Based Signature, Galindo-Garcia IBS, Multiple-Forking, Random-Oracle Dependency, Selective-Identity Model

Chapter 1

Introduction

The concept of identity-based cryptosystems (IBC) was introduced by Shamir in 1984 [Sha85]. In IBC, any arbitrary string such as an e-mail address can act as the public key. In traditional public-key cryptosystems (PKC), users have to exchange public-key certificates before being able to communicate securely. These certificates provide the external binding between the public key and the identity of a user. In some scenarios certificates can prove to be cumbersome. Using IBC one can avoid the complicated certificate management—this, precisely, was Shamir’s foresight.

Identity-based signatures (IBS) extend the notion of digital signatures to the identity-based setting. As in traditional signature schemes, the signer uses her secret key to sign a message. However, the signature can be verified by anyone using the signer’s identity and the “master” public key of the private-key generator (PKG)¹. Since IBS, as in IBC, does not require any certificates to be exchanged, it can be advantageous over the traditional public-key signature (PKS) systems in certain scenarios. IBS, in particular, turns out to be quite practical in wireless sensor networks [LBZ⁺10, TWZL08], BGP protocol [KLS00], MANET routing [HWS⁺06] *etc.*. Therefore, the question of designing efficient and secure IBS is important in the context of applied cryptography.

Although Shamir, in his 1984 paper, gave a simple RSA-based construction for IBS he was unable to establish the security of the scheme. Several “secure” RSA based IBS [FS87, GQ90] have been proposed in the literature since. In recent times, the advent of pairing ushered the way to an increased interest in research on IBC, and consequently IBS mushroomed [CC02, Her05, Hes03].

¹The PKG is a trusted third party whose duty is to create and then communicate the secret keys to the users in the system through a secure channel.

1.1 Identity-Based Signatures

Our aim, in this section, is to give an intuitive feel of IBS. A more formal definition follows in §1.3.2. But first we proceed to describe a basic IBC setup.

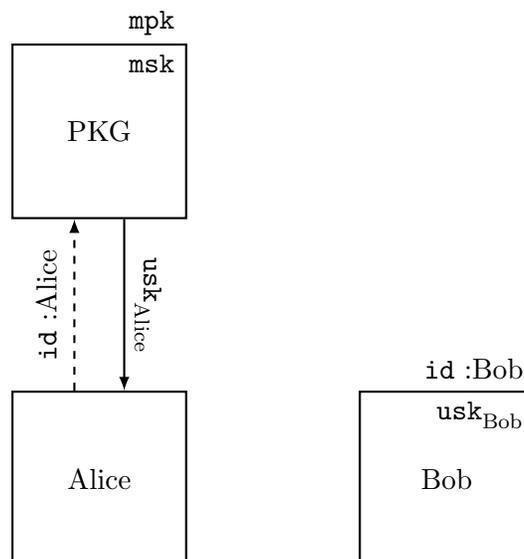


Figure 1.1: A schematic for IBC. The dashed line indicates an open channel, whereas the solid line indicates a secure channel.

A key component in IBC is the private-key generator (PKG). The role of the PKG, as the name suggests, is to generate the private keys² corresponding to the users in the system. This is done through a secure channel as shown for the user Alice with the identity “Alice” in **Figure 1.1**. It also generates the master public key mpk which can, in some sense, be considered to be the “public key” of the PKG. Therefore, on a high level, the PKG can be regarded as a trusted third party that provides the binding between the identity of a user and her secret key—passably analogous to the certifying authority in the traditional public-key infrastructure (PKI).

The key-escrow problem. An inherent problem with IBC, as one could make out, is the “key-escrow” problem. As the PKG can compute secret key corresponding to any user, it can decrypt ciphertext meant for any user (or sign message on behalf of any user) in the system. Some of the alternatives are distributed PKG [BF01] and certificateless PKC [ARP03].

²The terms ‘private key’ and ‘secret key’ will be used interchangeably throughout the thesis.

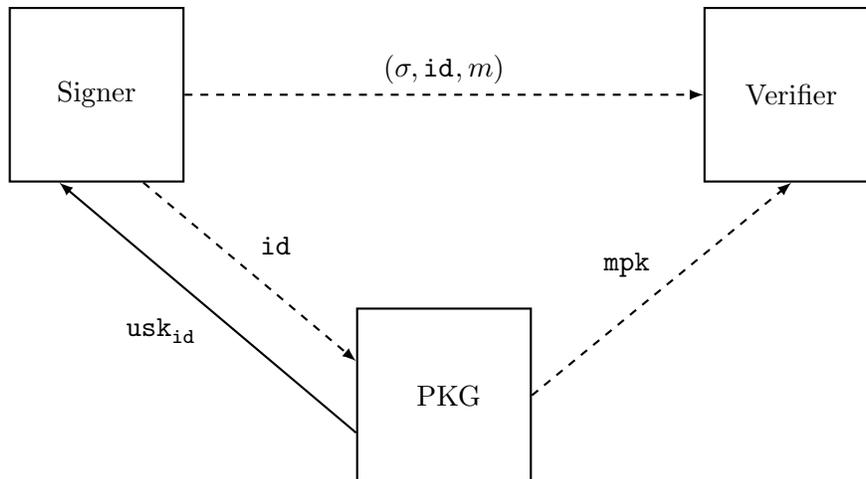


Figure 1.2: A schematic for IBS.

The notion of IBS, if you recall, is the extension of the idea of digital signatures to the identity-based setting. It consists of four probabilistic, polynomial-time (PPT) algorithms **Set-up**, **Key Extraction**, **Signing** and **Verification**. The **Set-up** algorithm is used to generate the master public key mpk with the exact level of security dictated by a security parameter³. **Key Extraction** is used by the PKG to generate the user secret key usk_{id} . The user secret key can be considered as the signature by the PKG with the identity of the user being the message and using the master secret key as the signing key. The algorithms **Signing** and **Verification** correspond to the namesake algorithms in the traditional PKS setting. A signer uses **Sign** to compute a signature σ on a message m with the user secret key being the signing key. Finally, this signature can be verified by the verifier using the **Verification** algorithm, using just the signer’s identity and the master public key.

1.2 Provable Security

Classical cryptography relied heavily on cryptanalysts. Thus, most of the classical cryptosystems that are still in use are the time-tested ones. But nowadays, provable security—the “art” of formally arguing the security of cryptosystems—is fast becoming a *de facto* standard [KM07, Sho04]. Protocol designers are expected to support their construction with an argument establishing its security, *i.e.*

³The security parameter is a variable which measures the level of security in which a cryptographic protocol is to be deployed. It is expressed using the *unary* notation; *e.g.*, for n -bit security level, $n = 1^n$. n determines the resource requirement of the protocol and, consequently, the capability adversary requires for breaking the protocol. For example, if a protocol is deployed in 80-bit security level, it would take an adversary $O(2^{80})$ operations to break the protocol.

the security argument. In a nutshell, security is argued through a “game” between a “challenger”—the party that strives to prove the protocol secure—and an “adversary”—the party that attempts to break the protocol. Let’s start with the description of the adversary, in particular, the abstraction that is used to capture its behaviour.

Modelling the adversary. The adversary is modelled as a probabilistic, polynomial-time (PPT) Turing machine (PTM).⁴ An adversary taking an input x and producing an output y is denoted by $y \stackrel{\$}{\leftarrow} A(x)$. The adversary may have access to some oracles. These are listed in the superscript; *e.g.*, $y \stackrel{\$}{\leftarrow} A^{O(\cdot)}(x)$. A PTM handles its internal decision making through a series of (internal) coin tosses (with some underlying probability distribution). It is desirable, sometimes, to control the internal decision making of an adversary (modelled as a PTM). This can be accomplished by passing the internal coins explicitly as an input, say, on a separate tape to the adversary. In this particular way of modelling, the adversary (in some sense) acts in a deterministic manner. However, it has to be, then, analysed over the space of these internal coins as in the case of a randomised algorithm. As for notation, these internal coins are distinguished from the normal input using a semi-colon, *i.e.* $y \leftarrow A(x; \rho)$. We are now equipped to formally define the notion of “an adversary breaking a protocol”.

Definition 1. Consider an adversary A trying to break a protocol \mathfrak{P} deployed in a n -bit security level. A is said to $(\epsilon(n), t(n))$ -break \mathfrak{P} if it takes time *at most* $t(n)$ and is successful with a probability *at least* $\epsilon(n)$. We say that A is an $(\epsilon(n), t(n))$ -adversary against \mathfrak{P} .⁵

Security models. The security models lay down the schema to be followed while giving a security argument. It is usually described using a security game between the challenger and adversary. The security model sets some ground rules regarding the powers the adversary has and the obligation the challenger has, to the adversary, in the security game. Security models are characterised [PV05] through an *adversarial goal* and an *attack model*—*e.g.*, in the EU-CMA model, the adversary has to produce an existential forgery (the adversarial goal) under the chosen-message attack (the attack model). A detailed discussion on security models and security games is deferred to §1.3.

⁴We consider only computationally-bounded adversaries.

⁵Whenever the context is clear, we abuse the notation by omitting n ; *i.e.*, we say that A is (ϵ, t) -adversary against \mathfrak{P} .

1.2.1 Reductionist Security Arguments

Now that we are aware of the notions relevant to security arguments, let's see how the security is actually argued. A prevalent approach in provable security is argument using the *reductio ad absurdum* principle. Suppose \mathfrak{P} is the cryptographic protocol that we want to argue secure. Let Π be the hard problem⁶ that \mathfrak{P} is based on. Our objective is to argue that $\Pi \implies M[\mathfrak{P}]$: if Π is a hard problem, then the protocol \mathfrak{P} is secure in the security model M .⁷ Another way to establish the same would be through the contrapositive $\neg M[\mathfrak{P}] \implies \neg \Pi$. Thus, the strategy is to assume a hypothetical (polynomial-time and possibly probabilistic) adversary A against \mathfrak{P} and construct an algorithm B —the so-called *reduction* algorithm—that harnesses A , to solve the hard problem Π . Since we presumed Π to be hard, this leads to a contradiction and there cannot exist such an adversary. Or, in classical complexity theory jargon, solving the hard problem Π is *reduced* to breaking the protocol \mathfrak{P} , *i.e.*, $\Pi \leq M[\mathfrak{P}]$ (or, $\Pi \Leftarrow M[\mathfrak{P}]$). This is illustrated⁸ in **Figure 1.3**.

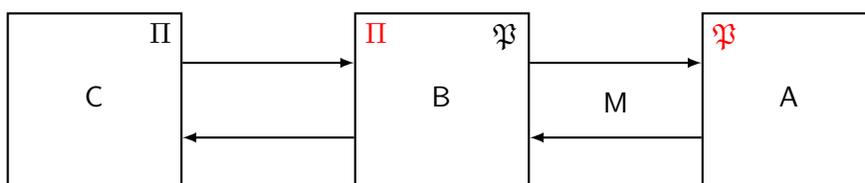


Figure 1.3: Argument through contradiction.

The notion can be easily extended to base security of one protocol on another. In order to base the security of a protocol \mathfrak{P} (in a security model M) on another protocol \mathfrak{P}' (which is secure in a security model M'), one has to establish $M'[\mathfrak{P}'] \leq M[\mathfrak{P}]$: given an adversary against \mathfrak{P} in the security model M , one has to construct an algorithm that breaks \mathfrak{P}' in the security model M' (resulting in a contradiction).

The reduction algorithm. Let's see how the reduction algorithm B works. Recall, from the preceding discussion, that B 's objective is to harness A to solve Π . Suppose that B receives the instance of the hard problem from its challenger C . It initiates the security game M with the adversary A by sending a challenge instance of \mathfrak{P} . Next, B *simulates* the environment of protocol

⁶Computationally infeasible problem. See §1.3.3 for an example of hard problem.

⁷Recall that we use long arrow to denote logical implication and short arrow to denote reductions.

⁸We follow this convention of block diagrams throughout the thesis. Moreover, oracle access is indicated in the bottom right corner of the blocks.

\mathfrak{P} to the adversary A . Note that the simulation has to be faithful⁹ to the actual execution of the protocol. At the end of the simulation, if A succeeds in breaking \mathfrak{P} then B has to (somehow) to solve Π . If B is indeed successful, then we would have reduced Π to \mathfrak{P} . Thus, B plays the role of *both* challenger and adversary (in the respective security games).

Tightness of security reductions. Returning to the preceding discussion, let's suppose that A is an (ϵ, t) -adversary against \mathfrak{P} . Also, let the reduction B that we designed (ϵ', t') -solve Π . The tightness gap for B is defined to be

$$\gamma_B = \frac{t'/\epsilon'}{t/\epsilon}.$$

We say that the security of \mathfrak{P} “degrades” by a factor of γ_B . A reduction B is said to be “tight” if $t' \approx t$ and $\epsilon' \approx \epsilon$ and, by implication, γ_B is of constant order ($O(1)$). In other words, B has to simulate A only a constant number of times before solving Π with a “considerable” advantage. Likewise, B has polynomial degradation if γ_B is a polynomial in n (and so forth).

1.2.2 The Random-Oracle Methodology

Hash functions are indispensable to cryptography. The random-oracle methodology was introduced with the aim of simplifying the security arguments of cryptographic protocols that use hash functions. Although the methodology was first used, in the context of security arguments, by Fiat and Shamir in their zero-knowledge proof [FS87], it was formalised (and popularised) by Bellare and Rogaway [BR93].

The idea is to, first, design an ideal system in which all parties, including the adversary A , have oracle access to a truly random function that represents the hash function—the so-called *random oracle*—and prove the security of this ideal system. During the simulation, the random oracle is under the control of the challenger C (roughly speaking, see [FLR⁺10] for refinements) and the adversary is allowed to make (a bounded number of) queries (Q_i) to this oracle. The random oracle is then instantiated by using a cryptographic hash function, say MD5 or SHA-2, giving us an instance of the ideal system in the real world (see **Figure 1.4**).

Protocols without random oracles—in the so-called *standard model*, that is—turn out to be comparatively involved; random oracle assumption is known to simplify the construction as well as the security argument. Plus, the security arguments are tighter (cf. [BF01] and [Wat05]). Moreover,

⁹Technically speaking, the adversary should not be able to distinguish the simulation of the protocol environment from the actual execution of the protocol.

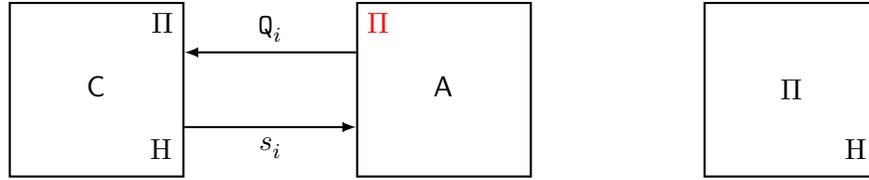


Figure 1.4: Random Oracle Methodology

there exist protocols (Schnorr signature [Sch91] for example) for which the only known security arguments are the ones based on the random oracle assumption. However, proofs in the random-oracle model provide only a *heuristic* support for actual security [FLR⁺10].

1.3 Preliminaries

In this section, we give the formal definition of PKS and IBS, along with their relevant security models. We also define the discrete-logarithm assumption—the hardness assumption on which the objects of our interest (Schnorr signature and Galindo-Garcia IBS) are based on. We cover some of the existing (and relevant) techniques for construction of IBS as well.

1.3.1 Public-Key Signatures

Definition 2 (Public-Key Signature). A PKS scheme consists of three PPT algorithms $\{\mathsf{K}, \mathsf{S}, \mathsf{V}, \dots\}$.

Key Generation, $\mathsf{K}(1^n)$: It takes as input the security parameter n (in unary) and outputs the public key pk and the secret key sk .

Signing, $\mathsf{S}(m, \mathsf{sk})$: It takes as input a message m and the secret key of the user sk to generate a signature σ .

Verification, $\mathsf{V}(\sigma, m, \mathsf{pk})$: It takes as input the signature σ , the message m and the public key of the user pk . It outputs a bit op which is 1 if σ is valid signature on m or 0 if the signature is invalid.

The standard *correctness* condition: $1 \leftarrow \mathsf{V}(\mathsf{S}(m, \mathsf{sk}), m, \mathsf{pk})$, should be satisfied for all m and $(\mathsf{pk}, \mathsf{sk}) \stackrel{\$}{\leftarrow} \mathsf{K}(1^n)$. The standard security notion for PKS schemes: existential unforgeability under chosen-message attack (EU-CMA) [GMR88], is defined below.

Definition 3 (EU-CMA Game). The security of a PKS scheme in the EU-CMA model is argued in terms of the following game between a challenger C and an adversary A .

Set-up: C runs K to obtain a public key pk and a secret key sk . It passes pk as the challenge public key to A .

Signature query, $O_s(m)$: A can adaptively make signature queries to an oracle O_s . C responds by running the algorithm S to obtain a signature corresponding to m . It then sends σ to A .

Forgery: A outputs a signature $\hat{\sigma}$ on a message \hat{m} and *wins* the game if the forgery is *i) valid*: $\hat{\sigma}$ passes the verification on \hat{m} ; and *ii) non-trivial*: A has not queried the signature oracle with \hat{m} .

The advantage that A has in the above game, denoted by $\mathbf{Adv}_A^{\text{EU-CMA}}(n)$, is defined as the probability with which it wins the above game, *i.e.*

$$\mathbb{P}(1 \leftarrow V(\hat{\sigma}, \hat{m}, pk) : (sk, pk) \xleftarrow{\$} K(1^n); (\hat{\sigma}, \hat{m}) \xleftarrow{\$} A^{O_s(\cdot)}(pk))$$

provided $\hat{\sigma}$ is a non-trivial forgery on \hat{m} . A is said to be an (ϵ, t, q_s) -forger of an PKS scheme if it has advantage of at least ϵ in the above game, runs in time at most t and queries the signature oracle at most q_s times.

Definition 4 (Full-Security of PKS). We say that a PKS is existentially unforgeable against a chosen-message attack (fully-secure, in short) if for any PPT adversary A the function $\mathbf{Adv}_A^{\text{EU-CMA}}(n)$ is negligible¹⁰ in n .

Remark 1 (No-Message Attack). The adversarial task of producing an existential forgery under the no-message attack (EU-NMA), also known as the key-only attack (EU-KOA), can be regarded as a particular, weaker, case of the task of producing an existential forgery under chosen-message attack, with $q_s = 0$.

Weakly-secure PKS. A weaker notion of PKS security that we are interested in is existential forgery under *generic* chosen-message attack (EU-GCMA) [GMR88]. The distinguishing feature of EU-GCMA model is that the adversary commits, beforehand, to a set of messages $\{m_1, \dots, m_{q_s}\}$ to the challenger. In response, the challenger gives it, along with challenge public key, a set of signatures corresponding to the committed messages. At the end of the game, the adversary has to forge on a message that is *not* part of the committed set. A more formal definition follows.

¹⁰A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is *negligible* if for any $n > 0$, we have $|f(x)| < 1/k^n$ for sufficiently large x [BF01].

Definition 5 (EU-GCMA Game). The security of a PKS scheme in the EU-GCMA model is argued in terms of the following game between a challenger C and an adversary A .

Commitment: A commits to a set of messages $\tilde{M} := \{m_1, \dots, m_{q_s}\}$.

Response: C runs G to obtain the public key \mathbf{pk} and the secret key \mathbf{sk} . It runs the signing algorithm S to obtain a set of signatures $\{\sigma_1, \dots, \sigma_{q_s}\}$ corresponding to \tilde{M} . Finally, it sends the challenge public key \mathbf{pk} along with $\{\sigma_1, \dots, \sigma_{q_s}\}$ to A .

Forgery: A outputs a signature $\hat{\sigma}$ on a message \hat{m} and *wins* the game if the forgery is *i) valid*: $\hat{\sigma}$ passes the verification on \hat{m} ; and *ii) non-trivial*: \hat{m} is not a part of the committed set \tilde{M} .

The advantage that A has in the above game, denoted by $\mathbf{Adv}_A^{\text{EU-GCMA}}(n)$, is defined as the probability with which it wins the above game, *i.e.*

$$\mathbb{P}(1 \leftarrow V(\hat{\sigma}, \hat{m}, \mathbf{pk}) \wedge \hat{m} \notin \tilde{M} : \tilde{M} \xleftarrow{\$} A(q_s); (\mathbf{sk}, \mathbf{pk}) \xleftarrow{\$} K(1^n); (\hat{\sigma}, \hat{m}) \xleftarrow{\$} A(\mathbf{pk}, \sigma_1, \dots, \sigma_{q_s}))$$

An adversary A is said to be an (ϵ, t, q_s) -forger of an PKS scheme if it has advantage of at least ϵ in the above game, runs in time at most t , after initially having committed to a set of q_s messages.

Definition 6 (Weak-Security of PKS). We say that a PKS is existentially unforgeable against a generic chosen-message attack (*weakly-secure*, in short) if for any PPT adversary A the function $\mathbf{Adv}_A^{\text{EU-GCMA}}(n)$ is negligible in n .

1.3.2 Identity-Based Signatures

Definition 7 (Identity-Based Signature). An IBS scheme consists of four PPT algorithms $\{G, E, S, V\}$ described below.

Set-up, $G(1^n)$: It takes as input the security parameter n (in unary). It outputs the master secret key \mathbf{msk} and the master public key \mathbf{mpk} .

Key Extraction, $E(\text{id}, \mathbf{msk})$: It takes as input the user's identity id , the master secret key \mathbf{msk} to generate a secret key \mathbf{usk} for the user.

Signing, $S(\text{id}, m, \mathbf{usk})$: It takes as input the user's identity id , a message m and the user's secret key \mathbf{usk} to generate a signature σ .

Verification, $V(\sigma, \text{id}, m, \text{mpk})$: It takes as input a signature σ , a message m , an identity id and the master public key mpk . It outputs a bit op which is 1 if σ is a valid signature on (id, m) or 0 if the signature is invalid.

The standard *correctness* condition: $1 \leftarrow V(S(\text{id}, m, \text{usk}), \text{id}, m, \text{mpk})$, should be satisfied for all $\text{id}, m, (\text{msk}, \text{mpk}) \xleftarrow{\$} G(1^n)$ and $\text{usk} \xleftarrow{\$} E(\text{id}, \text{msk})$. As for security, we use the detailed EU-ID-CMA model given in [BNN04].

Definition 8 (EU-ID-CMA Game¹¹). The security of an IBS scheme in the EU-ID-CMA model is argued in terms of the following game between a challenger C and an adversary A .

Set-up: C runs G to obtain the master public key mpk and the master secret key msk . It passes mpk as the challenge master public key to A .

Queries: A can adaptively make extract queries to an oracle O_E and signature queries to an oracle O_s . These queries are handled as follows.

- **Extract query**, $O_E(\text{id})$: A asks for the secret key of a user with identity id . If there has already been an extract query on id , C returns the user secret key that was generated during the earlier query. Otherwise, C uses the knowledge of msk to run E and generate the user secret key usk , which is then passed on to A .
- **Signature query**, $O_s(\text{id}, m)$: A asks for the signature of a user with identity id on a message m . C first obtains, as specified the extract query, a user secret key usk corresponding to id . Next, it uses the knowledge of usk to run S and generate a signature σ , which is passed to A .

Forgery: A outputs a signature $\hat{\sigma}$ on an identity $\hat{\text{id}}$ and a message \hat{m} , and *wins* the game if the forgery is *i) valid*: $\hat{\sigma}$ passes the verification on $(\hat{\text{id}}, \hat{m})$; and *ii) non-trivial*: A has not queried the extract oracle with $\hat{\text{id}}$, *nor* has it queried the signature oracle with $(\hat{\text{id}}, \hat{m})$.

The advantage that A has in the above game, denoted by $\mathbf{Adv}_A^{\text{EU-ID-CMA}}(n)$, is defined as the probability with which it wins the above game, *i.e.*

$$\mathbb{P}(1 \leftarrow V(\hat{\sigma}, \hat{\text{id}}, \hat{m}, \text{mpk}) : (\text{msk}, \text{mpk}) \xleftarrow{\$} G(1^n); (\hat{\sigma}, \hat{\text{id}}, \hat{m}) \xleftarrow{\$} A^{O_E(\cdot), O_s(\cdot, \cdot)}(\text{mpk}))$$

¹¹The security game in [BNN04], *i.e.* $\mathbf{Exp}_{\text{IBS}, \mathbb{F}}^{\text{uf-cma}}$, is explained in terms of the three oracles: INIT, CORR and SIGN. Here we use an *equivalent* formulation in terms of Extract and Signature queries.

provided $\hat{\sigma}$ is a non-trivial forgery on $(\hat{\text{id}}, \hat{m})$. An adversary A is said to be an $(\epsilon, t, q_\epsilon, q_s)$ -forger of an IBS scheme if it has advantage of at least ϵ in the above game, runs in time at most t and makes at most q_ϵ (resp. q_s) extract (resp. signature) queries.

Definition 9 (ID-Security of IBS). We say that an IBS is EU-ID-CMA-secure (ID-secure, in short) if for any PPT adversary A the function $\text{Adv}_A^{\text{EU-ID-CMA}}(n)$ is negligible in n .

The notion of selective-identity. The selective-identity (sID) model for identity-based cryptographic schemes was introduced in [CHK03]. The distinguishing feature of this model is that the adversary has to commit, beforehand, to a “target” identity—*i.e.*, the identity which it intends to eventually forge on.

Definition 10 (EU-sID-CMA Game). The security of an IBS scheme in the EU-sID-CMA model is argued in terms of the following game between a challenger C and an adversary A .

Commitment: A commits to a target identity $\tilde{\text{id}}$.

Set-up: C runs G to obtain the master keys (mpk, msk) . It passes mpk as the challenge master public key to A .

Queries: A can adaptively make extract queries to an oracle O_E and signature queries to an oracle O_s . It is restricted though from making extract query on the target identity $\tilde{\text{id}}$. These queries are handled as in the EU-ID-CMA game.

Forgery: A outputs a signature $\hat{\sigma}$ on a message \hat{m} and an identity $\tilde{\text{id}}$, and *wins* the game if the forgery is *i)* *valid*: $\hat{\sigma}$ passes the verification on $(\tilde{\text{id}}, \hat{m})$; and *ii)* *non-trivial*: A has queried the signature oracle with $(\tilde{\text{id}}, \hat{m})$.

The advantage that A has in the above game, denoted by $\text{Adv}_A^{\text{EU-sID-CMA}}(n)$, is defined as the probability with which it wins the game, *i.e.*

$$\mathbb{P}(1 \leftarrow V(\hat{\sigma}, \tilde{\text{id}}, \hat{m}, \text{mpk}) : \tilde{\text{id}} \xleftarrow{\$} A; (\text{msk}, \text{mpk}) \xleftarrow{\$} G(1^n); (\hat{\sigma}, \tilde{\text{id}}, \hat{m}) \xleftarrow{\$} A^{O_E(\cdot), O_s(\cdot, \cdot)}(\text{mpk}))$$

provided $\hat{\sigma}$ is a non-trivial forgery on $(\tilde{\text{id}}, \hat{m})$. An adversary A is said to be an $(\epsilon, t, q_\epsilon, q_s)$ -forger of an IBS scheme in the EU-sID-CMA model if it has advantage of at least ϵ in the above game, runs in time at most t and makes at most q_ϵ (resp. q_s) extract (resp. signature) queries.

Definition 11 (sID-Security of IBS). We say that an IBS is EU-sID-CMA-secure (sID-secure, in short) if for any PPT adversary A the function $\mathbf{Adv}_A^{\text{EU-sID-CMA}}(n)$ is negligible in n .

Remark 2. If the security argument models the hash functions as random oracles, the security games are to be modified, appropriately, to accommodate queries to the random oracles.

1.3.2.1 Existing Techniques for Constructing IBS

The task of constructing IBS is generally considered to be much easier than constructing IBE. [BNN04] contains a comprehensive list of such techniques, along with the security arguments. We discuss, although briefly, three techniques—one concrete and two generic—that are relevant to the thesis.

From Schnorr signatures. The concrete technique that we are interested in is based on the Schnorr signature scheme [Sch91]. Galindo and Garcia [GG09] devised a way concatenate two Schnorr signatures to produce an identity-based signature. The user secret key can be considered as the Schnorr signature by the PKG on the identity of the user, using the master secret key as the signing key; analogously, the signature on a message by a user is the Schnorr signature, by that user, on the message using her user secret key as the signing key. The resultant IBS is simple and efficient. We analyse its construction and security argument in the first part of our thesis.

The “folklore” construction. Although the cryptographic research community was long aware of this (generic) technique, it was formalised by Bellare *et al.* [BNN04]. It involves two applications of PKS (certificate). In a nutshell, the PKG generates a (public-private) key-pair for a user and then *binds* it to the identity of the user by using a certificate generated using its (master) secret key. Thus the PKG, to a large extent, plays the role of a certifying authority in PKI. Interestingly, no random oracles are involved.

The (existence of) folklore construction reduces the task of constructing an IBS to that of constructing a PKS. A PKS, on the other hand, can be derived from a weakly-secure PKS using a chameleon hash function (CHF) [KR00]. This approach was used implicitly in [ST01] and, later, formalised in [HW09]. This, in turn, implies the task of constructing IBS is reduced to that of constructing a weakly-secure PKS and a CHF. It also reiterates the fact that it is much easier (even in the standard model) to construct IBS schemes.

From sID IBS. Recall that the selective-identity (sID) model for IBC is considered to be weaker than the full-identity (ID) security model. However, it is much easier to design a protocol that is sID-secure than one that is ID-secure. Therefore, transforming an sID IBS to an ID IBS, in a *generic* manner, is an interesting problem to pursue, both from theoretical and practical point of view. An efficient (comparatively) black-box method to convert a sID-secure IBE scheme to ID security was suggested in [BF01]. But the method relies on random oracles [BR93]. This was followed by [BB04a], in which the problem is solved in the standard model, albeit with an *exponential* loss of tightness. Both these methods can be adapted to IBS. To the best of our knowledge, there doesn't exist an efficient generic transformation from sID IBS to ID IBS in the standard model. In the second half of our thesis we concentrate on this exact problem.

1.3.3 Discrete-Logarithm Assumption

Let G_{DL} be a (randomised) group generator: it takes as input a security parameter n (in unary) and outputs (\mathbb{G}, g, p) , where \mathbb{G} is a cyclic group of prime order p generated by g .¹² The discrete-logarithm problem (DLP) is defined through the following game between a challenger C and an adversary A (see **Figure 1.5**).

Definition 12 (DLP Game). C invokes the group generator to generate a cyclic group: $(\mathbb{G}, g, p) \xleftarrow{\$} G_{DL}(n)$. Next, it sends to A the challenge DLP instance (\mathbb{G}, g, p, h) , where $h := g^\alpha$ with $\alpha \xleftarrow{U} \mathbb{Z}_p$. A wins the game if it correctly finds the discrete logarithm of h (with respect to the generator g), *i.e.*, $\alpha' = \alpha$. The advantage A has solving the DLP is the probability with which it wins the game, *i.e.*,

$$\mathbf{Adv}_A^{\text{DLP}}(n) \stackrel{\text{def}}{=} \mathbb{P}(\alpha' = \alpha : (\mathbb{G}, g, p) \xleftarrow{\$} G_{DL}(n); \alpha \xleftarrow{U} \mathbb{Z}_p; \alpha' \xleftarrow{\$} A(\mathbb{G}, p, g, g^\alpha)).$$

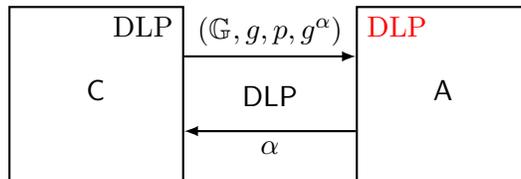


Figure 1.5: The DLP Game.

Definition 13 (Discrete-Logarithm Assumption). The (ϵ, t) -discrete-log assumption holds if no adversary that takes time at most t has advantage at least ϵ in the DLP game. In general, the

¹²We sometimes omit the description of the group generator for sake of convenience and, also, consistency.

(asymptotic) discrete-log assumption holds if for all PPT adversaries \mathbf{A} the function $\mathbf{Adv}_{\mathbf{A}}^{\text{DLP}}(n)$ is negligible in n ; we say that \mathbf{G}_{DL} generates DL-secure groups [GG09].

1.4 Organisation of the Thesis

The Galindo-Garcia IBS (GG-IBS) is based on the Schnorr signature scheme [Sch91]. Hence, we devote **Chapter 2** to the particulars of Schnorr signature: its construction and security argument. This, in turn, lends us some insight into GG-IBS, especially its security argument. The machinery of (oracle) replay attack¹³ [PS00] plays a pivotal role in the security argument of Schnorr signature and, therefore, we place emphasis on its development. This includes a discourse on the two techniques that are used to launch the replay attack: the Rewinding Technique and the Replicating Technique. Plus, we discuss an abstraction of the replay attack, called General Forking [BN06], that aids in simplifying security arguments. Finally, we augment the “elementary” replay attack (which involves only *one* random oracle) to include multiple random oracles and nested forkings. This yields the Multiple-Forking Algorithm [BPW12].

With the know-how on Schnorr signature in our arsenal, we revisit GG-IBS itself in **Chapter 3**. GG-IBS is derived through a simple and elegant concatenation of two Schnorr signatures. It works in the discrete-log setting but does not require pairing. The security is established by reducing the problem of computing the discrete-log to breaking the IBS, through two (reduction) algorithms \mathbf{B}_1 and \mathbf{B}_2 . Both \mathbf{B}_1 and \mathbf{B}_2 employ the Multiple-Forking (MF) Algorithm. The focus of the chapter is on this security argument: we find that it is riddled with flaws and ambiguities. In particular, we show that the reduction \mathbf{B}_1 in [GG09] *fails* in the standard security model for IBS [BNN04], whereas the reduction \mathbf{B}_2 is *incomplete*. As a remedy, we provide a new, detailed, security argument which allows tighter reductions [CKK13].

In **Chapter 4**, we explore possible avenues for further improvement in the tightness of GG-IBS security argument. As already stated, the machinery of MF Algorithm plays a central role in the security argument of GG-IBS. However, it is also responsible for the huge degradation in the concrete security of the IBS. Thus, a natural question arises: *Can the nested replay attack (on the GG-IBS adversary) be launched more effectively than in the MF Algorithm?* To this end, we introduce the two notions pertaining to random oracles: *dependency* and *independency*. Independency follows naturally for GG-IBS, whereas, dependency has to be induced through a

¹³We will use the terms forking and (the process of launching) oracle replay attack interchangeably.

tweak in construction. As it turns out, incorporating dependency and independency, in conjunction, results in a cleaner, (significantly) tighter security argument for GG-IBS completing the final piece of the GG-IBS jigsaw.

In **Chapter 5**, the second part of the thesis, we look into the concept of *selective-identity* (sID) [CHK03] in the context of IBS. Since its induction, the sID model for identity-based cryptosystems and its relationship with various other notions of security has been extensively studied. As a result, it is a general consensus that the sID model is much weaker than the full-identity (ID) model. The main focus is on the problem of constructing an ID-secure IBS given an sID-secure IBS in the *standard model* and with reasonable security degradation. We accomplish this by devising a generic construction which uses as black-box: *i*) a chameleon hash function and *ii*) a weakly-secure public-key signature [CK13b]. We argue that the resulting IBS is ID-secure but with a tightness gap of $O(q_s)$, where q_s is the upper bound on the number of signature queries that the adversary is allowed to make.

Chapter 2

Schnorr Signature and the Oracle Replay Attack

2.1 Introduction

We devote this chapter to the particulars of Schnorr signature scheme: its construction and security argument. We also emphasise on developing the machinery of *oracle replay attack* [PS00] which plays a pivotal role in the security argument of Schnorr signature scheme (as well as other class¹ of signature schemes [ELG85, Oka93]). Since the Galindo-Garcia IBS (GG-IBS) scheme is based on Schnorr signature, the discussion lends us some insight into GG-IBS. Plus, we discuss an abstraction of the replay attack, called General Forking [BN06], that aids in simplifying security arguments. Finally, we augment the “elementary” replay attack (which involves only *one* random oracle) to include multiple random oracles and nested forkings. This yields the Multiple-Forking Algorithm [BPW12].

2.2 Schnorr Signature

The Schnorr signature scheme is based on the discrete-log problem (see **Definition 12**). It can be concocted from the Schnorr identification scheme [Sch90] by using the Fiat-Shamir transformation [FS87]. We now describe a (minor) variant of this signature scheme that helps us in studying GG-IBS.

¹To be precise, the signature schemes obtained from three-round identification schemes (Σ -protocols) through the Fiat-Shamir transformation [FS87].

2.2.1 Construction

Parameter Generation^a, $G(1^n)$: Invoke the group generator G_{DL} (on 1^n) to obtain (\mathbb{G}, g, p) . In addition, let H be a hash function $H : \{0, 1\}^* \mapsto \mathbb{Z}_p$. Return $\mathbf{pp} := (\mathbb{G}, g, p, H)$ as public parameters.

Key Generation, $K(\mathbf{pp})$: Select $z \xleftarrow{\mathbb{U}} \mathbb{Z}_p$ and set $Z = g^z$. Return $\mathbf{sk} := (z, \mathbf{pp})$ as the secret key and $\mathbf{pk} := (Z, \mathbf{pp})$ as the public key.

Signing, $S(m, \mathbf{sk})$: Select $r \xleftarrow{\mathbb{U}} \mathbb{Z}_p$ and set $R := g^r$. Return $\sigma := (y, R) \in \mathbb{Z}_p \times \mathbb{G}$ as a signature on the message m , where

$$y := r + zc \quad \text{and} \quad c := H(m, R).$$

Verification, $V(\sigma, m, \mathbf{pk})$: Parse σ as (y, R) and compute $c := H(m, R)$. The signature is valid if

$$g^y = R \cdot Z^c.$$

Figure 2.1: The Schnorr Signature Scheme.

^aFor convenience, we have split the Key Generation algorithm into a Parameter Generation algorithm G and an actual Key Generation algorithm K . This aids in generating multiple public-private key-pairs in the same group setting.

Remark 3. We use the convention S_H to denote the usage of hash function H by the Signing algorithm S . We use this convention as the Signing algorithm may be used with different hash functions. Simply put, the hash function too is considered to be a parameter to the Signing algorithm. This applies to the Verification algorithm V as well.

2.2.2 Security of Schnorr Signature: An Intuition

Suppose that we are working in a group \mathbb{G} where the discrete-log assumption holds (see **Definition 13**). Consider an adversary A with the capability to carry out chosen-message attack on the Schnorr signature scheme. Let $\{\sigma_1, \dots, \sigma_n\}$ be the signatures that A procures. Moreover, (for each $i \in \{1, \dots, n\}$) let: *i*) σ_i be a signature on m_i , and *ii*) be of the form (y_i, R_i) with, $R_i = g^{r_i}$, $y_i = r_i + c_i z$ and $c_i = H(m_i, R_i)$. The adversary's view is captured by the following system of n

linear congruences (modulo p) in the $(n + 1)$ unknowns $\{r_1, \dots, r_n, z\}$.

$$\begin{pmatrix} 1 & 0 & \cdots & 0 & c_1 \\ 0 & 1 & \cdots & 0 & c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & c_n \end{pmatrix} \times \begin{pmatrix} r_1 \\ r_2 \\ \vdots \\ r_n \\ z \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} \quad (2.1)$$

It is evident from basic linear algebra that the system in (2.1) has p solutions. Besides, the discrete-log assumption ensures that it is infeasible for A to extract information regarding the r_i s from the corresponding R_i s. Thus, it is hard for A to learn the value of the secret key z . *However*, if (by some means) it secures two equations that use the same randomiser r but with different hash value c , it can easily solve for z (as the system is reduced to n linear congruences in n unknowns). Keeping this in mind, we argue the security of Schnorr signature.

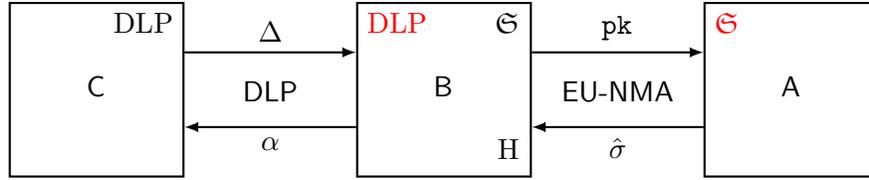
2.2.3 Basic Security

Let's start with the simpler (but weaker, see **Remark 1**) EU-NMA model for PKS. For the time being, we restrict ourselves to a particular class of (strong) adversaries which we refer to as “perfect” adversaries. A perfect adversary, against a particular scheme, has the maximum (possible) advantage in breaking the scheme (in some appropriate security model). Therefore, for the perfect adversary A that we consider against the Schnorr signature scheme, in the EU-NMA model, we have $\text{Adv}_A^{\text{EU-NMA}}(n) = 1$.

Claim 0.1 (Basic security of Schnorr signature). *Given a perfect adversary A against the Schnorr signature scheme in the EU-NMA model, it is possible to construct an algorithm B that solves the DLP, provided the hash function H is modelled as a random oracle with an upper bound of q queries. In other words,*

$$\text{DLP} \leq_{O(1)} \text{EU-NMA} [\text{Schnorr Signature}].$$

Argument. Our objective is to construct an algorithm B that simulates the protocol environment

Figure 2.2: (Basic) Security Argument for Schnorr PKS (denoted by \mathfrak{S}).

for the Schnorr signature adversary A in the EU-NMA model. At the end of the simulation, A (being a perfect adversary) forges successfully and B , in turn, has to somehow break the DLP. Thus, B plays the role of a challenger to A in the EU-NMA game; at the same time, it plays the role of an adversary in the DLP game (see **Figure 2.2**).

First round of simulation. The simulator B receives a DLP instance $\Delta = (\mathbb{G}, g, p, g^\alpha)$ from its own challenger C and its objective is to extract α . B embeds the problem instance g^α in Z and, thus, passes $\text{pk} := (\mathbb{G}, g, p, g^\alpha)$ as the challenge public key to A . A is allowed access to a random oracle H but it *does not* have access to the signature oracle. The adversary A makes a series of hash oracle queries $\{Q_1, \dots, Q_q\}$ (with each Q_i of the form $H(m_i, R_i)$) which the challenger responds to with $\{c_1, \dots, c_q\} \stackrel{U}{\leftarrow} \mathbb{Z}_p$. At the end of the simulation, it successfully forges $\hat{\sigma} = (y, R)$ on some message \hat{m} (with $R = g^r$, $c = H(\hat{m}, R)$ and $y = r + \alpha c$). Since H is a random oracle, for A to have produced the forgery $\hat{\sigma}$ with a non-negligible advantage, it must have made the query $H(\hat{m}, R)$ at some juncture during its simulation². In other words, there *must* exist an index $1 \leq I \leq q$ such that the query Q_I is $H(\hat{m}, R)$ (see **Figure 2.3**). This index is termed the “target” H-index and the query, the “target” H-oracle query. To paraphrase, the random oracle query that is involved in the fabrication of the forgery, by the adversary, is called the target oracle query; the index of this query is called the target index.

² Let S denote the event that A successfully forges $\hat{\sigma} = (y, R)$ on some message \hat{m} . In addition, let H denote the event that A queries the random oracle with $H(\hat{m}, R)$ —the “target” H-oracle query. Thus, we have

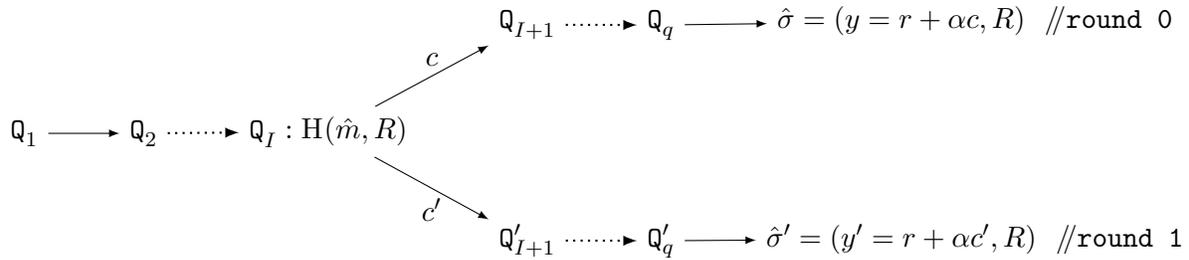
$$\begin{aligned} \mathbb{P}(S) &= \mathbb{P}(S \wedge H) + \mathbb{P}(S \wedge \neg H) \\ &= \mathbb{P}(S \wedge H) + \mathbb{P}(S \mid \neg H)\mathbb{P}(\neg H). \end{aligned}$$

The term $S \mid \neg H$ indicates that A has forged successfully, but, without actually having made the target H-oracle query $H(\hat{m}, R)$. However, the value of $H(\hat{m}, R)$ is necessary in order to forge. Therefore, in this situation, the best that the adversary can do to conjure up with the value of $H(\hat{m}, R)$ is to guess it. However, the probability with which the adversary could have *correctly* anticipated the output is negligible ($1/p$ to be precise). Thus the $\mathbb{P}(S \mid \neg H)$ term (weighed down by the $1/p$ factor) is negligible and as a result $\mathbb{P}(S) \approx \mathbb{P}(S \wedge H)$. Simply put, A *has* to make the target random oracle query to be able to forge with a non-negligible probability.

$$Q_1 \xrightarrow{c_1} Q_2 \cdots \rightarrow Q_I : H(\hat{m}, R) \xrightarrow{c_I} Q_{I+1} \cdots \rightarrow Q_q \xrightarrow{c_q} \hat{\sigma} = (y = r + \alpha c, R)$$

Figure 2.3: The “target” random oracle query $H(\hat{m}, R)$.

Second round of simulation. Now, let’s suppose that the simulator \mathbf{B} (by some means) manages to “turn back the clock” to the juncture when \mathbf{A} made the target oracle query Q_I . Next, it simulates \mathbf{A} again from Q_I but, this time, with a different set of outputs to the random oracle queries (starting with Q_I)³. This constitutes \mathbf{B} carrying out an *oracle replay attack* on the adversary \mathbf{A} —we say that the adversary \mathbf{A} is “forked” at the point Q_I .⁴ Let’s denote this new round of simulation of the

Figure 2.4: A successful oracle replay attack on \mathbf{A} .

adversary by **round 1** and the initial round by **round 0**. Suppose, at the end of **round 1**, the adversary forges $\hat{\sigma}'$ on some message \hat{m}' (with $R' = g^{r'}$, $c' = H(\hat{m}', R')$ and $y' = r' + \alpha c'$). Thus the challenger has secured two forgeries, $\hat{\sigma}$ and $\hat{\sigma}'$ with $y = r + \alpha c$ and $y' = r' + \alpha c'$. Let I' be the target H-index in **round 1**. Now, there are two possibilities with respect to I and I' : *i*) they could be different; or *ii*) they could match ($I' = I$). In the first case, \mathbf{B} cannot extract α as it possesses (only) two congruences in three unknowns. However, in the second case the adversary has made the same commitment (\hat{m}, R) , but received two, different, oracle outputs (see **Figure 2.4**). Thus, the simulator ends up with two forgeries $\{\hat{\sigma}, \hat{\sigma}'\}$ with $\{y = r + \alpha c, y' = r + \alpha c'\}$ —a system of two linear congruences in two unknowns—and it easily solves for α (by computing $(y - y')/(c - c')$). This completes our argument. \square

With a bit more effort, it is possible to cater to the signature queries as well. Thus, the above

³Note that a random oracle is distinguished by its outputs—the random function that represents the oracle. Therefore, we can conceive a “new” random oracle by changing the outputs from the “original” random oracle. However, in this particular case, the new random oracle has to *agree* with the original random oracle up to the target oracle query. Simply put, one has to use two different random functions that agree up to a particular point (the point being the target query). We elaborate on how to accomplish this in the next section.

⁴We will use the terms forking and (the process of launching) oracle replay attack interchangeably.

argument can be extended to the full model for PKS, *i.e.*, EU-CMA. We omit this discussion to avoid digression. However, we *are* interested in security of Schnorr signature against more “general” probabilistic, polynomial-time (PPT) adversaries. But, this transition requires some additional tools. We discuss this in the next section, along with the discourse on the oracle replay attack and its analysis (which we had omitted in the proof).

2.3 The Oracle Replay Attack

In the previous section, we demonstrated how the (oracle) replay attack was used in order to carefully manipulate the adversary and obtain two related signatures. However, we had refrained from divulging the details on the replay attack. It was used at a “black-box” level. In this section, our focus is on the replay attack itself. We discuss two approaches: the “Replicating” Technique and the “Rewinding” Technique, that are used to carry out replay attacks. In addition to this, we “lift” the (basic) security argument given in §2.2.3 to accommodate a more “general” PPT adversary. We end up with:

Theorem 1 (The Forking Lemma, Pointcheval-Stern [PS00]⁵). *Given an ϵ -adversary A against the Schnorr signature scheme in the EU-CMA model, it is possible to construct an algorithm B that solves the DLP with an advantage $\epsilon' \geq \epsilon^2/4q$, provided the hash function H is modelled as a random oracle with an upper bound of q queries. In other words,*

$$\text{DLP} \leq_{O(q)} \text{EU-CMA [Schnorr Signature]}.$$

The security argument is enabled, primarily, by the “Splitting” Lemma [PS00] which we discuss next.

2.3.1 The Splitting Lemma

Let \mathbb{X} and \mathbb{Y} be two finite sets (with some underlying probability distribution) with $|\mathbb{X}| = m$ and $|\mathbb{Y}| = n$. Also, let \mathbb{T} denote the “universal” set $\mathbb{X} \times \mathbb{Y}$. A pair $(x, y) \in \mathbb{X} \times \mathbb{Y}$ is deemed to be **good** if it satisfies a certain property. Let \mathbb{V} denote the set of all such **good** pairs. Suppose that at least

⁵It is possible to give a simpler, cleaner, security argument using the notion of “general forking”.

γ fraction of the pairs in $\mathbb{X} \times \mathbb{Y}$ are **good**, *i.e.*, $\mathbb{P}(\mathbb{V}) \geq \gamma$. We define a function $k : \mathbb{X} \mapsto \mathfrak{P}(\mathbb{Y})$ as

$$k(x) \stackrel{\text{def}}{=} \{y \in \mathbb{Y} \mid (x, y) \in \mathbb{V}\}.$$

k can be used to count, for a particular x , the number of elements $y \in \mathbb{Y}$ such that (x, y) forms a **good** pair. For some $\beta < \gamma$, we define the (sub)set of **better** pairs of \mathbb{V} as

$$\mathbb{V}^* \stackrel{\text{def}}{=} \{(x, y) \in \mathbb{V} \mid |k(x)| \geq (\gamma - \beta)n\}.$$

Thus, a **good** pair (x, y) is **better** if x forms **good** pairs with at least $(\gamma - \beta)$ fraction of the elements of \mathbb{Y} .

Lemma 1 (The Splitting Lemma, Pointcheval-Stern [PS00]⁶). *The following propositions hold on the aforementioned assumptions:*

1. $\mathbb{P}(\mathbb{V}^*) \geq \beta$, *i.e.*, at least β fraction of pairs in \mathbb{T} are **better**.
2. $\mathbb{P}((x, y') \in \mathbb{V} \mid (x, y) \in \mathbb{V}^*; y' \stackrel{\$}{\leftarrow} \mathbb{Y}) \geq (\gamma - \beta)$, *i.e.*, given a **better** pair (x, y) , the probability with which the pair (x, y') (with y' sampled randomly from \mathbb{Y}) is **good** is at least $(\gamma - \beta)$.
3. $\mathbb{P}(\mathbb{V}^* \mid \mathbb{V}) \geq \beta/\gamma$, *i.e.*, a **good** pair is **better** with a probability of at least β/γ .

Argument. We begin with **Proposition 2**. It is not difficult to see that it follows from the definition of **good** and **better** sets. Moving on, we establish **Proposition 1** through contradiction. Suppose $\mathbb{P}(\mathbb{V}) \geq \gamma$ and $\mathbb{P}(\mathbb{V}^*) < \beta$, then we have

$$\mathbb{P}(\mathbb{V}) = \mathbb{P}(\mathbb{V}^*) + \mathbb{P}(\mathbb{V} \setminus \mathbb{V}^*) < \beta + \mathbb{P}(\mathbb{V} \setminus \mathbb{V}^*). \quad (2.2)$$

Now, there can be at most m distinct x elements in $\mathbb{V} \setminus \mathbb{V}^*$. In addition, for each such $x \in \mathbb{V} \setminus \mathbb{V}^*$, $|k(x)| < (\gamma - \beta)n$ (by definition). Therefore (2.2) yields,

$$\mathbb{P}(\mathbb{V}) < \beta + \frac{|k(x)| \cdot m}{mn} < \beta + \frac{(\gamma - \beta)n \cdot m}{mn} = \gamma.$$

This contradicts our assumption that at least γ fraction of elements in \mathbb{T} are **good** ($\mathbb{P}(\mathbb{V}) \geq \gamma$), thus, establishing **Proposition 1**. Finally, we establish the **Proposition 3** using Bayes' theorem

⁶Also refer to [Kia07].

as shown below.

$$\begin{aligned} \mathbb{P}(\mathbb{V}^* \mid \mathbb{V}) &= \mathbb{P}(\mathbb{V} \mid \mathbb{V}^*)\mathbb{P}(\mathbb{V}^*)/\mathbb{P}(\mathbb{V}) \text{ (by Bayes' theorem)} \\ &\geq \beta/\gamma \text{ (using **Proposition 1**)} \end{aligned} \tag{2.3}$$

Note that the second step ($\mathbb{P}(\mathbb{V}^* \mid \mathbb{V}) = 1$) follows from the fact that a **better** pair is always **good**. This concludes the argument. \square

The Splitting Lemma: An alternative formulation. Koblitz-Menezes [KM07] explained the Splitting Lemma using a different perspective from that in **Lemma 1**. This comes in handy in certain scenarios. The sets \mathbb{X} , \mathbb{Y} (hence \mathbb{T}) and \mathbb{V} , as well as the function k , are defined as in the previous formulation. For some $\beta < \gamma$, we define the (sub)set of **better** elements of \mathbb{X} as

$$\mathbb{X}^* = \{x \in \mathbb{X} \mid |k(x)| \geq (\gamma - \beta)n\}.$$

Thus, an element $x \in \mathbb{X}$ is **better** if it forms **good** pairs with at least $(\gamma - \beta)$ fraction of the elements of \mathbb{Y} .

Lemma 2 (The Splitting Lemma, Koblitz-Menezes [KM07]). *The following propositions hold on the aforementioned assumptions:*

1. $\mathbb{P}(\mathbb{X}^*) \geq \beta$, *i.e.*, at least β fraction of elements in \mathbb{X} are **better**.
2. $\mathbb{P}((x, y) \in \mathbb{V} \mid x \in \mathbb{X}^* \wedge y' \stackrel{\$}{\leftarrow} \mathbb{Y}) \geq \gamma - \beta$, *i.e.*, given a **better** element x , the probability with which the pair (x, y) (with y sampled randomly from \mathbb{Y}) is **good** is at least $\gamma - \beta$.

Argument. As in **Lemma 1**, **Proposition 2** follows by definition. **Proposition 1** is, again, established through contradiction. Provided that $\mathbb{P}(\mathbb{V}) \geq \gamma$ and $\mathbb{P}(\mathbb{X}^*) < \beta$, let's count the number of **good** pairs in \mathbb{T} . For any **good** pair (x, y) , the coordinate x comes from either \mathbb{X}^* or $\mathbb{X} \setminus \mathbb{X}^*$. Keeping this in mind, we define the two sets \mathbb{V}_1 and \mathbb{V}_2 as follows.

$$\mathbb{V}_1 = \{(x, y) \mid x \in \mathbb{X}^* \wedge y \in k(x)\} \quad \mathbb{V}_2 = \{(x, y) \mid x \in \mathbb{X} \setminus \mathbb{X}^* \wedge y \in k(x)\}$$

It follows from the definition that

$$\mathbb{P}(\mathbb{V}) = \mathbb{P}(\mathbb{V}_1) + \mathbb{P}(\mathbb{V}_2) < \beta \cdot 1 + 1 \cdot (\gamma - \beta) = \beta. \tag{2.4}$$

(2.4) contradicts our initial assumption that $\mathbb{P}(\mathbb{V}) \geq \gamma$ and that concludes the argument. \square

Remark 4. We refer to the Splitting Lemma by (γ, β) -Splitting Lemma. For a particular value of γ , the optimal value of β depends on the way the replay attack is carried out. This is demonstrated in §2.3.2.3.

The Splitting Lemma in the context of signature schemes. Recall the basic security argument for Schnorr signature scheme from §2.2.3. We had assumed a perfect adversary—an adversary capable of forging every time. In other words, the adversary is successful irrespective of the input (which includes the output of the random oracle queries it made). Next, let’s consider a more general adversary with a success probability of ϵ . The **good** set can be regarded as (the set of) those particular inputs for which the adversary is able to forge a signature successfully⁷. Thus for an adversary with success probability ϵ , at least ϵ fraction of the inputs results in a valid forgery; in other words, $\gamma = \epsilon$.⁸ As for the perfect adversary, we can consider any input to be **good**, *i.e.*, $\gamma = \epsilon = 1$. Although the discussion pertains to Schnorr signature scheme, it applies to other signature schemes as well.

2.3.2 Launching the Oracle Replay Attack

Let’s go back to the (basic) security argument of the Schnorr signature scheme, especially, the way the replay attack was carried out. The simulator “turned back the clock” to a “critical” juncture and then continued with a new simulation with a different “setting”. In this section, we discuss two approaches by which one can actually carry out the replay attack: the *Replicating Technique* and the *Rewinding Technique*. This would explain how to turn back the clock, what comprises a critical point and what we mean by a different setting. Finally, we analyse the two techniques using the Splitting Lemma. We stick to Schnorr signature and will be constantly referring back to the basic security argument from §2.2.3.

⁷*Caveat:* Recall our modelling of the adversary. We had assumed the adversary to be a probabilistic (polynomial-time) Turing machine. Now, consider two simulations of the adversary on a specific input. Its behaviour, during these two simulations, depends on the internal coin tosses. Therefore, it may not necessarily be successful during both the simulations. This means that our classification of an input as **good** is a bit ambiguous. However, it serves the purpose of lending an intuitive explanation. A sound definition of **good** input has to take into account these internal coin tosses as well. We adopt this approach from the next section.

⁸We may, at times, use a more refined definition of **good**. See §2.3.2.3 for an example of such a definition.

Replicating and rewinding: an analogy. Before delving into the technical details, we would like to give an intuitive overview of the two approaches. We adopt popular science-fiction subjects for analogy. Intuitively, the Rewinding Technique can be likened to the concept of *time travel*. The adversary (modelled as a probabilistic Turing machine) is executed once. At the end of this round, we have sufficient knowledge about the critical point. The “rewinding” attribute of a Turing machine (which enables us to trace its actions back to a previous state of progress) is, then, used to restore the adversary to this critical point. The adversary is then simulated again, but, with a setting which is different from the initial round of simulation. The clock is thus, literally, turned back on the adversary.

On the other hand, the Replicating Technique can be likened to the concept of *parallel universes*⁹. At some (random) point during its simulation, the adversary is paused and we produce another, identical, copy of it¹⁰. Then, these two instances are simulated independently, but, each with a different setting. Finally, we hope the point at which we diverged turns out to be the critical point. To sum it all up, we have two adversaries, identical up to a certain point, but with behaviour differing from that point, as if in parallel universes¹¹. We now move on the technical details.

2.3.2.1 The Replicating Technique

Let’s return to the security argument of the Schnorr signature scheme. The critical point in the simulation is precisely the juncture when the adversary A makes the target oracle query Q_I . The simulator starts by making a guess i of this target index I . Next, B continues with the simulation until the adversary A makes the i^{th} query. At this juncture, B produces a replica of A which we denote by A' . The simulator, then, proceeds with the simulation of these two instances of the adversary independently—*forking, by replicating the adversary at Q_I* . However, in order to induce a different setting in the two rounds (from Q_i), B responds to the H -oracle queries with different outputs: A ’s queries are responded to with $\{c_i, \dots, c_q\}$, whereas, those of A' with $\{c'_i, \dots, c'_q\}$, as shown in **Figure 2.5** below.

At the end of the two rounds, the simulator is in possession of two forgeries, one each due to A and A' . Let I and I' be the target indices for the two rounds. The forking is successful if the simulator’s guess of the target index, *i.e.* i , turns out to be correct for both the rounds, *i.e.*,

⁹It is also referred to as a *multiverse*. See the introductory chapter of [Teg03] for a intuitive explanation of parallel universes.

¹⁰This is possible as the adversary is modelled as a Turing machine.

¹¹Interestingly, the concepts of time travel and parallel universes are closely related.

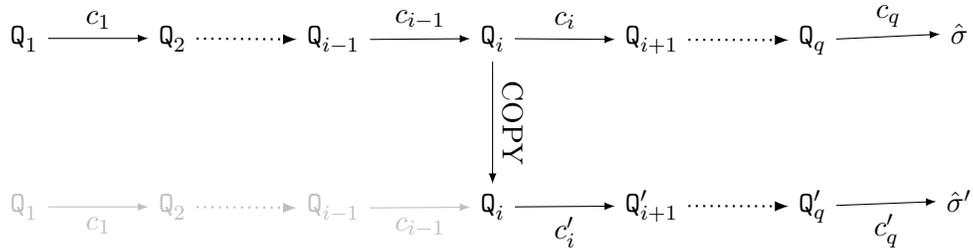


Figure 2.5: Oracle replay attack carried out by replicating.

$I' = I = i$.

2.3.2.2 The Rewinding Technique

The approach adopted to launch the replay attack using the Rewinding Technique is fundamentally different from that in the Replicating Technique. Nevertheless, it still involves two simulations of the adversary. But in the Rewinding Technique, contrary to the Replicating Technique, the challenger first completes a full simulation of the adversary as shown in round 0 in the figure below.

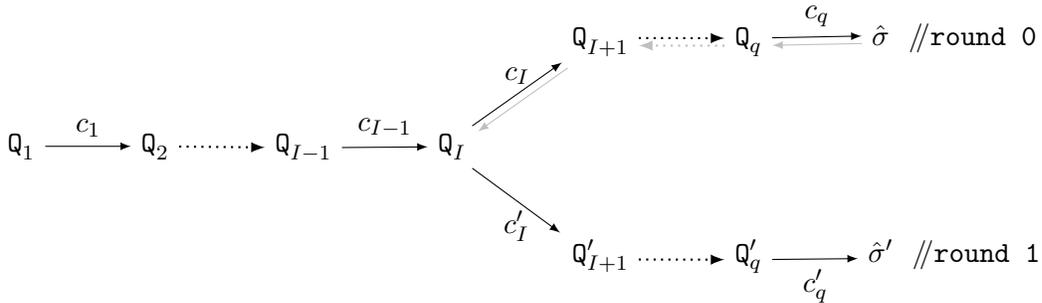


Figure 2.6: Oracle replay attack carried out by rewinding.

At the end of the round, the simulator is in possession of a forgery $\hat{\sigma} = (y, R)$ on some message \hat{m} . The forgery, in turn, contains sufficient information about the critical point for that particular round—it is, precisely, the juncture when the adversary made the target oracle query $H(\hat{m}, R)$, which is denoted in **Figure 2.6** by Q_I . The simulator, next, rewinds¹² the adversary A (which is modelled

¹²We discuss a rather naïve approach that could be used to rewind the adversary, which is modelled as a probabilistic Turing machine. An elegant, more abstract, approach which involves the alternative modelling the adversary (as a deterministic Turing machine) is discussed later.

The state of progress of a Turing machine is (completely) determined by its current state (from the state register), the contents of the tape and the position of the head on the tape. The basic idea involved in rewinding is as follows. The simulator explicitly saves the state of progress of the adversary every time it queries the random oracle. To be precise, it saves the current state (from the state register), the contents of the tape and the position of the head on a separate, external, tape. The simulator, thus, may have to save at most q states of progress. At the end of the round,

as a Turing machine) to this critical point and then continues with the simulation of the adversary, but with a different setting—*forking, by rewinding the adversary to Q_I* . Let’s call this **round 1**. A different setting is induced in **round 1**, as in the Replicating Technique, by using a different set of outputs in response to the random oracle queries— $\{c'_I, \dots, c'_q\}$ compared to $\{c_I, \dots, c_q\}$ in **round 0**. At the end of **round 1**, the simulator secures another forgery $\hat{\sigma}'$. Let I' be the target index for this particular round. The forking is successful if the target index for **round 1** matches the one during **round 0**, *i.e.*, $I' = I$.

2.3.2.3 Analysis

Forking, when carried out using the Rewinding Technique, is successful if the target indices during the two rounds of simulation of the adversary match. In the Replicating Technique, on the other hand, there is an additional requirement that the simulator’s guess of the target index has to be correct. This leads to additional degradation. Thus, at first glance, the Rewinding Technique seems to be the more efficient of the two approaches. The analysis concurs with this observation. The Splitting Lemma, discussed in §2.3.1, plays a central role in establishing the lower bound. In addition, we also use the following result.

Lemma 3 (Hölder’s inequality¹³). *Let $q \in \mathbb{Z}^+$, $1 \leq n < \infty$ and $x_1, \dots, x_q \geq 0$ be real numbers.*

Then

$$\sum_{k=1}^q x_k^n \geq \frac{1}{q^{n-1}} \left(\sum_{k=1}^q x_k \right)^n.$$

Defining the notions of “good”. Consider a general adversary A with a non-negligible advantage of ϵ . Let \mathbb{T} denote the set of all random tapes participating in a single round of simulation of the adversary—the universe of tapes. This includes the internal coins of the adversary as well as the randomness from the random functions associated with the random oracle and the signature oracle. Provided that the adversary has a non-negligible advantage of ϵ , at least ϵ fraction of the tapes in \mathbb{T} lead to a successful forgery. These constitute the **good** tapes in \mathbb{T} .

the simulator has sufficient information about the critical point. Besides, the state of progress of the adversary at this critical point—the critical state of progress—is highly likely to be one of the saved states of progress. Therefore, in order to rewind the adversary, the simulator has to just look up for this critical state of progress and restore the adversary to it.

¹³Although, the result is a corollary to a more general Hölder’s inequality (see [BPW12, Lemma C.3]), another way to proving the bound is by viewing it an optimisation problem. Let $f(x_1, \dots, x_q) := \sum_{k=1}^q x_k^n$ be the objective function under the set of constraints: *i)* $\sum_{k=1}^q x_k = x$, and *ii)* $(0 \leq x_k \leq 1)$ for $k \in \{1, \dots, q\}$. The function f attains a minima of x^n/q^{n-1} at the point $(x/q, \dots, x/q)$ which completes the proof.

Recall, from §2.2.3, that there must exist a target H-index I for the adversary to have produced the forgery with a non-negligible advantage. On this assumption, we define a more refined notion of **good**—termed $\text{good}_{(i)}$ —for the tapes in \mathbb{T} . For an arbitrary index $i \in \{1, \dots, q\}$, a tape in \mathbb{T} is deemed to be $\text{good}_{(i)}$ if it leads to the adversary forging successfully *and* with a target H-index of i . It follows that a tape T is **good** is $\text{good}_{(i)}$ for some i . Let's assume that ϵ_i fraction of the tapes in \mathbb{T} are $\text{good}_{(i)}$. Thus, we have

$$\sum_{i=1}^q \epsilon_i = \sum_{i=1}^q \mathbb{P}(I = i) = \mathbb{P}(I > 0) = \epsilon. \quad (2.5)$$

Replay attack, as already stated, involves simulating the adversary on two tapes that agree up to a certain juncture—the target oracle query. In order to accommodate this, we split the universe \mathbb{T} into the two sets $\mathbb{T}_{(i-)}$ and $\mathbb{T}_{(i+)}$ using the Cartesian product. Here, $\mathbb{T}_{(i-)}$ denotes the (sub)universe random tapes involved in the simulation before the adversary makes the query Q_i , whereas $\mathbb{T}_{(i+)}$, those after the query Q_i . We assume that \mathbb{T} is bijective to $\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ through a “join” function (denoted by \parallel) and a “split” function.¹⁴ The notion of $\text{good}_{(i)}$, which was originally defined for the set \mathbb{T} , can be easily adapted for the set $\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ using the bijection: a tape $(T_{i-}, T_{i+}) \in \mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ is deemed to be $\text{good}_{(i)}$ if its counterpart in \mathbb{T} (*i.e.*, $T_{i-} \parallel T_{i+}$) is $\text{good}_{(i)}$. We denote the (sub)set of all such tapes in $\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ by \mathbb{V}_i .

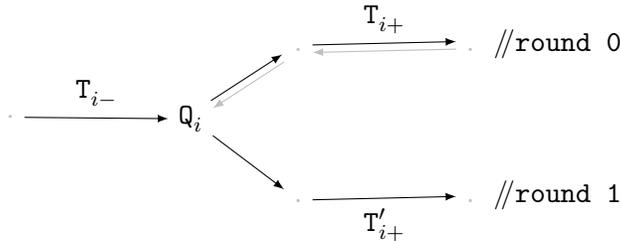


Figure 2.7: Tapes involved in the replay attack.

Applying the Splitting Lemma to the Replicating Technique. The Kobitz-Menezes version of the Splitting Lemma (**Lemma 2**), seems to be the more natural choice for analysing replay attack launched using the Replicating Technique. Let's assume that an arbitrary β_i fraction of the

¹⁴The “join” function f_i concatenates the two constituent tape segments from the set $\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ into a single tape in set \mathbb{T} , *i.e.*, $f_i(T_{i-}, T_{i+}) := T_{i-} \parallel T_{i+}$. Provided that f_i is bijective, the inverse of f_i would be the function which “splits” a tape T into the two constituent tape segments T_{i-} and T_{i+} depending upon the random oracle query Q_i , *i.e.*, $f_i^{-1}(T) := (T_{i-}, T_{i+})$. Note that the definition of f_i does not depend of i and hence we consider a single join function f .

tape segments in $\mathbb{T}_{(i-)}$ are $\text{better}_{(i)}$ in the sense that is described in **Lemma 2**. We denote this set by $\mathbb{T}_{(i-)}^*$. The exact value of β_i is determined later. Let $(\mathbf{T}_{i-}, \mathbf{T}_{i+})$ and $(\mathbf{T}_{i-}, \mathbf{T}'_{i+})$ be the tapes involved in the two rounds of simulation in a particular instance of the replay attack, with

$$\mathbf{T}_{i-} \stackrel{\$}{\leftarrow} \mathbb{T}_{(i-)} \quad \text{and} \quad (\mathbf{T}_{i+}, \mathbf{T}'_{i+}) \stackrel{\$}{\leftarrow} \mathbb{T}_{(i+)}. \quad (2.6)$$

Recall that i , here, is the simulator's guess of the target index. The forking is successful if both $(\mathbf{T}_{i-}, \mathbf{T}_{i+})$ and $(\mathbf{T}_{i-}, \mathbf{T}'_{i+})$ turn out to be $\text{good}_{(I)}$ pairs and the guess of the target index turns out to be correct, *i.e.*, $I = i$. Let's denote the probability of this event by ϵ'_i . Thus, the probability with which the forking is successful (with the tapes sampled as shown in (2.6)) is given by:

$$\begin{aligned} \epsilon'_i &= \mathbb{P}((\mathbf{T}_{i-}, \mathbf{T}_{i+}) \in \mathbb{V}_I \wedge (\mathbf{T}_{i-}, \mathbf{T}'_{i+}) \in \mathbb{V}_I \wedge I = i) \\ &= \mathbb{P}((\mathbf{T}_{i-}, \mathbf{T}_{i+}) \in \mathbb{V}_i \wedge (\mathbf{T}_{i-}, \mathbf{T}'_{i+}) \in \mathbb{V}_i \mid I = i) \cdot \mathbb{P}(I = i). \end{aligned} \quad (2.7)$$

As the index i is chosen uniformly at random from $\{1, \dots, q\}$, the probability that it matches I is at least $1/q$. As per our definition, the tape $(\mathbf{T}_{i-}, \mathbf{T}_{i+})$ is $\text{good}_{(i)}$ if \mathbf{T}_{i-} is $\text{better}_{(i)}$ and \mathbf{T}_{i+} belongs to $\mathbf{k}(\mathbf{T}_{i-})$. The same applies to the tape $(\mathbf{T}_{i-}, \mathbf{T}'_{i+})$. Thus, (2.7) can be rewritten as follows:

$$\begin{aligned} \epsilon'_i &= \mathbb{P}(\mathbf{T}_{i-} \in \mathbb{T}_{(i-)}^* \wedge \mathbf{T}_{i+} \in \mathbf{k}(\mathbf{T}_{i-}) \wedge \mathbf{T}'_{i+} \in \mathbf{k}(\mathbf{T}_{i-}) \mid I = i) \cdot 1/q \\ &= \mathbb{P}(\mathbf{T}_{i+} \in \mathbf{k}(\mathbf{T}_{i-}) \wedge \mathbf{T}'_{i+} \in \mathbf{k}(\mathbf{T}_{i-}) \mid \mathbf{T}_{i-} \in \mathbb{T}_{(i-)}^*) \cdot \mathbb{P}(\mathbf{T}_{i-} \in \mathbb{T}_{(i-)}^*) \cdot 1/q \\ &= \mathbb{P}(\mathbf{T}_{i+} \in \mathbf{k}(\mathbf{T}_{i-}) \mid \mathbf{T}_{i-} \in \mathbb{T}_{(i-)}^*) \cdot \mathbb{P}(\mathbf{T}'_{i+} \in \mathbf{k}(\mathbf{T}_{i-}) \mid \mathbf{T}_{i-} \in \mathbb{T}_{(i-)}^*) \cdot \beta_i \cdot 1/q \\ &\quad (\text{using **Proposition 1** of **Lemma 2**}) \\ &\geq (\epsilon_i - \beta_i)^2 \beta_i / q (\text{using **Proposition 2** of **Lemma 2**}) \end{aligned} \quad (2.8)$$

The above expression attains a maxima¹⁵ of $4\epsilon_i^3/27q$ at the point $\beta_i = \epsilon_i/3$. Thus, the probability with which the forking is successful for any i is given by

$$\epsilon' = \sum_{i=1}^q \epsilon'_i = \sum_{i=1}^q \frac{4}{27q} \epsilon_i^3 = \frac{4}{27q} \sum_{i=1}^q \epsilon_i^3. \quad (2.9)$$

¹⁵The objective function is $f(\beta_i) = (\epsilon_i - \beta_i)^2 \beta_i / q$ under the constraint $\beta_i < \epsilon_i$. The function attains maximum value at the point $\beta_i = \epsilon_i/3$.

Under the constraint given in (2.5), (2.9) attains a minima¹⁶ at the point $\epsilon_i = \epsilon/q$ for all $1 \leq i \leq q$ and thus, we get the lower bound

$$\epsilon' \geq \frac{4}{27q} \sum_{i=1}^q (\epsilon/q)^3 = \frac{4}{27} (\epsilon/q)^3. \quad (2.10)$$

Applying the Splitting Lemma to the Rewinding technique. For analysing the replay attack launched using the Rewinding Technique, we use **Lemma 1**. Let's presume that an arbitrary β_i fraction of the tape pairs in \mathbb{T} are $\text{better}_{(i)}$ in the sense described in **Lemma 1**. We denote this set by \mathbb{V}_i^* . The exact value of β_i is, again, determined later. Let \mathbb{T} (sampled at random from \mathbb{T}) be the tape involved in the simulation of the adversary in **round 0**. If the adversary forges successfully, the simulator can determine the target index I . We split \mathbb{T} into the two tapes \mathbb{T}_{I-} and \mathbb{T}_{I+} depending on this index I . Here, \mathbb{T}_{I-} is the tape involved in the simulation before the adversary makes the query \mathbb{Q}_I , whereas \mathbb{T}_{I+} , that after the query \mathbb{Q}_I . Thus, $(\mathbb{T}_{I-}, \mathbb{T}_{I+})$ is a $\text{good}_{(I)}$ pair. The simulator, next, rewinds the adversary to the point \mathbb{Q}_I and simulates it on a fresh tape \mathbb{T}'_{I+} (sampled at random from $\mathbb{T}_{(I+)}$). The replay attack is successful if $(\mathbb{T}_{I-}, \mathbb{T}'_{I+})$ also turns out to be a $\text{good}_{(I)}$ pair. However, this is a necessary condition. For ease of analysis, we use the sufficient condition that $(\mathbb{T}_{I-}, \mathbb{T}_{I+})$ be better and $(\mathbb{T}_{I-}, \mathbb{T}'_{I+})$ be good . Thus, the probability with which the Rewinding Technique is successful (with the tapes sampled as described above) is given by:

$$\begin{aligned} \epsilon'_I &= \mathbb{P}((\mathbb{T}_{I-}, \mathbb{T}_{I+}) \in \mathbb{V}_I^* \wedge (\mathbb{T}_{I-}, \mathbb{T}'_{I+}) \in \mathbb{V}_I) \\ &= \mathbb{P}((\mathbb{T}_{I-}, \mathbb{T}'_{I+}) \in \mathbb{V}_I \mid (\mathbb{T}_{I-}, \mathbb{T}_{I+}) \in \mathbb{V}_I^*) \cdot \mathbb{P}((\mathbb{T}_{I-}, \mathbb{T}_{I+}) \in \mathbb{V}_I^*) \\ &= (\epsilon_I - \beta_I)\beta_I \text{ (using **Proposition 1** and **2** of **Lemma 1**)} \end{aligned} \quad (2.11)$$

¹⁶The objective function is

$$f(\epsilon_1, \dots, \epsilon_q) = \frac{4}{27q} \sum_{i=1}^q \epsilon_i^3$$

under the set of constraints $\sum_{i=1}^q \epsilon_i = \epsilon$ and $\epsilon_i \in [0, 1]$ for each $i \in \{1, \dots, q\}$. The function attains minimum value at the point $(\epsilon/q, \dots, \epsilon/q)$. This can also be established using Hölder's inequality (**Lemma 3**).

The above expression attains a maxima¹⁷ of $\epsilon_I^2/4$ at the point $\beta_I = \epsilon_I/2$. Thus, the probability that the oracle replay attack is successful for any I is given by

$$\begin{aligned}\epsilon' &= \sum_{I=1}^q \epsilon'_I \geq \sum_{I=1}^q \frac{\epsilon_I^2}{4} \\ &\geq \frac{\epsilon^2}{4q} \text{ (using Hölder's inequality)}\end{aligned}\tag{2.12}$$

That concludes the analysis. It is evident that the Rewinding Technique is more *efficient* than the Replicating Technique—the security degradation incurred (by the former) is just $O(q)$ compared to $O(q^3)$ (for the latter). Thus, it is the Rewinding Technique that yields **Theorem 1**. Henceforth, by forking we *always* refer (unless explicitly mentioned) to replay attack carried out by rewinding the adversary.

We wrap up the discussion by introducing the notion of “characteristic” expressions and, also, with some remarks on the analyses. In the next section, we describe how the intricacy of forking (and also its analysis) can be separated out from the simulation of the adversary through the notion of general forking. We also describe the nested oracle replay attack which, in simple terms, involves two random oracles and (as the name suggests) multiple, nested, forkings.

The *characteristic* expression. The replay attack launched using Replicating Technique is captured, essentially, by the expression

$$\epsilon' \geq \frac{4}{27q} \sum_{i=1}^q \epsilon_i^3\tag{2.13}$$

under the set of constraints: *i*) $\sum_{i=1}^q \epsilon_i = \epsilon$, and *ii*) $(0 \leq \epsilon_i \leq 1)_{i \in \{1, \dots, q\}}$. We say that (2.13) is the *characteristic* expression for the Replicating Technique. Similarly, the characteristic expression for the Rewinding Technique is

$$\epsilon' \geq \frac{1}{4} \sum_{i=1}^q \epsilon_i^2\tag{2.14}$$

under the same set of constraints. This abstraction comes in handy (in coming chapters) for analysing nested replay attacks.

Remark 5. After optimisation, we arrived at two different values of β_i for the Splitting Lemma:

¹⁷The objective function is $f(\beta_I) = (\epsilon_I - \beta_I)\beta_I$ under the constraint $\beta_I < \epsilon_I$. The function attains maximum value at the point $\beta_I = \epsilon_I/2$.

for the Replicating Technique, we used an $(\epsilon_i, \epsilon_i/3)$ -Splitting Lemma, whereas, for the Rewinding Technique, an $(\epsilon_i, \epsilon_i/2)$ -Splitting Lemma.

Remark 6. We denote the universe of tapes for two rounds of simulations, associated with a forking of the adversary, by $\mathbb{T}_{(1)}$ ('1' in the subscript denotes one forking), *i.e.*,

$$\begin{aligned} \mathbb{T}_{(1)} &= \bigcup_{i=1}^q \mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}^2 \\ &= \bigcup_{i=1}^q \mathbb{T}_{(i-)} \times \mathbb{T}_{(1,i+)} \text{ (using shorthand notation)} \end{aligned}$$

The introduction of $\mathbb{T}_{(1)}$ allows us to define the ‘‘higher’’ notion of $\mathbf{good}_{(1)}$ based on the notions of \mathbf{good} and $\mathbf{good}_{(i)}$. The triplet of tape segments $(\mathbb{T}_{I-}, (\mathbb{T}_{I+}, \mathbb{T}'_{I+})) \in \mathbb{T}_{(1)}$ is considered to be $\mathbf{good}_{(1)}$ if both the tapes $(\mathbb{T}_{I-}, \mathbb{T}_{I+})$ and $(\mathbb{T}_{I-}, \mathbb{T}'_{I+})$ are $\mathbf{good}_{(I)}$, *i.e.*, if the forking is successful. From the analyses, we had concluded the forking is successful with probability at least ϵ' . This result can be interpreted in terms of $\mathbf{good}_{(1)}$ as follows: a randomly sampled triplet from $\mathbb{T}_{(1)}$ is $\mathbf{good}_{(1)}$ with probability at least ϵ' . These notions come into play in the analysis of modified GG-IBS in §4.3.2.

2.4 General Forking

In the previous section, we saw how the machinery of oracle replay attack was used in the security argument of Schnorr signature scheme. However, we also saw that the probability analysis of the security argument, which culminated with the Forking Lemma (**Lemma 1**), was quite *involved*. Bellare and Neven [BN06] observed that

‘‘the Forking Lemma is something purely probabilistic, not about signatures’’

and proposed a more abstract version called the General Forking (GF) Lemma. The motivation is to demarcate the probability analysis of the rewinding from that of the actual simulation, allowing for more modular security arguments. The concept of general forking is formulated in terms of randomised algorithms and its outputs, leaving out the notions of signature scheme as well as random oracles altogether. Thus, the alternative modelling of the adversary (as a deterministic Turing machine with the internal coins being passed explicitly as input.) which we had discussed in §1.2, comes into play. The claimed advantage is to allow for more modular and easy to verify proof of cryptographic schemes that apply the notion of forking in their security argument.

The modularity of the GF Lemma allows one to abstract out the probability analysis of the rewinding process from the actual simulation in the security argument. The gap between the abstract and the concrete is, then, bridged using the so-called “wrapper” algorithm. While the GF Algorithm takes care of the replay attack, it is the wrapper that handles the simulation of the protocol environment to the actual adversary. The reduction involves invoking the General Forking Algorithm (on the associated wrapper) and utilising its outputs to solve the underlying hard problem. So the design of the wrapper is central to any security argument involving GF Algorithm. In fact, the design depends on the actual protocol and the security model used—see, e.g., [BN06, BPW12] for the concrete design of the wrappers in their respective contexts.

General Forking Algorithm. Fix $q \in \mathbb{Z}^+$ and a set \mathbb{S} such that $|\mathbb{S}| \geq 2$. Let W be a randomised algorithm that on input a string x and elements $s_1, \dots, s_q \in \mathbb{S}$ returns a pair (I, σ) consisting of an integer $0 \leq I \leq q$ and a string σ . The forking algorithm F_W associated to W is defined as **Algorithm 1** below.

Algorithm 1 $F_W(x)$

Pick coins ρ for W at random

$\{s_1^0, \dots, s_q^0\} \xleftarrow{\text{U}} \mathbb{S}; (I_0, \sigma_0) \leftarrow W(x, s_1^0, \dots, s_q^0; \rho)$ //round 0

if $(I_0 = 0)$ **then** $(0, \perp, \perp)$

$\{s_{I_0}^1, \dots, s_q^1\} \xleftarrow{\text{U}} \mathbb{S}; (I_1, \sigma_1) \leftarrow W(x, s_1^0, \dots, s_{I_0-1}^0, s_{I_0}^1, \dots, s_q^1; \rho)$ //round 1

if $(I_1 = I_0 \wedge s_{I_0}^1 \neq s_{I_0}^0)$ **then** $(1, \sigma_0, \sigma_1)$

else $(0, \perp, \perp)$

Role of the wrapper. Let’s take a *simplistic* look¹⁸ at how the GF Algorithm, together with the wrapper W , is used to launch the oracle replay attack. The input to W comprises of some external randomness (s_1, \dots, s_q) and the internal coins (ρ) for the adversary, whereas the output, an index I . Consider the first invocation of W (on $s_1, \dots, s_q; \rho$) within the GF Algorithm: W simulates the protocol environment to the actual adversary A having access to ρ . W responds to the random oracle queries of A by using s_1, \dots, s_q . At the end of the simulation, W outputs an index I that

¹⁸For the time being, we do not consider the input string x or the side-output σ .

refers to the target query¹⁹.

Next, the GF Algorithm invokes W on an input $(s_1, \dots, s_{I-1}, s'_I, \dots, s'_q; \rho)$ that is *related* to the first invocation. The behaviour of A remains identical to the first round of simulation, right up to the I^{th} random oracle query, at which point it diverges (assuming $s'_I \neq s_I$). This is tantamount to forking A at the index I . The forking is successful, if the target index for the second round of simulation is the same as that for the first, *i.e.*, $I' = I$. The probability of this event is governed by the General Forking Lemma given below. One can clearly see how the wrapper acts as an intermediary between the abstract GF Algorithm and the adversary in the concrete setting of the reduction.

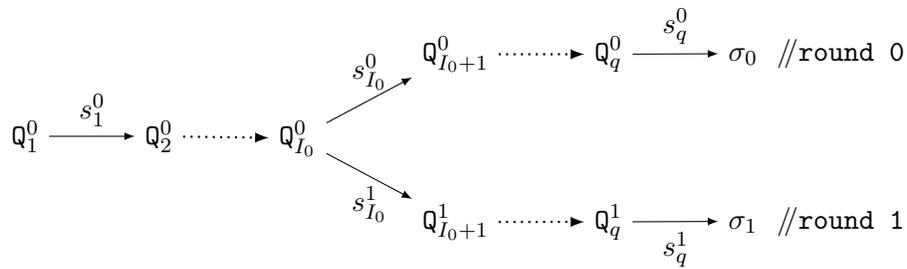


Figure 2.8: A successful forking of the wrapper W by F_W .

The cost of forking. The GF Lemma gives us a lower bound on the probability of success of the forking algorithm in terms of the success probability of the associated wrapper (and hence, the underlying adversary). Roughly speaking, the cost of forking can be measured in terms of the degradation incurred in the forking process. Let q denote the upper bound on the number of random oracle queries, then the cost of general forking is roughly $O(q)$.

Lemma 4 (General Forking Lemma [BN06]). *Let G_I be a randomised algorithm that takes no input and returns a string. Let*

$$gfrk := \mathbb{P}((op = 1) : x \stackrel{\$}{\leftarrow} G_I; (op, \sigma, \sigma) \stackrel{\$}{\leftarrow} F_W(x)) \quad \text{and}$$

$$acc := \mathbb{P}(I \geq 1 : x \stackrel{\$}{\leftarrow} G_I; \{s_1, \dots, s_q\} \stackrel{U}{\leftarrow} \mathbb{S}; (I, \sigma) \stackrel{\$}{\leftarrow} W(x, s_1, \dots, s_q)),$$

¹⁹Recall that, the random oracle query that is used by A to produce its desired output is termed the target query and the index of this query is the target index.

then

$$gfrk \geq acc \cdot \left(\frac{acc}{q} - \frac{1}{|\mathbb{S}|} \right). \quad (2.15)$$

Remark 7. The analysis of the GF Algorithm is quite different from that of Rewinding Technique. It is done through a random variable that models W as a randomised algorithm. However, the resulting characteristic expression is the *same* as that for the Rewinding Technique given in (2.14). Thus, the characteristic expression *links* the two approaches.

2.5 Nested Oracle Replay Attacks and Multiple Forking

So far, we have seen replay attacks involving *one* random oracle and a *single* forking of the adversary—an *elementary* oracle replay attack. Boldyreva *et al.* [BPW12] generalised the concept of forking, further, leading to the Multiple-Forking (MF) Lemma. The immediate motivation behind this new abstraction was to argue the security of a proxy signature scheme that uses more than one hash functions. The hash functions are modelled as random oracles and the MF Algorithm allows one to mount the nested replay attacks by rewinding the adversary several times on related inputs. In particular, a nested replay attack involves two random oracles and multiple forkings on the two oracles. Hence, unlike elementary replay attack, there are *two* critical points for a round.

A (successful) nested replay attack with three forkings is illustrated in **Figure 2.9**. The adversary is forked at the critical points Q_i^0 , Q_j^0 and Q_i^2 (in that order²⁰) and the forking is successful if the target indices (of the two random oracles involved) in all the four rounds is *same* (see (2.18) for the exact condition). The more abstract notion of multiple forking still retains the modularity advantage of general forking and has been applied in several other security arguments [GG09, CMW12, CKK13] in a more-or-less black-box fashion.

We begin with a Pointcheval-Stern-style analysis of the nested replay attack (given in **Figure 2.9**) and then describe the original MF Algorithm [BPW12] (but, with some notational changes).

2.5.1 Analysis

The nested replay attack given in **Figure 2.9** cannot be direct analysed using the Splitting Lemma (at least, not at the moment: see §4.3.2). We need to extend the Splitting Lemma (**Lemma 1**)

²⁰Further forkings can be carried out at Q_j^0 and Q_i^4 (in that order) and so forth.

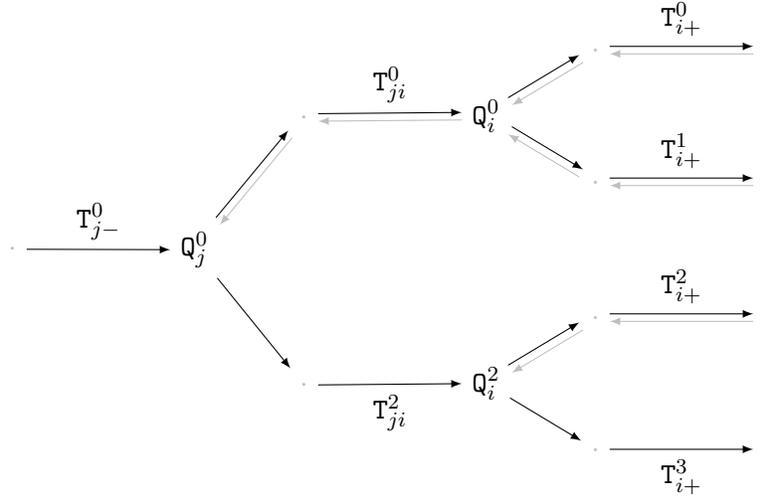


Figure 2.9: Nested replay attack with three forkings.

to suit the purpose.

2.5.1.1 The Extended Splitting Lemma

Let \mathbb{W} , \mathbb{X} and \mathbb{Y} be finite sets (with some underlying probability distribution) with $|\mathbb{W}| = \ell$, $|\mathbb{X}| = m$ and $|\mathbb{Y}| = n$. Also, let \mathbb{T} denote the “universal” set $\mathbb{W} \times \mathbb{X} \times \mathbb{Y}$. A triple $(w, x, y) \in \mathbb{W} \times \mathbb{X} \times \mathbb{Y}$ is deemed to be **good** if it satisfies a certain property. Let \mathbb{V} denote the set of all such **good** triples. Suppose that at least γ fraction of the triples in $\mathbb{W} \times \mathbb{X} \times \mathbb{Y}$ are **good**, *i.e.*, $\mathbb{P}(\mathbb{V}) \geq \gamma$. We define a function $k : \mathbb{W} \times \mathbb{X} \mapsto \mathfrak{P}(\mathbb{Y})$ as

$$k(w, x) \stackrel{\text{def}}{=} \{y \in \mathbb{Y} \mid (w, x, y) \in \mathbb{V}\}.$$

k can be used to count, for a particular (w, x) , the number of elements $y \in \mathbb{Y}$ such that (w, x, y) forms a **good** triple. For some $\beta < \gamma$, we define the (sub)set of **better** triples of \mathbb{V} as

$$\mathbb{V}^* \stackrel{\text{def}}{=} \{(w, x, y) \in \mathbb{V} \mid |k(w, x)| \geq (\gamma - \beta)\}.$$

Thus, a **good** triple (w, x, y) is **better** if (w, x) forms **good** triples with at least $(\gamma - \beta)$ fraction of the elements of \mathbb{Y} . We define another function $g : \mathbb{W} \mapsto \mathfrak{P}(\mathbb{X} \times \mathbb{Y})$ as

$$g(w) \stackrel{\text{def}}{=} \{(x, y) \in \mathbb{X} \times \mathbb{Y} \mid (w, x, y) \in \mathbb{V}^*\}.$$

g can be used to count, for a particular w , the number of elements $(x, y) \in \mathbb{X} \times \mathbb{Y}$ such that (w, x, y) forms a better triple. For some $\alpha < \beta$, we define the best triples of \mathbb{V}^* as

$$\mathbb{V}^{**} \stackrel{\text{def}}{=} \{(w, x, y) \in \mathbb{V}^* \mid |g(w)| \geq (\beta - \alpha)\}.$$

Thus, a better triple (w, x, y) is best if (w) forms better triples with at least $(\beta - \alpha)$ fraction of the elements of $\mathbb{X} \times \mathbb{Y}$.

Lemma 5 (The Extended Splitting Lemma). *The following propositions hold on the aforementioned assumptions:*

1. *i) $\mathbb{P}(\mathbb{V}^*) \geq \beta$, i.e., at least β fraction of triples in \mathbb{T} are better; and ii) $\mathbb{P}(\mathbb{V}^{**}) \geq \alpha$, i.e., at least α fraction of triples in \mathbb{T} are best.*
2. *i) $\mathbb{P}((w, x, y') \in \mathbb{V} \mid (w, x, y) \in \mathbb{V}^*; y' \stackrel{\$}{\leftarrow} \mathbb{Y}) \geq (\gamma - \beta)$, i.e., given a better triple (w, x, y) , the probability with which the triple (w, x, y') (with y' sampled randomly from \mathbb{Y}) is good is at least $(\gamma - \beta)$; and ii) $\mathbb{P}((w, x', y') \in \mathbb{V}^* \mid (w, x, y) \in \mathbb{V}^{**}; x' \stackrel{\$}{\leftarrow} \mathbb{X}; y' \stackrel{\$}{\leftarrow} \mathbb{Y}) \geq (\beta - \alpha)$ (interpreted as in 2.i).*
3. *i) $\mathbb{P}(\mathbb{V}^* \mid \mathbb{V}) \geq \beta/\gamma$, i.e., a good pair is better with a probability of at least β/γ ; and ii) $\mathbb{P}(\mathbb{V}^{**} \mid \mathbb{V}^*) \geq \alpha/\beta$, i.e., a better pair is best with a probability of at least β/γ ; and iii) $\mathbb{P}(\mathbb{V}^{**} \mid \mathbb{V}) \geq \alpha/\gamma$, i.e., a good pair is best with a probability of at least α/γ .*

Argument. The argument is quite similar to that of the Splitting Lemma. It is not difficult to see that **Proposition 2.i** and **2.ii** follow from the definitions (of better and best sets, respectively). The proof of **Proposition 1.i** is quite similar to that of **Proposition 1** of the Splitting Lemma (i.e., through contradiction). But it is given next for the sake of completeness. Suppose $\mathbb{P}(\mathbb{V}) \geq \gamma$ and $\mathbb{P}(\mathbb{V}^*) < \beta$, then we have

$$\mathbb{P}(\mathbb{V}) = \mathbb{P}(\mathbb{V}^*) + \mathbb{P}(\mathbb{V} \setminus \mathbb{V}^*) < \beta + \mathbb{P}(\mathbb{V} \setminus \mathbb{V}^*). \quad (2.16)$$

Now, there can be at most ℓm distinct (w, x) elements in $\mathbb{V} \setminus \mathbb{V}^*$. In addition, for each such $(w, x) \in \mathbb{V} \setminus \mathbb{V}^*$, $|k(w, x)| < (\gamma - \beta)n$ (by definition). Therefore (2.16) yields,

$$\mathbb{P}(\mathbb{V}) < \beta + \frac{|k(x, w)| \cdot \ell m}{\ell m n} < \beta + \frac{(\gamma - \beta)n \cdot \ell m}{\ell m n} = \gamma.$$

This contradicts our assumption that at least γ fraction of elements in \mathbb{T} are **good** ($\mathbb{P}(\mathbb{V}) \geq \gamma$), thus, establishing the proposition. The argument for **Proposition 2.ii** is similar as well. Suppose $\mathbb{P}(\mathbb{V}^*) \geq \beta$ and $\mathbb{P}(\mathbb{V}^{**}) < \alpha$, then we have

$$\mathbb{P}(\mathbb{V}^*) = \mathbb{P}(\mathbb{V}^{**}) + \mathbb{P}(\mathbb{V}^* \setminus \mathbb{V}^{**}) < \alpha + \mathbb{P}(\mathbb{V}^* \setminus \mathbb{V}^{**}). \quad (2.17)$$

Now, there can be at most ℓ distinct w elements in $\mathbb{V}^* \setminus \mathbb{V}^{**}$. In addition, for each such $w \in \mathbb{V} \setminus \mathbb{V}^*$, $|\mathbf{g}(w)| < (\gamma - \beta)mn$ (by definition). Therefore (2.17) yields,

$$\mathbb{P}(\mathbb{V}^*) < \alpha + \frac{|\mathbf{g}(w)| \cdot \ell}{\ell mn} < \alpha + \frac{(\beta - \alpha)mn \cdot \ell}{\ell mn} = \beta.$$

This contradicts our assumption ($\mathbb{P}(\mathbb{V}^*) \geq \gamma$), thus, establishing **Proposition 1.ii**. Finally, we turn to **Proposition 3.i** through **3.iii**. They can be established, as in **Proposition 3** of the Splitting Lemma, using Bayes' theorem. This concludes the argument. \square

2.5.1.2 Probability Analysis

Let H and G be the two random oracles involved in the nested replay attack. Nested forking, when carried out using rewinding, is successful if the target indices (of both G and H) during all four rounds of simulation of the adversary *match*. Thus, if (J_k, I_k) denote the target H and G indices, respectively, for round k , then the success condition is:

$$(I_3, J_3) = (I_2, J_2) = (I_1, J_1) = (I_0, J_0) \quad (2.18)$$

Defining the notions of “good”. Consider a general adversary A with the universe of tapes denoted by \mathbb{T} (as in §2.3.2.3). Provided that A has a non-negligible advantage of ϵ , at least ϵ fraction of the tapes in \mathbb{T} lead to its success (*e.g.*, it forging successfully). These constitute the good tapes in \mathbb{T} . The refined notion for **good** is a bit different as there are *two* random oracles in consideration. Let j [resp. i] denote the target H-index [resp. G-index] (see **Footnote 2**). On this assumption, for arbitrary indices $(j, i) \in \{1, \dots, q\}^2$, a tape in \mathbb{T} is deemed to be $\mathbf{good}_{(j,i)}$ if it leads to the adversary being successful *and* with a target H-index of j and a target G-index of i . It follows that a tape T is **good** is $\mathbf{good}_{(j,i)}$ for some $(j, i) \in \{1, \dots, q\}^2$. Let's assume that $\epsilon_{j,i}$

fraction of the tapes in \mathbb{T} are $\text{good}_{(j,i)}$. Thus, we have

$$\sum_{j=1}^q \sum_{i=1}^q \epsilon_{j,i} = \sum_{j=1}^q \sum_{i=1}^q \mathbb{P}(J = j \wedge I = i) = \mathbb{P}(J > 0 \wedge I > 0) = \epsilon. \quad (2.19)$$

To analyse forking involving two oracles, we need split the universe \mathbb{T} into the three sets $\mathbb{T}_{(j-)}$, $\mathbb{T}_{(ji)}$ and $\mathbb{T}_{(i+)}$ using the Cartesian product. As usual, we assume that \mathbb{T} is bijective to $\mathbb{T}_{(j-)} \times \mathbb{T}_{(ji)} \times \mathbb{T}_{(i+)}$ through a join and a split function and use this bijection to tailor the notion of $\text{good}_{(j,i)}$ for this set. We denote the (sub)set of all such tapes by $\mathbb{V}_{j,i}$.

Applying the Extended Splitting Lemma. Let's presume that an arbitrary $\beta_{j,i}$ and $\alpha_{j,i}$ fraction of the tape triples in \mathbb{T} are $\text{better}_{(j,i)}$ and $\text{best}_{(j,i)}$, respectively. We denote these sets, respectively, by $\mathbb{V}_{j,i}^*$ and $\mathbb{V}_{j,i}^{**}$. The exact value of $\beta_{j,i}$ and $\alpha_{j,i}$ is, again, determined later. The tapes involved are shown in **Figure 2.9**. The probability with which the nested replay attack is successful is calculated as follows. Refer to **Table 2.1** for notation.

Symbol	Denotes
\mathbb{T}^0	$(\mathbb{T}_{j-}^0, \mathbb{T}_{ji}^0, \mathbb{T}_{i+}^0)$
\mathbb{T}^1	$(\mathbb{T}_{j-}^0, \mathbb{T}_{ji}^0, \mathbb{T}_{i+}^1)$
\mathbb{T}^2	$(\mathbb{T}_{j-}^0, \mathbb{T}_{ji}^2, \mathbb{T}_{i+}^2)$
\mathbb{T}^3	$(\mathbb{T}_{j-}^0, \mathbb{T}_{ji}^2, \mathbb{T}_{i+}^3)$

Table 2.1: Shorthand for tapes.

$$\begin{aligned}
\epsilon'_{j,i} &= \mathbb{P}(\mathbf{T}^0 \in \mathbb{V}_{j,i}^{**} \wedge \mathbf{T}^1 \in \mathbb{V}_{j,i} \wedge \mathbf{T}^2 \in \mathbb{V}_{j,i}^* \wedge \mathbf{T}^3 \in \mathbb{V}_{j,i}) \\
&= \mathbb{P}(\mathbf{T}^1 \in \mathbb{V}_{j,i} \mid \mathbf{T}^0 \in \mathbb{V}_{j,i}^{**}) \cdot \mathbb{P}(\mathbf{T}^3 \in \mathbb{V}_{j,i} \wedge \mathbf{T}^2 \in \mathbb{V}_{j,i}^* \mid \mathbf{T}^0 \in \mathbb{V}_{j,i}^{**}) \cdot \mathbb{P}(\mathbf{T}^0 \in \mathbb{V}_{j,i}^{**}) \\
&\geq \mathbb{P}(\mathbf{T}^1 \in \mathbb{V}_{j,i} \mid \mathbf{T}^0 \in \mathbb{V}_{j,i}^*) \cdot \mathbb{P}(\mathbf{T}^3 \in \mathbb{V}_{j,i} \mid \mathbf{T}^2 \in \mathbb{V}_{j,i}^*) \cdot \mathbb{P}(\mathbf{T}^2 \in \mathbb{V}_{j,i}^* \mid \mathbf{T}^0 \in \mathbb{V}_{j,i}^{**}) \cdot \alpha_{j,i} \\
&\quad \text{(using **Proposition 1.ii**)} \\
&= \alpha_{j,i}(\beta_{j,i} - \alpha_{j,i})(\gamma_{j,i} - \beta_{j,i})^2 \text{(using **Proposition 2.i** and **2.ii**)}
\end{aligned}$$

The expression attains a maxima²¹ of $\epsilon_{j,i}^4/4^3$ at the point $(\epsilon_{j,i}/4, \epsilon_{j,i}/2, \epsilon_{j,i})$. Thus, the probability that the nested replay attack is successful for any (j, i) is given by

$$\begin{aligned}
\epsilon' &= \sum_{j=1}^q \sum_{i=1}^q \epsilon'_{j,i} \geq \sum_{j=1}^q \sum_{i=1}^q \frac{\epsilon_{j,i}^4}{4^3} \\
&\geq \frac{\epsilon^4}{4^3 q^6} \text{(using Hölder's inequality)} \tag{2.20}
\end{aligned}$$

Remark 8. The result can be extended for an arbitrary n number of forkings and we get $\epsilon' \geq \epsilon^{n+1}/4^n q^{2n}$.

The characteristic expression. The characteristic expression, evidently, is:

$$\epsilon' \geq \frac{1}{4^n} \sum_{j=1}^q \sum_{i=1}^q \epsilon_{ij}^{n+1} \tag{2.21}$$

under the set of constraints

$$\sum_{j=1}^q \sum_{i=1}^q \epsilon_{ij} = \epsilon \quad \text{and} \quad (0 \leq \epsilon_{ij} \leq 1)_{(i,j) \in \{1, \dots, q\}^2}. \tag{2.22}$$

The expression attains a minima of $\epsilon^{n+1}/4^n q^{2n}$ at the point $\epsilon_{ij} = \epsilon/q^2$ for $1 \leq i, j \leq q$, and hence the success probability of $\Omega(\epsilon^{n+1}/q^{2n})$.

²¹The objective function is $f(\beta_{j,i}, \alpha_{j,i}) = \alpha_{j,i}(\beta_{j,i} - \alpha_{j,i})(\epsilon_{j,i} - \beta_{j,i})^2$ under the constraint $(0 < \alpha_{j,i} < \beta_{j,i} < \epsilon_{j,i}) < 1$. The function attains maximum value at the point $(\epsilon_{j,i}/4, \epsilon_{j,i}/2, \epsilon_{j,i})$.

2.5.2 Multiple Forking

The Multiple-Forking Algorithm. Fix $q \in \mathbb{Z}^+$ and a set \mathbb{S} such that $|\mathbb{S}| \geq 2$. Let W be a randomised algorithm that on input a string x and elements $s_1, \dots, s_q \in \mathbb{S}$ returns a triple (I, J, σ) consisting of two integers $0 \leq J < I \leq q$ and a string σ . Let $n \geq 1$ be an odd integer. The multiple-forking algorithm $M_W n$ associated to W and n is defined as **Algorithm 2** below.

Algorithm 2 $M_W n(x)$

Pick coins ρ for W at random

$\{s_1^0, \dots, s_q^0\} \xleftarrow{U} \mathbb{S};$

$(I_0, J_0, \sigma_0) \leftarrow W(x, s_1^0, \dots, s_q^0; \rho)$ //round 0

if $((I_0 = 0) \vee (J_0 = 0))$ **then** $(0, \perp)$ //Condition $\neg B$

$\{s_{I_0}^1, \dots, s_q^1\} \xleftarrow{U} \mathbb{S};$

$(I_1, J_1, \sigma_1) \leftarrow W(x, s_1^0, \dots, s_{I_0-1}^0, s_{I_0}^1, \dots, s_q^1; \rho)$ //round 1

if $((I_1, J_1) \neq (I_0, J_0) \vee (s_{I_0}^1 = s_{I_0}^0))$ **then** $(0, \perp)$ //Condition $\neg C_0$

$k := 2$

while $(k < n)$ **do**

$\{s_{J_0}^k, \dots, s_q^k\} \xleftarrow{U} \mathbb{S};$

$(I_k, J_k, \sigma_k) \leftarrow W(x, s_1^0, \dots, s_{J_0-1}^0, s_{J_0}^k, \dots, s_q^k; \rho)$ //round k

if $((I_k, J_k) \neq (I_0, J_0) \vee (s_{J_0}^k = s_{J_0}^{k-1}))$ **then** $(0, \perp)$ //Condition $\neg D_k$

$\{s_{I_k}^{k+1}, \dots, s_q^{k+1}\} \xleftarrow{U} \mathbb{S};$

$(I_{k+1}, J_{k+1}, \sigma_{k+1}) \leftarrow W(x, s_1^0, \dots, s_{J_0-1}^0, s_{J_0}^k, \dots, s_{I_k-1}^k, s_{I_k}^{k+1}, \dots, s_q^{k+1}; \rho)$ //round $k+1$

if $((I_{k+1}, J_{k+1}) \neq (I_0, J_0) \vee (s_{I_0}^{k+1} = s_{I_0}^k))$ **then** $(0, \perp)$ //Condition $\neg C_k$

$k := k + 2$

end while

return $(1, \{\sigma_0, \dots, \sigma_n\})$

While the role of wrapper remains the same in the MF Algorithm, there are a few significant

changes in its actual structure. The wrapper now simulates two random oracles and hence its output contains a pair of indices (I, J) with $J < I$. The two indices are usually associated to the target queries made to the two random oracles involved in the replay attack. For reductions employing the MF Algorithm, the design of the wrapper becomes a bit more involved because of the additional index in its output. In particular, the relative “order” among the two target oracle queries must be taken into consideration in the design of the wrapper. (We’ll later see how neglecting the order of the indices, or even worse, using the MF Algorithm as a black-box may lead the reductions to fail.)

The cost of multiple forkings. According to the MF Lemma (see **Lemma 6**), the cost of n forkings (so the wrapper is called $n + 1$ times), carried out as per the specifications in **Algorithm 2**, is roughly $O(q^{2n})$, where q is the sum of the upper bound on the queries to the random oracles involved. The degradation, in turn, can be attributed to the set of conditions $\mathbb{A} := \{B, C_0, \dots, C_{n-1}, D_2, \dots, D_{n-1}\}$ where

$$\begin{aligned} B &: (I_0 \geq 1) \wedge (J_0 \geq 1) \\ C_k &: (I_{k+1}, J_{k+1}) = (I_k, J_k) \wedge (s_{I_k}^{k+1} \neq s_{I_k}^k) \quad (\text{for } k = 0, 2, \dots, n-1) \\ D_k &: (I_k, J_k) = (I_0, J_0) \wedge (s_{J_0}^k \neq s_{J_0}^{k-1}) \quad (\text{for } k = 2, 4, \dots, n-1) \end{aligned} \tag{2.23}$$

To be precise, the MF Algorithm is successful in the event E that all of the conditions in \mathbb{A}_0 are satisfied, *i.e.*,

$$E : B \wedge (C_0 \wedge C_2 \wedge \dots \wedge C_{n-1}) \wedge (D_2 \wedge D_4 \wedge \dots \wedge D_{n-1}). \tag{2.24}$$

The probability of this event, which is denoted by $mfrk$, is bounded by the MF lemma given below.

Lemma 6 (Multiple-Forking Lemma [BPW12]). *Let G_I be a randomised algorithm that takes no input and returns a string. Let*

$$\begin{aligned} mfrk &:= \mathbb{P}((op = 1) : x \stackrel{\$}{\leftarrow} G_I; (op, \{\sigma_0, \dots, \sigma_n\}) \stackrel{\$}{\leftarrow} M_{Wn}(x)) \quad \text{and} \\ acc &:= \mathbb{P}((I \geq 1) \wedge (J \geq 1) : x \stackrel{\$}{\leftarrow} G_I; \{s_1, \dots, s_q\} \stackrel{U}{\leftarrow} \mathbb{S}; (I, J, \sigma) \stackrel{\$}{\leftarrow} W(x, s_1, \dots, s_q)) \end{aligned}$$

then

$$mfrk \geq acc \cdot \left(\frac{acc^n}{q^{2n}} - \frac{n}{|\mathbb{S}|} \right). \tag{2.25}$$

The characteristic expression. The approach used for analysing the MF Algorithm is quite similar to that used for the GF Algorithm—*i.e.*, through random variables. However, as in the case of General Forking/Rewinding, the characteristic expression that we end up with is the same as for the nested replay attack (*i.e.*, (2.21) under the set of constraints given in (2.22)).

Remark 9 (On the degradation). Note that the MF Algorithm causes a substantial amount of degradation to the reduction. The dominant part of degradation incurred is due to the condition

$$(I_n, J_n) = (I_{n-1}, J_{n-1}) = \cdots = (I_0, J_0) \quad (2.26)$$

which is necessary for the success of the MF Algorithm. We call (2.26) the “core” condition for the MF Algorithm. Each of the n equality checks in the expression (on a high level) contributes a factor of $O(q^2)$, leading to the overall degradation of $O(q^{2n})$. In **Chapter 4**, we will revisit the expression with the objective of minimising the degradation on mind.

Remark 10 (Forking in terms of congruences and unknowns.). In the security argument of the Schnorr signature (which was discussed in the previous section), forking was used, in a vague sense, to obtain two linear congruences in two unknowns, given a linear congruence in two unknowns. In the same vein, Multiple-Forking (for $n=3$) can be considered to be a mechanism to obtain four linear congruence in four unknowns given a linear congruence in three unknowns.

Chapter 3

Galindo-Garcia IBS, Revisited

3.1 Introduction

In Africacrypt 2009, Galindo and Garcia [GG09] proposed a lightweight IBS scheme based on the Schnorr signature scheme. The construction is simple and claimed to be the most efficient IBS till date. The security argument consists of two reductions, B_1 and B_2 , the choice of which is determined by an event E . The authors construct B_1 to solve the DLP when the event E occurs. Likewise, B_2 is used to solve the DLP in case the complement of E occurs. Both the reductions use the MF Lemma (**Lemma 6**) to show that the DLP is reduced to breaking the IBS scheme. The authors suggest to implement their scheme in a suitable elliptic-curve group and, after a comparative study, concluded that the proposed construction has an overall better performance than the existing RSA-based and pairing-based schemes.

Revisiting the security argument. Critical examination of the security argument of a cryptographic construction to see whether the claimed security is indeed achieved or not is an important topic in cryptographic research. Two such well-known examples are Shoup’s work on OAEP [Sho01] and Galindo’s work on Boneh-Franklin IBE [Gal05]. Another important question in the area of provable security is to obtain tighter security reduction for existing constructions. One such classical example is Coron’s analysis of FDH [Cor00].

With this in mind, in this chapter, we revisit the security argument given in [GG09]. Our contributions are two fold; we: *i*) identify several problems in the original argument, and *ii*) provide a detailed new security argument that allows significantly tighter reductions. In particular, we show that the reduction B_1 in [GG09] *fails* in the standard security model for IBS [BNN04], while the

reduction B_2 is *incomplete*. Moreover, the details of the wrapper algorithms has been *neglected* in both the reductions—the forking algorithms have been used in a more or less black-box manner. To remedy these problems, we adopt a two-pronged approach. First, we sketch ways to fill the gaps by making minimal changes to the structure of the original security argument; then, we provide a new security argument. The new argument consists of three reductions: R_1 , R_2 and R_3 , and in each of them, solving the DLP is reduced to breaking the IBS. The reduction R_1 uses the GF Lemma (**Lemma 4**) and the technique first introduced by Coron [Cor00] to prove the security of FDH. We show that this results in a significantly tighter security reduction. On the other hand, both R_2 and R_3 are structurally similar to B_2 but uses two different versions of the MF Lemma, together with an algebraic technique similar to one adopted by Boneh-Boyen in [BB04a]. The security reduction R_2 is also significantly tighter than the original B_2 . All the three reductions use the programmability of the random oracles in a crucial way.

As the *bottom-up* approach—building complex systems from simpler, more primitive, systems—is prevalent in cryptography, sound security arguments are crucial in preventing the logical fallacies from getting carried over. Galindo-Galindo IBS (GG-IBS), due to its efficiency and simplicity, has been used as a building block for numerous systems [RS11, XW12]. Thus, to ensure that they rest on solid ground, it is important to give a sound argument for GG-IBS.

3.2 Revisiting the Galindo-Garcia Security Argument

We first reproduce the construction of GG-IBS and then identify several problems with the original security argument in [GG09].

3.2.1 The Construction

The scheme is based on the Schnorr signature scheme [Sch91] discussed in the previous chapter. The user secret key can be considered as the Schnorr signature by the PKG on the identity of the user, using the master secret key as the signing key. Analogously, the signature on a message by a user is the Schnorr signature, by that user, on the message using her user secret key as the signing key. The construction is given below (for further details see [GG09, §3]).

Set-up, $G(1^n)$: Invoke the group generator G_{DL} (on 1^n) to obtain (G, g, p) . Select $z \xleftarrow{U} \mathbb{Z}_p$ and set $Z = g^z$. Return z as the master secret key \mathbf{msk} and (G, p, g, Z, H, G) as the master public

key mpk , where H and G are hash functions

$$H : \{0, 1\}^* \mapsto \mathbb{Z}_p \quad \text{and} \quad G : \{0, 1\}^* \mapsto \mathbb{Z}_p.$$

Key Extraction, $E(\text{id}, \text{msk})$: Select $r \xleftarrow{U} \mathbb{Z}_p$ and set $R := g^r$. Return $\text{usk} := (y, R) \in \mathbb{Z}_p \times \mathbb{G}$ as the user secret key, where

$$y := r + zc \quad \text{and} \quad c := H(R, \text{id}).$$

Signing, $S(\text{id}, m, \text{usk})$: Let $\text{usk} = (y, R)$ and $c = H(R, \text{id})$. Select $a \xleftarrow{U} \mathbb{Z}_p$ and set $A := g^a$. Return $\sigma := (b, R, A) \in \mathbb{G} \times \mathbb{Z}_p \times \mathbb{G}$ as the signature, where

$$b := a + yd \quad \text{and} \quad d := G(\text{id}, A, m).$$

Verification, $V(\sigma, \text{id}, m, \text{mpk})$: Let $\sigma = (b, R, A)$, $c := H(R, \text{id})$ and $d := G(\text{id}, A, m)$. The signature is valid if

$$g^b = A(R \cdot Z^c)^d.$$

Figure 3.1: The (Original) Galindo-Garcia IBS.

break

Remark 11. Note that, although R is a part of the secret key of a user it is actually public information. In fact, R also forms a part of the signatures given by that user. Hence, by construction, all signatures generated using the user secret key $\text{usk} = (y, R)$ will contain the same R . This also means that, in the security argument, the simulator *has* to maintain the same R for a particular user; otherwise, the simulation will diverge from the actual protocol execution.

Remark 12. There are two hash functions in consideration: H and G . H is used to generate the user secret key which, in turn, is required to sign on a message (using G). Hence $H < G$ ($<$ denotes ‘followed by’) constitutes the *logical* order for calling the hash functions.

GG-IBS in terms of Schnorr Signature. In the beginning of the section, we had given a (vague) description of GG-IBS in terms of the Schnorr signature. Let’s elaborate on the idea. Let $\{G_s, K_s, S_s, V_s\}$ be the Schnorr signature scheme. The Key Extraction algorithm of GG-IBS can be realised using the Schnorr Signing algorithm S_s and the hash function H . The identity of the user

acts as the message, the master secret key as the signing key and, *i.e.*,

$$E(\text{id}, \text{msk}) = S_{s,H}(\text{id}, \text{msk}).$$

Similarly, if usk denotes the secret key for the user corresponding to id , the Signing algorithm of GG-IBS is realised using S_s as shown below.

$$S(\text{id}, m, \text{usk}) = S_{s,G}(m, \text{usk})$$

3.2.2 The Security Argument and Problems with it

The original security argument involves the construction of two algorithms B_1 and B_2 for complementary events E and NE respectively. E is the event that the attempted forgery $\hat{\sigma} = (\hat{A}, \hat{b}, \hat{R})$ is valid, non-trivial and \hat{R} is contained in the response to some query to the signature oracle. The event E (and its complement NE) is defined in this particular way in order to guarantee that the forgeries are in a format that helps the respective reductions to solve the DLP challenge.

In reduction B_2 , the problem instance is embedded as a part of the master public key and hence the master secret key is not known to B_2 . The extract queries are answered by using an *algebraic technique* similar to the one in [BB04a]. The signature queries, on the other hand, are answered by invoking the signing algorithm S after the user secret key has been generated as in the extract query. Finally, B_2 uses the MF Algorithm M_{W3} (see §2.5) to obtain four related forgeries and uses these forgeries to solve the DLP challenge.

On the other hand, the strategy adopted in B_1 is quite different from B_2 . The central idea is to embed the problem instance in the randomiser R while choosing its own master keys. In the simulation, B_1 randomly chooses one of the identities involved in G -oracle query as the target identity. For signature queries involving this target identity, B_1 embeds the problem instance in R and then uses the aforementioned algebraic technique to give the signature. The circularity involved is resolved by programming the random oracles. If A makes an extract query on the target identity, B_1 fails; for all other identities, B_1 uses the knowledge of the master secret key to respond to the queries. Finally it hopes that A returns a forgery containing R (in which the problem instance is embedded). In order to solve the DLP, B_1 needs to obtain two such forgeries. This is accomplished with the help of the MF Algorithm M_{W1} .

The original argument. We now reproduce the original reductions from [GG09, §4] using our notation (for ease of reference, the bullets are replaced by numeric values). Each reduction is followed, immediately, by the observations that cause its failure. In the following, $B_{i,j}$ refers to the j^{th} step in the construction of B_i , $i \in \{1, 2\}$. The description of the forking algorithms F_W and M_{Wn} is given in §2.4 and §2.5 respectively. (Some of the “typos” in the original security argument, that were corrected, are mentioned in the footnotes.)

Let A be an adversary against GG-IBS in the EU-ID-CMA model. Eventually, A outputs an attempted forgery of the form $\sigma = (A, B, R)$. Let E be the event that σ is a valid signature and R was contained in an answer of the signature oracle O_s . Let NE be the event that σ is a valid signature and R was never part of an answer of O_s . Galindo and Garcia construct an algorithm B_1 (resp. B_2) that breaks the DLP in case of event E (resp. NE).

3.2.2.1 Reduction B_1

B_1 takes as argument the description of a group (\mathbb{G}, p, g) and a challenge g^α with $\alpha \xleftarrow{U} \mathbb{Z}_p$ and tries to extract the discrete logarithm α . The protocol environment is simulated as shown below.

$\mathcal{B}_{1.1}$ B_1 picks $\hat{i} \xleftarrow{U} \{1, \dots, q_G\}$, where q_G is the maximum number of queries that the adversary A performs to the G -oracle. Let \hat{id} (the target identity) be the identity in the \hat{i}^{th} query to the G -oracle. Next, B_1 chooses $z \xleftarrow{U} \mathbb{Z}_p$ and sets $(\text{mpk}, \text{msk}) := ((\mathbb{G}, g, p, G, H, g^z), z)$, where G, H are descriptions of hash functions modelled as random oracles. As usual, B_1 simulates these oracles with the help of two tables \mathfrak{L}_G and \mathfrak{L}_H containing the queried values along with the answers given to A .

$\mathcal{B}_{1.2}$ Every time A queries the key extraction oracle O_E , for user id , B_1 chooses $c, y \xleftarrow{U} \mathbb{Z}_p$, sets $R := g^{-zc}g^y$ and adds $\langle R, id, c \rangle$ to the table \mathfrak{L}_H . Then it returns the key (y, R) to A .

$\mathcal{B}_{1.3}$ When A makes a query to the signature oracle O_s for (id, m) with $id \neq \hat{id}$, B_1 simply computes id 's secret key as described in the previous bullet. Then it invokes the signing algorithm S and returns the produced signature to A .

$\mathcal{B}_{1.4}$ When A makes a query to the signature oracle O_s for (id, m) with $id = \hat{id}$, B_1 chooses $t \xleftarrow{U} \mathbb{Z}_p, B \xleftarrow{U} \mathbb{G}$, sets $R := g^{-zc}(g^\alpha)^t, c := H(id, R)$,¹ and $A := B(g^\alpha g^{zc})^{-d}$.² Then it returns

¹In the original reduction, c was set to $H(\hat{id}, g^\alpha)$ instead of $H(\hat{id}, R)$. This is most likely a typo as it leads to the signatures on \hat{id} fundamentally failing the verification.

²Here, d is not assigned a value, though from the protocol we may infer that $d := G(id, A, m)$. But this leads to

the signature (A, B, R) to A .

$\mathcal{B}_1.5$ \mathcal{B}_1 invokes the algorithm $M_W1(\text{mpk})$ as described in Lemma 1 ([GG09, §4]). Here algorithm W is simply a wrapper that takes as explicit input, the answers from the random oracles. Then it calls A and returns its output together with two integers I, J . These integers are the indices of A 's queries to the random oracles G, H with the target identity $\hat{\text{id}}$.

Algorithm 3 $M_W1(\text{mpk})$

Pick random coins ρ for W

$$s_1^0, \dots, s_{q_G}^0 \xleftarrow{\text{U}} \mathbb{Z}_p$$

$$(I_0, J_0, \sigma_0) \leftarrow W(\text{mpk}, s_1^0, \dots, s_{q_G}^0; \rho)$$

If $(I_0 = 0 \vee J_0 = 0)$ then return \perp

$$s_{I_0}^1, \dots, s_{q_G}^1 \xleftarrow{\text{U}} \mathbb{Z}_p$$

$$(I_1, J_1, \sigma_1) \leftarrow W(\text{mpk}, s_1^0, \dots, s_{I_0-1}^0, s_{I_0}^1, \dots, s_{q_G}^1; \rho)$$

If $((I_1, J_1) \neq (I_0, J_0) \vee s_{I_0}^1 = s_{I_0}^0)$ then return \perp

Otherwise return (σ_0, σ_1)

In this way we get two forgeries of the form $\sigma_0 = (\text{id}, m, (A, B_0, R))$ and $\sigma_1 = (\text{id}, m, (A, B_1, R))$. Let d_0 be the answer from the G -oracle given to A in the first simulation, $s_{I_0}^0$ in M_W1 and let d_1 be the second answer $s_{I_0}^1$. If the identity id is not equal to the target identity $\hat{\text{id}}$ then \mathcal{B}_1 aborts. Otherwise it terminates and outputs the attempted discrete logarithm

$$\alpha = \frac{(B_0 - B_1)}{td_0 - td_1}.$$

Observations on \mathcal{B}_1 . We now note the following points about the reduction \mathcal{B}_1 given above. We also mention ways to fix the problems.

Observation 1 (Correctness of signatures on $\hat{\text{id}}$). *In $\mathcal{B}_1.4$, when A makes a signature query on $\hat{\text{id}}$, \mathcal{B}_1 returns $(A, B, R) \in \mathbb{G}^3$ as the signature. However, in the protocol definition, the signatures are elements of $\mathbb{G} \times \mathbb{Z}_p \times \mathbb{G}$. Therefore, the signatures on $\hat{\text{id}}$ will fail the verification in the general*

a circularity as the value of A depends on d . To avoid this circularity, \mathcal{B}_1 has to program G -oracle as follows: choose $d \xleftarrow{\text{U}} \mathbb{Z}_p$, compute $A = B(g^\alpha g^{zc})^{-d}$ and then set $G(\text{id}, A, m) := d$.

group setup—i.e., \mathbb{G} is any cyclic group of prime order p , and in particular, in the elliptic curve setting—as the operation g^B is not defined in \mathbb{G} .

What the authors *could* have intended in $\mathsf{B}_1.4$ is

- When A queries the signature oracle O_s with (id, m) where $\mathsf{id} = \hat{\mathsf{id}}$, B_1 chooses $t, b \xleftarrow{\mathsf{U}} \mathbb{Z}_p$, sets $B := g^b, R := g^{-zc}(g^\alpha)^t, c := \mathsf{H}(\mathsf{id}, R)$ and $A := B(g^\alpha g^{zc})^{-d}$. Then it returns the signature (A, b, R) to A .

Even after the above correction is applied, the signatures on $\hat{\mathsf{id}}$ fail the verification algorithm. For the signatures to verify, the following equality should hold.

$$\begin{aligned} g^b &= A(R \cdot (g^\alpha)^c)^d \\ &= g^b (g^\alpha g^{zc})^{-d} (g^{-zc} (g^\alpha)^t g^{zc})^d \\ 1 &= g^{(\alpha+zc)(-d)} g^{\alpha t d} \end{aligned}$$

However, it holds only if

$$(\alpha t - zc - \alpha) d \equiv 0 \pmod{p}. \quad (3.1)$$

It is easy to check that the LHS in (3.1) is a random element of \mathbb{Z}_p . Hence, the signatures on $\hat{\mathsf{id}}$ given by B_1 will fail to verify with an overwhelming probability of $1 - 1/p$. The equality holds if we set $t := 1 + zc/\alpha$, instead of selecting t uniformly at random from \mathbb{Z}_p .³ However, setting $t := 1 + zc/\alpha$ results in R being set to the problem instance g^α , removing t from the picture altogether. Thus, $\mathsf{B}_1.4$ would finally look like:

- When A queries the signature oracle O_s with (id, m) where $\mathsf{id} = \hat{\mathsf{id}}$, B_1 chooses $b, d \xleftarrow{\mathsf{U}} \mathbb{Z}_p$, sets $B := g^b, R := g^\alpha, c := \mathsf{H}(\mathsf{id}, R), A := B(g^\alpha g^{zc})^{-d}$ and programs the random oracle in such a way that $d := \mathsf{G}(\mathsf{id}, A, m)$. Then it returns the signature (A, b, R) to A .

Although it may appear that the reduction B_1 can be rescued with the modification mentioned above, the line of argument in B_1 has another inherent—much more serious—problem, which we describe next.

Observation 2 (Ambiguity due to the choice of $\hat{\mathsf{id}}$). B_1 sets the identity involved in the \hat{i}^{th} G -oracle query as the target identity $\hat{\mathsf{id}}$ (see $\mathsf{B}_1.1$). Hence, the target identity can be fixed only after

³This modification was pointed out by an anonymous reviewer.

the \hat{i}^{th} query to the G -oracle has been made. However, whenever a signature query is made on any identity, B_1 has to decide whether the identity is the target identity or not. Therefore, when A makes a signature query before the \hat{i}^{th} G -oracle query, B_1 has no way to decide whether to proceed to $B_{1.3}$ or $B_{1.4}$ (as it depends on whether $\text{id} = \hat{\text{id}}$ or not).

B_1 can provide a proper simulation of the protocol environment only if *no* signature query is made on the target identity $\hat{\text{id}}$ before the \hat{i}^{th} G -oracle query. However, B_1 cannot really restrict the adversarial strategy this way. In fact, B_1 will fail to give a proper simulation of the protocol environment if A makes one signature query on $\hat{\text{id}}$ before the \hat{i}^{th} G -oracle query and one more signature query on $\hat{\text{id}}$ after the \hat{i}^{th} G -oracle query.

One way to fix the problem noted above is to guess the “index” of the target identity *instead* of guessing the target G -index. Suppose n distinct identities are involved in the queries to the G -oracle, where $1 \leq n \leq q_G$.⁴ The strategy would be to guess the index \hat{i} of the target identity $\hat{\text{id}}$ among all the identities, *i.e.* if $\{\text{id}_1, \dots, \text{id}_n\}$ were the *distinct* identities involved in the queries to the G -oracle (in that order), we set $\text{id}_{\hat{i}}$ with $1 \leq \hat{i} \leq n$ as the target identity. Now, by assumption no identity queried to the G -oracle prior to $\text{id}_{\hat{i}}$ can be the target identity. Hence, the ambiguity noted before can be avoided. Although this strategy works well with the “mended” reduction that we ended up in **Observation 1**, it will still incur a tightness loss of the order $O(q_G^3)$.

In our alternative security argument given in §3.3, we show how to get around the problem in B_1 by using Coron’s technique, together with some algebraic manipulation and non-trivial random oracle programming. In addition to correcting the errors in B_1 , we end up with a much tighter reduction as a result.

3.2.2.2 Reduction B_2

It takes as argument, the description of a group (\mathbb{G}, p, g) and a challenge g^α with $\alpha \xleftarrow{U} \mathbb{Z}_p$ and outputs the discrete logarithm α . To do so, it will run A simulating the environment as shown below.

$B_{2.1}$ At the beginning of the experiment, B_2 sets the master public key $\text{mpk} := (\mathbb{G}, p, g, G, H)$ and $\text{msk} := (g^\alpha)$, where G, H are description of hash functions modelled as random oracles. As usual, B_2 simulates these oracles with the help of two tables \mathfrak{L}_G and \mathfrak{L}_H containing the queried values together with the answers given to A .

⁴ B_1 will maintain a counter and increment it by 1 each time a new identity is queried to the G -oracle.

$\mathcal{B}_2.2$ Every time A queries the key extraction oracle \mathcal{O}_E , for user id , \mathcal{B}_2 chooses $c, y \xleftarrow{\text{U}} \mathbb{Z}_p$, sets $R := g^{-\alpha c} g^y$ and adds $\langle R, \text{id}, c \rangle$ to the table \mathcal{L}_H . Then it returns the key (y, R) to A.

$\mathcal{B}_2.3$ When A queries the signature oracle \mathcal{O}_s with (id, m) , \mathcal{B}_2 simply computes id 's secret key as described in the previous step. Then it computes a signature by calling \mathcal{S} , adding the respective call to the G-oracle, $((\text{id}, g^a, m), d)$ to the table \mathcal{L}_G and gives the resulting signature to the adversary.

$\mathcal{B}_2.4$ \mathcal{B}_2 invokes the algorithm $M_{\mathbb{W}3}(\text{mpk})$. In this way either \mathcal{B}_2 aborts prematurely or we get, for some identity id , some message m and some R , four forgeries $(\text{id}, m, (A_k, b_k, R))$ ⁵, $k := 0, \dots, 3$ with $A_0 = A_1$ and $A_2 = A_3$. As all these signatures are valid, the following equations hold.

$$\begin{aligned} \{b_0 = \log A_0 + (\log R + c_0\alpha)d_0, b_1 = \log A_1 + (\log R + c_0\alpha)d_1, \\ b_2 = \log A_2 + (\log R + c_1\alpha)d_2, b_3 = \log A_3 + (\log R + c_1\alpha)d_3\} \end{aligned} \quad (3.2)$$

with $c_0 \neq c_1, d_0 \neq d_1$ and $d_2 \neq d_3$. Since we know $c_0, c_1, d_0, \dots, d_3$, a simple computation yields

$$\alpha = \frac{b_2 + b_1 - b_0 - b_3}{c_1(d_2 - d_3) - c_0(d_0 - d_1)}. \quad (3.3)$$

Observations on \mathcal{B}_2 . We now note the following points about the reduction \mathcal{B}_2 given above. As in \mathcal{B}_1 , we discuss possible fixes.

Observation 3 (Incorrect solution of the DLP instance.). *In Step $\mathcal{B}_2.4$, the reduction obtains the solution of the DLP instance by solving the four equations given in (3.2). However, on substituting the values of b_k s from (3.2) in (3.3) we get*

$$\frac{b_2 + b_1 - b_0 - b_3}{c_1(d_2 - d_3) - c_0(d_0 - d_1)} = \alpha + \log gR \cdot \frac{d_2 + d_1 - d_0 - d_3}{c_1(d_2 - d_3) - c_0(d_0 - d_1)}, \quad (3.4)$$

which is not the correct solution to the DLP instance.

Note that the simulator does not know the value of $\log gR$ and hence cannot extract α from the above expression. However, it is not difficult to get the correct solution (as we show in (3.12) of

⁵We use b_k instead of B_k , throughout the reduction, to maintain consistency with the protocol description (in §3.2.1).

§3.3.3). The more fundamental problem is that \mathcal{B}_2 fails to capture all possible adversarial strategies as we show next.

Observation 4 (Incompleteness of \mathcal{B}_2). *In Step $\mathcal{B}_{2.4}$, \mathcal{B}_2 invokes M_{W3} to get four forged signatures with b_k s as given in (3.2). The b_k component of the forged signatures, though, need not always have this particular structure.*

The structure depends on the precise order in which A makes the target oracle queries $G(\text{id}, A, m)$ and $H(R, \text{id})$ during the simulation. (Here, (id, m) corresponds to the target identity and the message pair in the forgery while (A, R) are components of the forged signature.) Thus, (3.2) covers only one of the two possible adversarial behaviours: A querying the random oracles in the logical order $H < G$ (shown in **Figure A.3** where the first branching corresponds to the forking of the H -oracle). But one cannot rule out the complementary case of $G < H$: A querying the G -oracle before the H -oracle (see **Figure 3.2**) where the first branching corresponds to the forking of the G -oracle.⁶ Let's look into the structure of the forged signatures in the case $G < H$. As a result of

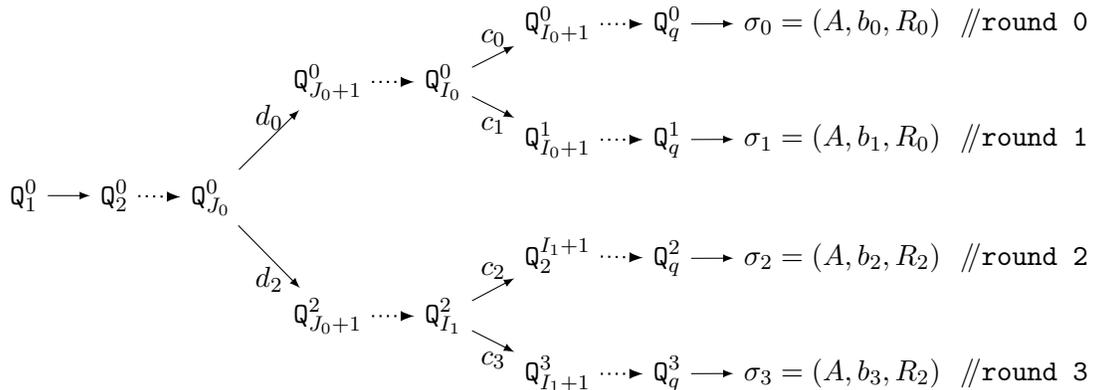


Figure 3.2: Structure of the forgeries in the case $G < H$. $Q_{J_0}^0$ denotes the target G -query $G(\text{id}, A, m)$; $Q_{I_0}^0$ (resp $Q_{I_0}^2$) denotes the target H -query $H(R_0, \text{id})$ (resp $H(R_2, \text{id})$).

the ordering of the oracle queries, W returns J_0 as the index of the G -oracle query on (id, A, m) and I_0 as the index of the H -oracle query on (R_0, id) , at the end of round 0. As G -oracle is forked

⁶This is captured by an adversary A with the following behaviour:

- Fix a target identity-message pair $(\hat{\text{id}}, \hat{m})$ and corresponding $\hat{R}, \hat{A} \in \mathbb{G}$.
- Make the two oracle queries: $G(\hat{\text{id}}, \hat{A}, \hat{m})$ and $H(\hat{R}, \hat{\text{id}})$, in that order.
- Produce a forgery $\sigma = (\hat{A}, \hat{B}, \hat{R})$ on $(\hat{\text{id}}, \hat{m})$.

before the H-oracle, we get $d_1 = d_0$, $d_3 = d_2$ and $R_1 = R_0$, $R_3 = R_2$ in the subsequent forkings, while all the c_i , $0 \leq i \leq 3$ will be different. On the other hand, the value A returned as part of the forged signature remains the same in all the four rounds. Hence, the signatures returned by M_W3 will contain b_k s of the form:

$$\begin{aligned} \{b_0 = \log A + (\log R_0 + c_0\alpha)d_0, b_1 = \log A + (\log R_0 + c_1\alpha)d_0, \\ b_2 = \log A + (\log R_2 + c_2\alpha)d_2, b_3 = \log A + (\log R_2 + c_3\alpha)d_2\} \end{aligned} \quad (3.5)$$

When the signatures have the structure as in (3.5), we cannot use (3.3) (more precisely, the corrected version as given in (A.18) of **Appendix A.2.3.2**) to get a solution of the DLP. This is because $d_1 = d_0$ and $d_3 = d_2$ makes the denominator part in the corresponding expression zero. As we cannot rule out this particular adversary, the reduction does not address all the cases, rendering it incomplete.

To summarize, the same strategy to solve the DLP will *not* work for the two aforementioned complementary cases. Still it is possible to distinguish between the two cases ($H < G$ and $G < H$) simply by looking at the structure of the forged signatures. In the case $H < G$, all the R s will be equal, *i.e.* $R_3 = R_2 = R_1 = R_0$; as for $G < H$, all the A s will be equal, *i.e.* $A_3 = A_2 = A_1 = A_0$. We could then use appropriate relations⁷ to solve for the DLP instance. However, this results in an *unnecessary* forking (the branch consisting of **round 2** and **round 3** in **Figure 3.2**) being carried out in the case $G < H$. We address this in §3.3 by *splitting* B_2 into two reductions R_2 and R_3 , with R_2 involving only a single forking. The single forking, in turn, leads to a tighter reduction (see **Table 3.1**).

3.3 New Security Argument

On the basis of the observations made in the previous section, we now proceed to provide a detailed security argument for GG-IBS. In a nutshell, we have effectively *modularised* the security argument into three mutually exclusive parts so that each of the three situations mentioned in the previous section can be studied in more detail. We also show that it is possible to obtain tighter reductions in two of the three cases.

In order to address the problem in B_1 we redefine the event E and to address the incompleteness

⁷*i.e.*, (3.12) and (3.10) derived in **Appendix A.2.3** and **A.2.2**, respectively. For the sake of completeness, we provide the modified security argument incorporating all the above mentioned fixes in the **Appendix A.1**.

of B_2 we introduce another event F . The security argument involves constructing three algorithms: R_1 , R_2 and R_3 and in each of them solving the DLP is reduced to breaking the IBS. R_1 , unlike its counterpart B_1 , uses the GF Algorithm, whereas R_2 and R_3 , the counterparts of B_2 , still use the MF algorithm. The new reductions R_1 and R_2 are also tighter than their counterparts in the original argument. We also use wrappers in all the three reductions.

Theorem 2. *Let A be an $(\epsilon, t, q_\epsilon, q_s, q_H, q_G)$ -adversary against the IBS in the EU-ID-CMA model. If the hash functions H and G are modelled as random oracles, we can construct either*

(i) *Algorithm R_1 that (ϵ_1, t_1) -breaks the DLP, where*

$$\epsilon_1 \geq \frac{\epsilon^2}{\exp(1)q_Gq_\epsilon} \quad \text{and} \quad t_1 \leq t + 2(q_\epsilon + 3q_s)\tau, \quad \text{or}$$

(ii) *Algorithm R_2 that (ϵ_2, t_2) -breaks the DLP, where*

$$\epsilon_2 \geq \epsilon \left(\frac{\epsilon}{(q_H + q_G)^2} - \frac{1}{p} \right) \quad \text{and} \quad t_2 \leq t + 2(2q_\epsilon + 3q_s)\tau, \quad \text{or}$$

(iii) *Algorithm R_3 that (ϵ_3, t_3) -breaks the DLP, where*

$$\epsilon_3 \geq \epsilon \left(\frac{\epsilon^3}{(q_H + q_G)^6} - \frac{3}{p} \right) \quad \text{and} \quad t_3 \leq t + 4(2q_\epsilon + 3q_s)\tau.$$

Here q_ϵ (resp. q_s) denotes the upper bound on the number of extract (resp. signature) queries that A can make; q_H (resp. q_G) denotes the upper bound on the number of queries to the H -oracle (resp. G -oracle). τ is the time taken for an exponentiation in the group \mathbb{G} and \exp is the base of natural logarithm.

Argument. A is successful if it produces a valid forgery $\hat{\sigma} = (\hat{A}, \hat{b}, \hat{R})$ on (\hat{id}, \hat{m}) . Consider the following event⁸ in the case that A is successful.

E: A makes at least one signature query on \hat{id} and \hat{R} was returned by the simulator as part of the output to a signature query on \hat{id} .

The complement of this event is

⁸Note that the definition of the new event E (and $\neg E$) is slightly different from the one given in the security argument of [GG09], i.e. event E (and NE) discussed in §3.2.2.

$\neg E$: Either A does not any make signature query on \hat{id} or \hat{R} was never returned by the simulator as part of the output to a signature query on \hat{id} .

In order to come up with the forgery $\hat{\sigma}$ with a non-negligible probability, the adversary, at some juncture during its simulation, has to make the two random oracle queries: $H(\hat{R}, \hat{id})$ and $G(\hat{id}, \hat{A}, \hat{m})$. Depending on the order in which A makes these calls, we further subdivide the event $\neg E$ into an event F and its complementary event $\neg F$, where

F : The event that A makes the oracle query $G(\hat{id}, \hat{A}, \hat{m})$ before the oracle query $H(\hat{R}, \hat{id})$ ($G < H$).

$\neg F$: The event that A makes the oracle query $H(\hat{R}, \hat{id})$ before the oracle query $G(\hat{id}, \hat{A}, \hat{m})$ ($H < G$).

In the case of the events E , $\neg E \wedge F$ and $\neg E \wedge \neg F$, we give the reductions R_1 , R_2 and R_3 respectively. They are described in the subsequent sections. \square

Simulating the random oracles. A random oracle query is defined to be *fresh* if it is the first query involving that particular input. If a query is not fresh for an input, in order to maintain consistency, the random oracle has to respond with the same output as in the previous query on that input. We say that a fresh query does not require *programming* if the simulator can simply return a random value as the response. The crux of most security arguments involving random oracles, including ours, is the way the simulator answers the queries that require programming. In our case, random oracle programming is used to resolve the circularity involved while dealing with the *implicit* random oracle queries. A random oracle query is said to be implicit if it is not an explicit query by the adversary or the simulator. As usual, to simplify the book-keeping, all implicit random oracle queries involved in answering the extract and signature queries are put into the account of A .

3.3.1 Reduction R_1

Let $\Delta := (\mathbb{G}, p, g, g^\alpha)$ be the given DLP instance. The reduction involves invoking the GF Algorithm on a wrapper Y as shown in **Algorithm 4**. As a result, it obtains a set of two congruences in two unknowns and solves for α . It can be verified, as we do later, that R_1 indeed returns the correct solution to the DLP instance. The novelty in the design of Y lies in the way the problem instance is embedded in the randomiser R instead of the master public key— R_1 generates its own master keys.

Algorithm 4 Reduction $R_1(\Delta)$

Select $z \xleftarrow{U} \mathbb{Z}_p^*$ as the **msk** and set $\text{mpk} := (\mathbb{G}, g, p, g^z)$.
 $(\text{op}, \sigma_0, \sigma_1) \xleftarrow{\mathbb{S}} F_Y((\text{mpk}, \text{msk}), g^\alpha)$
if $(\text{op} = 0)$ **then** \perp //abort_{1,2}
Parse σ_i as $(\hat{b}_i, c_i, r_i, \beta_i, d_i)$.
if $(\beta_0 = 1) \wedge (\beta_1 = 0)$ **then return** $(z(c_0d_0 - c_1d_1) + r_0d_0 - (\hat{b}_0 - \hat{b}_1))/r_1d_1$
else if $(\beta_0 = 0) \wedge (\beta_1 = 1)$ **then return** $(z(c_1d_1 - c_0d_0) + r_1d_1 - (\hat{b}_1 - \hat{b}_0))/r_0d_0$
else if $(\beta_0 = 0) \wedge (\beta_1 = 0)$ **then return** $((\hat{b}_0 - \hat{b}_1) - z(c_0d_0 - c_1d_1))/(r_0d_0 - r_1d_1)$
else return \perp //abort_{1,3}
end if

The Wrapper

The main ingredient is the so-called “partitioning strategy”, first used by Coron in the security argument of FDH [Cor00]. The basic idea is to divide the identity-space \mathbb{I} into two disjoint sets, \mathbb{I}_ε and \mathbb{I}_s , depending upon the outcome of a biased coin. Y is equipped to respond to both extract and signature queries on identities from \mathbb{I}_ε . But it fails if the adversary does an extract query on any identity from \mathbb{I}_s ; it can answer only to signature queries on identities from \mathbb{I}_s . The problem instance is embedded in the randomiser R , depending on the outcome of the biased coin. As Y maintains a unique R for each identity, the structure of R decides whether that identity belongs to \mathbb{I}_ε or to \mathbb{I}_s . The optimal size of the sets is determined on analysis. The details follow.

Suppose that $q := q_G$ and $\mathbb{S} := \mathbb{Z}_p$. Y takes as input the master keys (mpk, msk) , the problem instance g^α and $\{s_1, \dots, s_q\}$. It returns a pair (I, σ) where I is the target G-index and σ is the side-output. In order to track the index of the current G-oracle query, Y maintains a counter ℓ , initially set to 1. It also maintains a table \mathfrak{L}_H (resp. \mathfrak{L}_G) to manage the random oracle H (resp. G). Y initiates the EU-ID-CMA game by passing mpk as the challenge master public key to the adversary A . The queries by A are handled as per the following specifications.

(a) **Random oracle query**, $H(R, \text{id})$: \mathfrak{L}_H contains tuples of the form

$$\langle R, \text{id}, c, r, \beta \rangle \in \mathbb{G} \times \{0, 1\}^* \times \mathbb{Z}_p \times \mathbb{Z}_p \cup \{\perp\} \times \{0, 1, \phi\}.$$

Here, (R, id) is the query to the H-oracle and c is the corresponding output. Therefore, a query $H(R, \text{id})$ is fresh if there exists no tuple $\langle R_i, \text{id}_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(R_i = R) \wedge (\text{id}_i = \text{id})$.

If such a tuple exists, then the oracle has to return the corresponding c_i as the output.

The r -field is used to store additional information related to the R -field. The tuples corresponding to the explicit H-oracle queries, made by A , are tracked by storing ' \perp ' in the r -field. This indicates that Y does not have any additional information regarding R . In these tuples, the β -field is irrelevant and this is indicated by storing ' ϕ '. In tuples with $r \neq \perp$, the β -field indicates whether the DLP instance is embedded in R or not. If $\beta = 0$ then $R = (g^\alpha)^r$ for some known $r \in \mathbb{Z}_p$, which is stored in the r -field. On the other hand, $\beta = 1$ implies $R = g^r$ for some known $r \in \mathbb{Z}_p$, which is, again, stored in the r -field. We now explain how the fresh H-oracle queries are handled. The query may be

- (i) H_1 , Explicit query made by A : In this case Y returns $c \xleftarrow{U} \mathbb{Z}_p$ as the output. $\langle R, \text{id}, c, \perp, \phi \rangle$ is added to \mathfrak{L}_H .
- (ii) H_2 , Explicit query made by Y : As in the previous case, Y returns $c \xleftarrow{U} \mathbb{Z}_p$ as the output. As Y knows $r = \log gR$, $\langle R, \text{id}, c, r, 1 \rangle$ is added to \mathfrak{L}_H .
- (iii) H_3 , Implicit query by Y in order to answer a signature query made by A : See step (iii) of **Signature query** on how to program the random oracle in this situation.

(b) **Random oracle query**, $G(\text{id}, A, m)$: \mathfrak{L}_G contains tuples of the form

$$\langle \text{id}, A, m, d, \ell \rangle \in \{0, 1\}^* \times \mathbb{G} \times \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}^+.$$

Here, (id, A, m) is the query to the G-oracle and d is the corresponding output. The index of the query is stored in the ℓ -field. Therefore, a random oracle query $G(\text{id}, A, m)$ is fresh if there exists no tuple $\langle \text{id}_i, A_i, m_i, d_i, \ell_i \rangle$, in \mathfrak{L}_G such that $(\text{id}_i = \text{id}) \wedge (A_i = A) \wedge (m_i = m)$. If such a tuple exists, then the oracle has to return the corresponding d_i as the output. We now explain how the fresh G-oracle queries are handled. The query may be

- (i) G_1 , Explicit query made by either A or Y : In this case Y returns $d := s_\ell$ as the output. $\langle \text{id}, A, m, d, \ell \rangle$ is added to \mathfrak{L}_G and ℓ is incremented by one.
- (ii) G_2 , Implicit query by Y in order to answer a signature query made by A : See steps (i) and (iii) of **Signature query** on how to program the random oracle in this situation.

(c) **Extract query**, $O_E(\text{id})$: Y first checks if id has an associated R . This is done by searching for tuples $\langle R_i, \text{id}_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H with $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$. If such a tuple exists, Y checks

for the value of β_i in the tuple. $\beta_i = 0$ implies the identity belongs to \mathbb{I}_s and consequently the extract query fails, leading to Υ aborting the simulation: $\text{abort}_{1,1}$. On the other hand, $\beta_i = 1$ implies that there was a prior extract query on id and also that the identity belongs to \mathbb{I}_ε . Υ generates the user secret key (same as in prior extract query) using the information available in the tuple. On the other hand, if such a tuple does not exist, Υ selects a fresh r and assigns id to \mathbb{I}_ε . Υ has this freedom since the adversary *cannot* forge on this identity. A more formal description follows.

If there exists a tuple $\langle R_i, \text{id}_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$

- (i) If $\beta_i = 0$, Υ *aborts* the simulation ($\text{abort}_{1,1}$) and returns $(0, \perp, \perp)$.
- (ii) Otherwise, $\beta_i = 1$ and Υ returns $\text{usk} := (r_i + zc_i, R_i)$ as the user secret key.

Otherwise

- (iii) Υ chooses $r \xleftarrow{U} \mathbb{Z}_p$, sets $R := g^r$ and queries the H-oracle for $c := H(\text{id}, R)$. It returns $\text{usk} := (r + zc, R)$ as the secret key.

- (d) **Signature query**, $O_s(\text{id}, m)$: As in **Extract query**, Υ checks the identity for an associated R by searching tuples $\langle R_i, \text{id}_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H with $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$. If such a tuple exists, the identity has been assigned to either of \mathbb{I}_ε or \mathbb{I}_s , determined by the value of β_i . If such a tuple does not exist, then the identity is *unassigned* and Υ assigns the identity to either \mathbb{I}_ε or \mathbb{I}_s by tossing a (biased) coin β . If the outcome is 0, id is assigned to \mathbb{I}_s ; else it is assigned to \mathbb{I}_ε . Identities assigned to \mathbb{I}_s have the problem instance g^α embedded in the randomiser R . Although the private key cannot be calculated, an algebraic technique, similar to one adopted by Boneh-Boyen in [BB04a], coupled with random oracle programming enables us to give the signature. On the other hand, signature queries involving identities from \mathbb{I}_ε are answered by first generating usk as in **Extract query** and then invoking S . A more formal description follows.

If there exists a tuple $\langle R_i, \text{id}_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$

- (i) If $\beta_i = 0$, Υ selects $s \xleftarrow{U} \mathbb{Z}_p$ and sets $d := s_\ell$, $A := g^s(g^\alpha)^{-r_i d}$. It then adds $\langle \text{id}, A, m, d, \ell \rangle$ to \mathfrak{L}_G (deferred case G_2)⁹ and increments ℓ by one. The signature returned is $\sigma := (A, s + zcd, R_i)$.

⁹ In the unlikely event of there already existing a tuple $\langle \text{id}_i, A_i, m_i, d_i, \ell_i \rangle$ in \mathfrak{L}_G with $(\text{id}_i = \text{id}) \wedge (A_i = A) \wedge (m_i = m)$ but $(d_i \neq d)$ then $G(\text{id}, A, m)$ cannot be set to d . In that case Υ can simply increment ℓ and repeat step (i).

- (ii) Otherwise, $\beta_i = 1$ and the user secret key is $\text{usk} := (y, R_i)$, where $y = r_i + zc_i$ and $R_i = g^{r_i}$. Υ then selects $a \xleftarrow{\text{U}} \mathbb{Z}_p$, sets $A := g^a$ and queries the G-oracle with $d := G(\text{id}, A, m_i)$. The signature returned is $\sigma := (A_i, a + yd, R_i)$.

Otherwise, Υ tosses a coin β with a bias δ (i.e., $\Pr[\beta = 0] = \delta$). The value of δ will be quantified on analysis.

- (iii) If $\beta = 0$, Υ selects $c, s, r \xleftarrow{\text{U}} \mathbb{Z}_p$ and sets $d := s_\ell$, $R := (g^\alpha)^r$, $A := g^s (g^\alpha)^{-rd}$. Next, it adds $\langle \text{id}, (g^\alpha)^r, c, r, 0 \rangle$ to \mathfrak{L}_H (deferred case H_3), $\langle \text{id}, A, m, d, \ell \rangle$ to \mathfrak{L}_G (deferred case G_2) and increments ℓ by one.¹⁰ The signature returned is $\sigma := (A, s + zc_\ell d, R)$.
- (iv) Otherwise, $\beta = 1$ and Υ selects $a, r \xleftarrow{\text{U}} \mathbb{Z}_p$ and sets $A := g^a$, $R := g^r$. It then queries the respective oracles with $c := H(R, \text{id})$ and $d := G(\text{id}, A, m)$. The signature returned is $\sigma := (A, a + (r + zc)d, R)$.

Correctness of the signatures. For $\beta = 1$, the signatures are generated as in the protocol and hence they fundamentally verify. For $\beta = 0$, the signature given by Υ is of the form (A, b, R) , where $A = g^s (g^\alpha)^{-rd}$, $b = s + zcd$ and $R = (g^\alpha)^r$. Υ also sets $c := H(R, \text{id})$ and $d := G(\text{id}, A, m)$. As a result, the signature verifies as shown below.

$$\begin{aligned}
g^b &= g^{s+zcd} \\
&= g^{s-\alpha rd + \alpha rd + zcd} \\
&= g^s (g^\alpha)^{-rd} ((g^\alpha)^r (g^z)^c)^d \\
&= A(R(g^z)^c)^d.
\end{aligned}$$

At the end of the simulation, a successful adversary forges $\hat{\sigma} := (\hat{A}, \hat{b}, \hat{R})$ on $(\hat{\text{id}}, \hat{m})$. Let $\langle R_j, \text{id}_j, c_j, r_j, \beta_j \rangle$ be the tuple in \mathfrak{L}_H that corresponds to the target H-query. Similarly, let $\langle \text{id}_i, A_i, m_i, d_i, \ell_i \rangle$ be the tuple in \mathfrak{L}_G that corresponds to the target G-query. Υ returns $(\ell_i, (\hat{b}, c_j, r_j, \beta_j, d_i))$ as its own output. Note that the side-output σ consists of $(\hat{b}, c_j, r_j, \beta_j, d_i)$.

That concludes the description of the wrapper.

¹⁰ Υ chooses different randomisers if there is a collision as explained in **Footnote 9**.

3.3.1.1 Correctness of the Discrete-Log.

In the event of successful forking, R_1 obtains two (related) sets of side-outputs σ_0 and σ_1 , where σ_i (for $i = 1, 2$) is of the form $(\hat{b}_i, c_i, r_i, \beta_i, d_i)$. It aborts in the event that $\beta_1 = \beta_0 = 1$ ($\text{abort}_{1,3}$). In

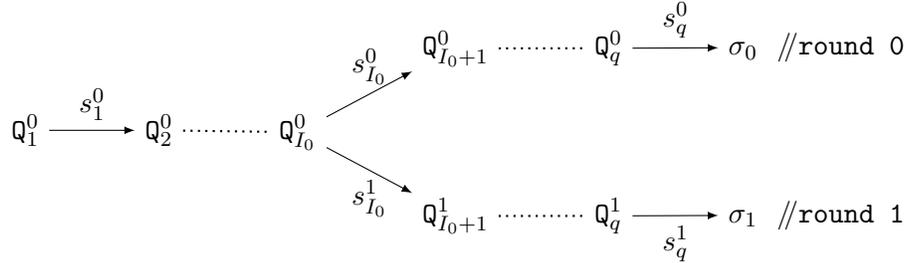


Figure 3.3: Successful forking by R_1 . $Q_{I_0}^0$ denotes the target query $G(\hat{\text{id}}, \hat{A}, \hat{m})$.

the rest of the cases, we claim that R_1 ends up with a system of two congruences in two unknowns $\{\hat{a}, \alpha\}$. In the following discussion, let \hat{a} denote $\log g \hat{A}_1 = \log g \hat{A}_0$.

- (i) $(\beta_0 = 1) \wedge (\beta_1 = 0)$: In this case, $\hat{R}_0 = g^{r_0}$ while \hat{R}_1 is of the form $g^{r_1 \alpha}$. As a result, we have $\{\hat{b}_0 = \hat{a} + (r_0 + zc_0)d_0, \hat{b}_1 = \hat{a} + (r_1 \alpha + zc_1)d_1\}$ —a system of two congruences in the two unknowns $\{\hat{a}, \alpha\}$. α can be solved for as shown below.

$$\alpha = \frac{z(c_0 d_0 - c_1 d_1) + r_0 d_0 - (\hat{b}_0 - \hat{b}_1)}{r_1 d_1} \quad (3.6)$$

- (ii) $(\beta_0 = 0) \wedge (\beta_1 = 1)$: In this case, \hat{R}_0 is of the form $g^{r_0 \alpha}$ while $\hat{R}_1 = g^{r_1}$. As a result, we have $\{\hat{b}_0 = \hat{a} + (r_0 \alpha + zc_0)d_0, \hat{b}_1 = \hat{a} + (r_1 + zc_1)d_1\}$. α can be solved for as shown below.

$$\alpha = \frac{z(c_1 d_1 - c_0 d_0) + r_1 d_1 - (\hat{b}_1 - \hat{b}_0)}{r_0 d_0} \quad (3.7)$$

- (iii) $(\beta_0 = \beta_1 = 0)$: In this case, \hat{R}_0 is of the form $g^{r_0 \alpha}$ and \hat{R}_1 is also of the form $g^{r_1 \alpha}$. As a result, we have $\{\hat{b}_0 = \hat{a} + (r_0 \alpha + zc_0)d_0, \hat{b}_1 = \hat{a} + (r_1 \alpha + zc_1)d_1\}$. α can be solved for as shown below.

$$\alpha = \frac{(\hat{b}_0 - \hat{b}_1) - z(c_0 d_0 - c_1 d_1)}{(r_0 d_0 - r_1 d_1)} \quad (3.8)$$

Notice that (3.6), (3.7) and (3.8) is precisely what R_1 outputs in **Algorithm 4**.

Remark 13. The equations (3.6), (3.7) and (3.8) hold even if $\hat{R}_1 = \hat{R}_0$ (and consequently $r_j = r_i$ and $c_j = c_i$). Note that this can happen if the adversary makes the random oracle query $H(\hat{R}_0, \hat{\text{id}})$

before the query $G(\hat{\text{id}}, \hat{A}, \hat{m})$ ($H < G$) in round 0. Hence, the order in which A makes the aforementioned random oracle queries is not relevant.

3.3.1.2 Analysis

The probability analysis is governed by the three events $\text{abort}_{1,1}$, $\text{abort}_{1,2}$ and $\text{abort}_{1,3}$. First, let's focus on the probability with which the wrapper Y successfully produces an output—the accepting probability acc_Y .

The accepting probability. Y aborts the simulation only when A does an extract query on an identity from \mathbb{I}_s , *i.e.* an identity with $\beta = 0$. Therefore, Y *does not* abort if all the extract queries correspond to identities from \mathbb{I}_ϵ , and consequently $\mathbb{P}(\neg\text{abort}_{1,1}) = (1 - \delta)^{q_\epsilon}$. Y accepts if the adversary produces a valid (non-trivial) forgery at the end of a *successful* simulation. Therefore $\text{acc}_Y \geq (1 - \delta)^{q_\epsilon} \epsilon$.

The reduction is successful in the event that neither $\text{abort}_{1,2}$ and $\text{abort}_{1,3}$ occurs. The first of the aborts ($\text{abort}_{1,2}$) pertains to the GF Algorithm: R_1 aborts in the event that the forking ended up being a failure. On applying the GF Lemma (**Lemma 4**) with $\text{acc} = \text{acc}_Y$, $|\mathbb{S}| = p$ and $q = q_G$, we get the following lower bound:

$$\mathbb{P}(\neg\text{abort}_{1,2}) \geq (1 - \delta)^{q_\epsilon} \epsilon \cdot \left(\frac{(1 - \delta)^{q_\epsilon} \epsilon}{q_G} - \frac{1}{p} \right).$$

The probability of event $\text{abort}_{1,3}$, on the other hand, is the same as that with which $(\beta_i = 1) \wedge (\beta_j = 1)$, *i.e.* $\mathbb{P}(\text{abort}_{1,3} \mid \neg\text{abort}_{1,2}) = (1 - \delta)^2$. On putting it all together, we get

$$\begin{aligned} \epsilon_1 &= \mathbb{P}(\neg\text{abort}_{1,3} \wedge \neg\text{abort}_{1,2}) \\ &\geq (1 - (1 - \delta)^2) \cdot (1 - \delta)^{q_\epsilon} \epsilon \cdot \left(\frac{(1 - \delta)^{q_\epsilon} \epsilon}{q_G} - \frac{1}{p} \right) \\ &= (2\delta - \delta^2) \cdot (1 - \delta)^{q_\epsilon} \epsilon \cdot \left(\frac{(1 - \delta)^{q_\epsilon} \epsilon}{q_G} - \frac{1}{p} \right) \end{aligned} \quad (3.9)$$

Assuming $p \gg 1$, (3.9) attains maximum value at the point $\delta = \left(1 - \sqrt{q_\epsilon / (q_\epsilon + 1)}\right)$, at which

$$\epsilon_1 \geq \frac{\epsilon^2}{\exp(1) q_G q_\epsilon}.$$

Here, \exp is the base of natural logarithm.

Remark 14. The above reduction is tighter than the original reduction B_1 given in [GG09]. This can be attributed to two reasons: *i*) R_1 uses the GF Algorithm F_W instead of the MF Algorithm M_W1 ; and *ii*) B_1 in [GG09] randomly chooses one of the identities involved in the G-oracle query as the target identity (refer to §3.2.2.1) which contributes a factor of q_G^2 to the degradation in B_1 . By contrast, we apply Coron’s technique in R_1 to partition the identity space in an optimal way.

Time complexity. If τ is the time taken for an exponentiation in \mathbb{G} , then the time taken by R_1 is $t_1 \leq t + 2(q_\varepsilon + 3q_s)\tau$. It takes at most one exponentiation for answering the extract query and three exponentiations for answering the signature query. This contributes the $(q_\varepsilon + 3q_s)\tau$ factor in the running time. The factor of two comes from the forking algorithm, since it involves running the adversary twice.

3.3.2 Reduction R_2

Let $\Delta := (\mathbb{G}, p, g, g^\alpha)$ be the given DLP instance. The reduction involves invoking the MF Algorithm on the wrapper W as shown in **Algorithm 5**. As a result, it obtains a set of two congruences in two unknowns and solves for α . It can be verified that R_2 indeed returns the correct solution to the DLP instance. The design of the wrapper W follows.

Algorithm 5 Reduction $R_2(\Delta)$

```

Set mpk :=  $\Delta$ 
 $(op, \{\sigma_0, \sigma_1\}) \stackrel{\$}{\leftarrow} M_W1(mpk)$ 
if  $(op = 0)$  then 0
Parse  $\sigma_i$  as  $(\hat{b}_i, c_i, d_i)$ ; let  $d$  denote  $d_1 = d_0$ 
return  $(\hat{b}_0 - \hat{b}_1)/(d_0(c_0 - c_1))$ 

```

The Wrapper

Suppose that $q := q_H + q_G$ and $\mathbb{S} := \mathbb{Z}_p$. W takes as input the master public key mpk and $\{s_1, \dots, s_q\}$. It returns a triple (I, J, σ) where J (resp. I) is the target H-index (resp. G-index) and σ is the side-output. In order to track the index of the current random oracle query, W maintains a counter ℓ , initially set to 1. It also maintains a table \mathcal{L}_H (resp. \mathcal{L}_G) to manage the random oracle H (resp. G). W initiates the EU-ID-CMA game by passing mpk as the challenge master public key to the adversary A . The queries by A are handled as per the following specifications.

(a) **Random oracle query**, $H(R, \text{id})$: \mathcal{L}_H contains tuples of the form

$$\langle R, \text{id}, c, \ell, y \rangle \in \mathbb{G} \times \{0, 1\}^* \times \mathbb{Z}_p \times \mathbb{Z}^+ \times \mathbb{Z}_p \cup \{\perp\}.$$

Here, (R, id) is the query to the H-oracle with c being the corresponding output. The index of the query is stored in the ℓ -field. Finally, the y -field stores either (a component of) the secret key for id , or a ' \perp ' in case the field is invalid. $H(R, \text{id})$ is fresh if there exists no tuple $\langle R_i, \text{id}_i, c_i, \ell_i, y_i \rangle$ in \mathcal{L}_H such that $(R_i = R) \wedge (\text{id}_i = \text{id})$. If such a tuple exists, then the oracle has to return c_i as the output. A fresh, explicit, H-oracle query is handled as follows: *i*) return $c := s_\ell$ as the output, and *ii*) add $\langle R, \text{id}, c, \ell, \perp \rangle$ to \mathcal{L}_H and increment ℓ by one.

(b) **Random oracle query**, $G(\text{id}, A, m)$: \mathcal{L}_G contains tuples of the form

$$\langle \text{id}, A, m, d, \ell \rangle \in \{0, 1\}^* \times \mathbb{G} \times \{0, 1\}^* \times \mathbb{Z}_p \times \mathbb{Z}^+.$$

Here, (id, A, m) is the query to the G-oracle with d being the corresponding output. The index of the query is stored in the ℓ -field. Therefore, a random oracle query $G(\text{id}, A, m)$ is fresh if there exists no tuple $\langle \text{id}_i, A_i, m_i, d_i \rangle$, in \mathcal{L}_G such that $(\text{id}_i = \text{id}) \wedge (A_i = A) \wedge (m_i = m)$. If such a tuple exists, then the oracle has to return d_i as the output. A fresh, explicit, G-oracle query is handled as follows: *i*) return $d := s_\ell$ as the output, and *ii*) add $\langle \text{id}, A, m, d, \ell \rangle$ to \mathcal{L}_G and increment ℓ by one.

(c) **Extract query**, $O_E(\text{id})$: Since the master secret key α is unknown to W , it has to carefully program the H-oracle in order to generate the user secret key usk .

- (i) If there exists a tuple $\langle R_i, \text{id}_i, c_i, \ell_i, y_i \rangle$ in \mathcal{L}_H such that $(\text{id}_i = \text{id}) \wedge (y_i \neq \perp)$, W returns $\text{usk} := (y_i, R_i)$ as the secret key.
- (ii) Otherwise, W chooses $y \xleftarrow{U} \mathbb{Z}_p$, sets $c := s_\ell$ and $R := (g^\alpha)^{-c} g^y$. It then adds $\langle R, \text{id}, c, \ell, y \rangle$ ¹¹ to \mathcal{L}_H and increments ℓ by one (an implicit H-oracle query). Finally, it returns $\text{usk} := (y, R)$ as the secret key.

(d) **Signature query**, $O_s(\text{id}, m)$: The signature queries are answered by first generating usk (by querying with O_E on id), followed by invoking S .

¹¹In the unlikely event of there already existing a tuple $\langle R_i, \text{id}_i, c_i, \ell_i, \perp \rangle$ in \mathcal{L}_H with $(R_i = R) \wedge (\text{id}_i = \text{id}) \wedge (c_i = c)$, W will simply increment ℓ and repeat step (ii).

- (i) If there exists a tuple $\langle R_i, \mathbf{id}_i, c_i, \ell_i, y_i \rangle$ in \mathfrak{L}_H such that $(\mathbf{id}_i = \mathbf{id}) \wedge (y_i \neq \perp)$, then $\mathbf{usk} = (y_i, R_i)$. W now uses the knowledge of \mathbf{usk} to run S and returns the signature.
- (ii) Otherwise, W generates \mathbf{usk} as in step (ii) of **Extract query** and runs S to return the signature.

At the end of the simulation, a successful adversary outputs a valid forgery $\hat{\sigma} := (\hat{A}, \hat{b}, \hat{R})$ on a $(\hat{\mathbf{id}}, \hat{m})$. Let $\langle \mathbf{id}_j, R_j, c_j, \ell_j, y_j \rangle$ be the tuple in \mathfrak{L}_H that corresponds to the target H-query. Similarly, let $\langle m_i, A_i, c_i, d_i, \ell_i \rangle$ be the tuple in \mathfrak{L}_G that corresponds to the target G-query. W returns $(\ell_i, \ell_j, (\hat{b}, c_j, d_i))$ as its own output. Note that the side-output σ consists of (\hat{b}, c_j, d_i) .

Structure of the forgery. Recall that the signature queries are answered by doing an extract query on the identity followed by calling S . Therefore, the resultant secret keys are of the form $\mathbf{usk} = (y, R)$, where $R = (g^\alpha)^{-c}g^y$ and we have $r = -\alpha c + y$. If a forgery is produced using the same R as given by R_2 as part of the signature query on \mathbf{id} , then b will be of the form $b = a + (-\alpha c + y + \alpha c)d = a + yd$. Therefore, it will not contain the solution to the DLP challenge α , and such forgeries are of no use to R_2 . But the event $\neg E$ guarantees that A does not forge using an R which was given as part of the signature query on \mathbf{id} and hence, for the forgery to be valid b will necessarily be of the form $b = a + (r + \alpha c)d$.

3.3.2.1 Correctness of the Discrete-Log.

In the event of successful forking, R_2 obtains two (related) sets of side-outputs $\{\sigma_0, \sigma_1\}$, where σ_i (for $i = 0, 1$) is of the form (\hat{b}_i, c_i, d_i) . Let \hat{a} denote $\log g \hat{A}_1 = \log g \hat{A}_0$; let \hat{r} denote $\log g \hat{R}_1 =$

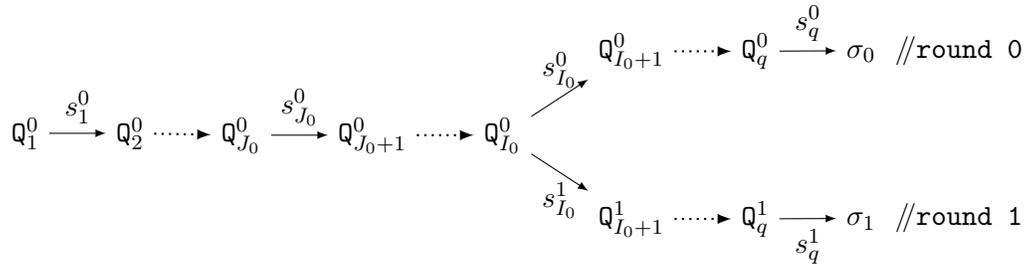


Figure 3.4: Successful forking by R_2 . $Q_{J_0}^0$ denotes the target G-query $G(\hat{\mathbf{id}}, \hat{A}, \hat{m})$ while $Q_{I_0}^0$ denotes the target H-query $H(\hat{R}, \hat{\mathbf{id}})$.

$\log g \hat{R}_0$; and let d denote $d_1 = d_0$. Since the multiple-forking was successful, we have: $\{\hat{b}_0 =$

$\hat{a} + (\hat{r} + \alpha c_0)d, \hat{b}_1 = \hat{a} + (\hat{r} + \alpha c_1)d$ —a system of two congruences in the two (effective) unknowns $\{\hat{a} + \hat{r}d, \alpha\}$. α can be solved for by using the expression given below.

$$\alpha := (\hat{b}_0 - \hat{b}_1)/(d(c_0 - c_1)) \quad (3.10)$$

Notice that (3.10) is precisely what R_2 outputs in **Algorithm 5**.

3.3.2.2 Analysis

Since there is no abort involved in the simulation of the protocol, we may conclude that the accepting probability of W is the same as the advantage of the adversary, *i.e.* $acc_W = \epsilon$. The probability of success of the reduction R_2 is computed by using MF Lemma (**Lemma 6**) with $q := q_H + q_G$, $|\mathbb{S}| := p$ and $n = 1$.¹² Hence, we get $\epsilon_2 = O(\epsilon^2/(q_H + q_G)^2)$.

Time complexity. Drawing analogy from the analysis of time complexity of R_1 , the time taken by R_2 is easily seen to be bounded by $t_2 \leq t + 2(2q_\epsilon + 3q_s)\tau$.

Remark 15. R_2 is similar in some aspects to the (incomplete) reduction B_2 in [GG09]. However, a *major* difference is that R_2 uses the MF Algorithm M_W1 instead of M_W3 to solve the DLP challenge. Therefore, only one forking is involved leading to a much tighter reduction than B_2 .

3.3.3 Reduction R_3

Let $\Delta := (\mathbb{G}, p, g, g^\alpha)$ be the given DLP instance. As in R_2 , the reduction involves invoking the MF Algorithm on the wrapper W . The only difference is that MF Algorithm is run for $n = 3$. As a result, it obtains a set of four congruences in four unknowns and solves for α . It can be verified that R_3 indeed returns the correct solution to the DLP instance.

¹² In the analysis of B_2 in [GG09], q was assumed to be $q_H \cdot q_G$. However, q actually denotes the size of the set of responses to the random oracle queries involved in the replay attack. As both H and G-oracle is involved in the replay attack in B_2 , the size of the set is $q_H + q_G$ rather than $q_H \cdot q_G$.

Algorithm 6 Reduction $R_3(\Delta)$ Set $\text{mpk} := \Delta$ $(\text{op}, \{\sigma_0, \sigma_1, \sigma_2, \sigma_3\}) \xleftarrow{\$} M_{W3}(\text{mpk})$ **if** $(\text{op} = 0)$ **then** 0Parse σ_i as (\hat{b}_i, c_i, d_i) .**return** $((\hat{b}_0 - \hat{b}_1)(d_2 - d_3) - (\hat{b}_2 - \hat{b}_3)(d_0 - d_1)) / (c_0 - c_1)(d_0 - d_1)(d_2 - d_3)$ **3.3.3.1 Correctness of the discrete-log.**

In the event of successful forking, R_3 obtains four (related) sets of side-outputs $\{\sigma_0, \sigma_1, \sigma_2, \sigma_3\}$, where σ_i (for $i = 0, \dots, 3$) is of the form (\hat{b}_i, c_i, d_i) . Let \hat{a}_0 (resp. \hat{a}_2) denote $\log g \hat{A}_1 = \log g \hat{A}_0$

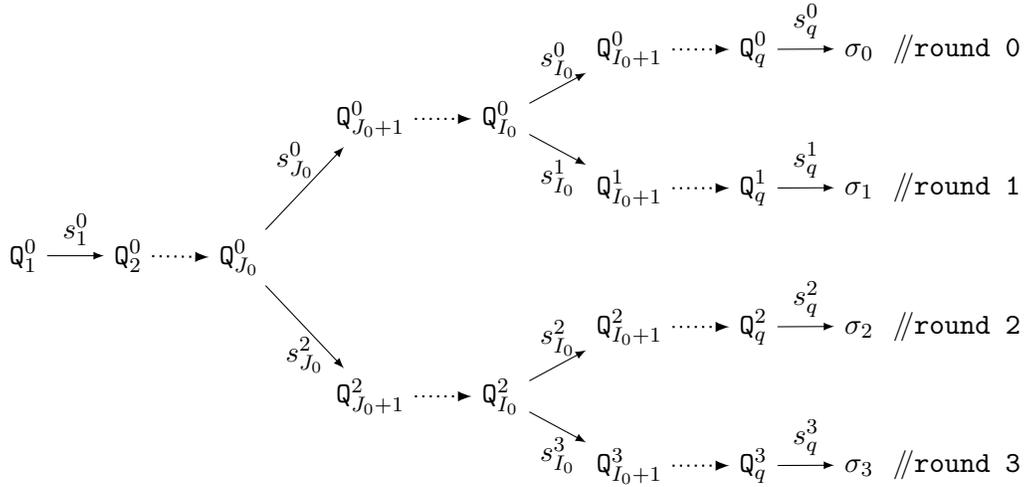


Figure 3.5: Successful multiple forkings by R_3 . $Q_{J_0}^0$ denotes the target H-query $H(\hat{R}, \hat{\text{id}})$; $Q_{I_0}^0$ (resp. $Q_{I_0}^2$) denotes the target G-query $G(\hat{\text{id}}, \hat{A}_0, \hat{m})$ (resp. $G(\hat{\text{id}}, \hat{A}_2, \hat{m})$).

(resp. $\log g \hat{A}_2 = \log g \hat{A}_3$); let \hat{r} denote $\log g \hat{R}_3 = \log g \hat{R}_2 = \log g \hat{R}_1 = \log g \hat{R}_0$. Since the multiple forkings were successful, we have:

$$\{\hat{b}_0 = \hat{a}_0 + (\hat{r} + \alpha c_0)d_0, \hat{b}_1 = \hat{a}_0 + (\hat{r} + \alpha c_0)d_1, \hat{b}_2 = \hat{a}_2 + (\hat{r} + \alpha c_2)d_2, \hat{b}_3 = \hat{a}_2 + (\hat{r} + \alpha c_2)d_3\} \quad (3.11)$$

We claim that (3.11) forms a system of four congruences in the four unknowns $\{\hat{r}, \hat{a}_0, \hat{a}_2, \alpha\}$. α can be solved for by using the expression given below.

$$\alpha = \frac{(\hat{b}_0 - \hat{b}_1)(d_2 - d_3) - (\hat{b}_2 - \hat{b}_3)(d_0 - d_1)}{(c_0 - c_1)(d_0 - d_1)(d_2 - d_3)} \quad (3.12)$$

Notice that (3.12) is precisely what R_3 outputs in **Algorithm 6**.

3.3.3.2 Analysis

Since there is no abort involved in the simulation of the protocol, we may conclude that the accepting probability of W is the same as the advantage of the adversary, *i.e.* $acc_W = \epsilon$. The probability of success of the reduction R_3 is computed by using MF Lemma (**Lemma 6**) with $q := q_H + q_G$, $|\mathcal{S}| := p$ and $n = 3$. Hence, we have $\epsilon_3 = O(\epsilon^4 / (q_H + q_G)^6)$.

3.3.4 A Comparison with the Original Reduction.

Recall that we replaced the reduction B_1 in the original security argument with the new reduction R_1 . Likewise, B_2 was replaced with the two reductions R_2 and R_3 . The resulting effect on tightness is tabulated below. The security degradation involved in original B_1 is of the order $O(q_G^3)$. In comparison, R_1 incurs a degradation of order $O(q_G q_\epsilon)$ which is much lower than that of B_1 . Note that $q_G \gg q_\epsilon$, *i.e.* the bound on the number of random oracle queries is much greater than the bound on the number of extract queries. For example, for 80-bit security one usually assumes $q_G \approx 2^{60}$ while $q_\epsilon \approx 2^{30}$. The degradation involved in the original B_2 would be of the order of $O((q_G + q_H)^6)$ (as pointed out in **Footnote 12**). In comparison, the security degradation involved in R_2 and R_3 is of order $O((q_H + q_G)^2)$ and $O((q_H + q_G)^6)$ respectively. Thus, the effective degradation¹³ is *still* of the order $O((q_H + q_G)^6)$.

Original reductions [GG09]	B_1	B_2	
Degradation	$O(q_G^3)$	$O((q_G + q_H)^6)$	
Our new reductions	R_1	R_2	R_3
Degradation	$O(q_G q_\epsilon)$	$O((q_H + q_G)^2)$	$O((q_H + q_G)^6)$

Table 3.1: A comparison of degradation in the original [GG09] and the new security argument.

¹³The effective degradation of a security argument involving multiple reductions from the same hard problem (as in the case of GG-IBS) is equal to the degradation of the reduction that incurs the worst security degradation.

Chapter 4

Galindo-Garcia IBS, Improved

4.1 Introduction

Even though we had managed to fix the security argument for GG-IBS in the previous chapter, the security bound that we ended up with in **Theorem 2** is still quite loose. The loss of tightness is *inherited* from the MF Algorithm (**Algorithm 2**) which is used to launch nested replay attack on the GG-IBS adversary. To be precise, the use of M_W3 in reduction R_3 leads to an effective degradation of $O(q^6)$ (where q denotes $q_H + q_G$). In this chapter, we contemplate a better security bound for GG-IBS. To that end, we introduce two notions pertaining to simulation of random oracles: “dependency” and “independency”. The notion of independency follows naturally for GG-IBS; dependency, on the other hand, has to be induced by modifying the construction (to be precise, the hash functions H and G) of the protocol in a clever manner. It turns out that the two notions can be applied in conjunction and this leads to the nested replay attack being launched far more *effectively* than specified in the MF Algorithm. As a result, the effective degradation is reduced to $O(q^3)$. The non-trivial aspect is to leverage these two notions in the security argument.

In the previous chapter, we had concluded that the security bound for GG-IBS is quite loose. The loss of tightness of $O(q^6)$ (where q denotes $q_H + q_G$) is *inherited*, primarily, from the MF Algorithm M_W3 which is used for launching a nested oracle replay attack on the GG-IBS adversary in reduction R_3 . In this chapter, we address the question of whether the nested oracle replay attack can be launched more effectively than specified in the MF Algorithm. To that end, we exploit the notions of “dependency” and “independency”–(in)dependency, in conjunction–among the target indices of the two random oracles H and G . The notion of independency follows naturally for GG-IBS; dependency, on the other hand, has to be induced by modifying the construction (to be

precise, the hash functions H and G) of the protocol in a clever manner. The non-trivial aspect is to leverage these two notions in the security argument.

4.2 Degradation: A Closer Look

Let's return to the primary source of degradation for GG-IBS: reduction R_3 . For $i = \{0, \dots, 3\}$, let J_i (resp. I_i) denote the target H -index (resp. G -index) for round i of simulation. (see **Figure 4.1**). The condition for the success of the forkings can be derived from (2.23) to be $E : B \wedge C_0 \wedge C_2 \wedge D_2$,

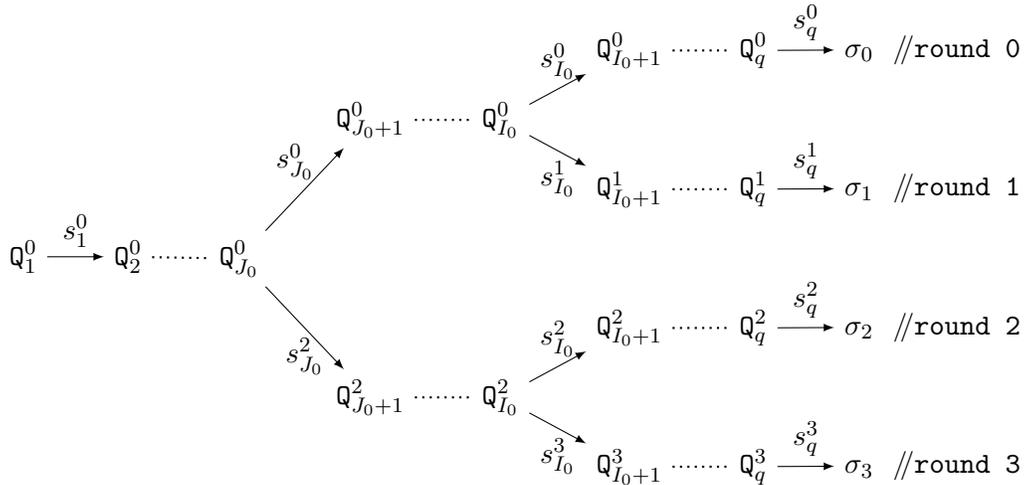


Figure 4.1: Successful Multiple-Forking for R_3 . $Q_{J_0}^0$ denotes the target H -query $H(\hat{R}, \hat{id})$ for round 0; $Q_{I_0}^0$ (resp $Q_{I_0}^2$) denotes the target G -query $G(\hat{id}, \hat{A}_0, \hat{m})$ (resp $G(\hat{id}, \hat{A}_1, \hat{m})$) for round 0 (resp round 2).

with the *core* condition of

$$(I_3, J_3) = (I_2, J_2) = (I_1, J_1) = (I_0, J_0). \quad (4.1)$$

Each of the equality checks, loosely speaking, contributes to a factor of $O(q^2)$ leading to the overall degradation of $O(q^6)$. Our objective of bringing down the overall degradation relies on relaxing the core condition. This, in turn, is owed to the following observations.

Observation 5 (Independency of I_2 and I_0). *It is not necessary for the target indices of the H oracle to be the same in round 0 and round 2 (i.e., I_2 needn't match I_0).*

In order to see this, let's return to the system of congruences in (3.11), *i.e.*,

$$\{\hat{b}_0 = \hat{a}_0 + (\hat{r} + \alpha c_0)d_0, \hat{b}_1 = \hat{a}_0 + (\hat{r} + \alpha c_0)d_1, \hat{b}_2 = \hat{a}_2 + (\hat{r} + \alpha c_2)d_2, \hat{b}_3 = \hat{a}_2 + (\hat{r} + \alpha c_2)d_3\} \quad (4.2)$$

What we have is a system of four congruences in the four unknowns $\{\hat{r}, \hat{a}_0, \hat{a}_2, \alpha\}$ with α being the solution to the DLP. From the structure of the hash function H , it follows that the index I corresponds to the unknown \hat{a} . The process of solving for α starts by eliminating the unknown \hat{a}_0 (resp. \hat{a}_2) from $\{\hat{b}_0, \hat{b}_1\}$ (resp. $\{\hat{b}_2, \hat{b}_3\}$). What is *necessary* at this point is that the I indices must match for **round 0** and **round 1** (resp. **round 2** and **round 3**). However, eliminating \hat{a}_0 from $\{\hat{b}_0, \hat{b}_1\}$ is not affected by the pair $\{\hat{b}_2, \hat{b}_3\}$ and vice versa. Hence, from the point of view of the reduction, it doesn't make any difference whether we relax the condition to accommodate independency—the system of congruences one ends up with is *exactly* the same as in (4.2). In fact, the reduction is unlikely to achieve anything by restricting the indices. Thus, incorporating independency simplifies the condition in (4.1) to

$$(I_3, J_3) = (I_2, J_2) \wedge (I_1, J_1) = (I_0, J_0) \wedge (J_2 = J_0). \quad (4.3)$$

Remark 16. The notion of independency can be better appreciated if we visualise the process of multiple-forking in terms of congruences and unknowns. At a high level, what the reduction algorithm R_3 secures from the MF Algorithm is a set of four congruences in four unknowns. One of these unknowns is the solution to the DLP. The MF Algorithm needs to ensure that the congruences are linearly independent of each other with a certain non-negligible probability. **Observation 5** can then be restated as: even if the condition on the I indices is relaxed, we *still* end up with a system of four congruences in four unknowns.

Observation 6 (Dependency between H and G). *It is possible to modify GG-IBS, the structure of the hash function G to be precise, such that $I_1 = I_0$ implies $J_1 = J_0$ (and similarly $I_3 = I_2$ implies $J_3 = J_2$) with a non-negligible probability.*

Recall that the hash functions in the construction of GG-IBS are of the form $H(\text{id}, R)$ and $G(\text{id}, A, m)$. Suppose that we modify the structure of G to $G(\text{id}, A, m, c)$ where $c := H(\text{id}, R)$. Let's consider **round 1** of simulation for the “modified” GG-IBS initiated by a forking at I_0 as in reduction R_3 . Suppose that the adversary is successful and, in addition, the target H -index for the **round 1** matches with that in **round 0** (*i.e.*, $I_1 = I_0$). We claim: due to the *binding* between the

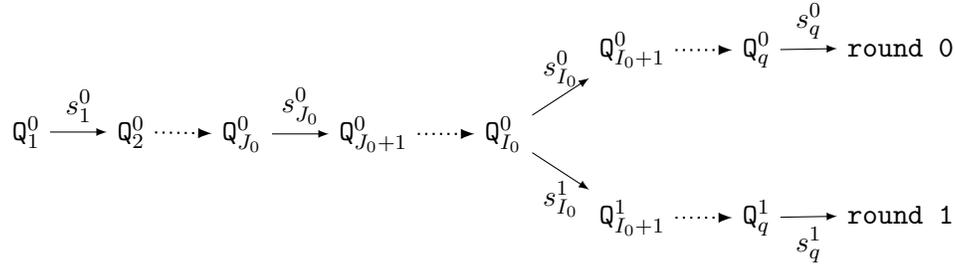


Figure 4.2: Oracle replay attack after setting up dependency between the hash functions. $Q_{J_0}^0$ denotes the target H-query $H(\hat{x})$ while $Q_{I_0}^0$ corresponds to the target $G(\hat{y}, \hat{c})$ where $\hat{c} = H(\hat{x})$.

hash functions that we have introduced, the target G-indices for the two rounds also *have to* match. The advantage with which an adversary can forge a signature having violated this condition is, in fact, negligible (see **Claim 2.1**). Hence, $(I_1 = I_0)$ *implies* $(J_1 = J_0)$. We say that the random oracle G is “dependent” on the random oracle H (denoted by $H \prec G$) over these two rounds¹. The property holds for round 2 and round 3 as well (*i.e.* $(I_3 = I_2)$ implies $(J_3 = J_2)$). As a result, the condition in (4.3) is further simplified to

$$(I_3 = I_2) \wedge (I_1 = I_0) \wedge (J_2 = J_0). \quad (4.4)$$

The introduction of binding, in fact, brings more to the table. In a particular round of simulation, to forge a signature, an adversary (except with a negligible probability of guessing) *has* to make the target random oracle queries in the order $H < G$ (see **Claim 2.1**). For example, it follows in round 0 that $(1 \leq J_0 < I_0 \leq q)$. Thus, as a result of the binding, the logical order is *imposed* on the adversary. That brings us to the formal definition of the notion of random-oracle dependency.

Definition 14 (Random-Oracle Dependency). Consider the oracle replay attack in the context of a cryptographic protocol that employs two hash functions H_1 and H_2 modelled as random oracles. Let J (resp. I) denote the target H_1 -index (resp. H_2 -index) for the first round of simulation of the protocol. Also, let J' (resp. I') denote the target H_1 -index (resp. H_2 -index) for the second round of simulation that was initiated by a forking at I . Suppose that the adversary was successful in both the rounds. The random oracle H_2 is defined to be η -dependent on the random oracle H_1 on the target query (denoted by $H_1 \prec H_2$) if the following criteria are satisfied: *i*) $(1 \leq J < I \leq q)$ or, in

¹Intuitively, if a protocol uses two hash functions H_1 and H_2 in such a way that the input to H_2 is a *function* of the output of H_1 , then we say that the H_2 -call is *dependent* on the H_1 -call. This is possible only when there is an inherent order, a *hierarchy*, among the hash functions used in the construction of the protocol.

other words $H_1 < H_2$; and *ii*) $\mathbb{P}((J' \neq J) \mid (I' = I)) \leq \eta$.

Inducing dependency. Let's return to the random oracles in **Definition 14**. A natural way to set up the dependency $H_1 \prec H_2$ is by redefining the random oracle H_2 as a function

$$H_2 : \mathbb{D}_2 \times \mathbb{R}_1 \mapsto \mathbb{R}_2,$$

where the second input to the oracle H_2 is chosen as the output of the H_1 -oracle for an appropriate input. The trick is to choose the appropriate input. This, in turn, will require changes in the protocol, to be precise, changes in the hash functions corresponding to the random oracles.

The second criterion, in other words, requires $I' = I \implies J' = J$ with probability at least $(1 - \eta)$; for the criterion to hold with overwhelming probability, η should be negligible in n . A little more specifically, H_2 is said to be *fully-dependent* on H_1 if $\eta = 0$, *i.e.* $(I' = I) \implies (J' = J)$. But for most applications, it suffices that J be η -dependent on I for some η which is negligible.

Although dependency forces an order among the hash functions, the reverse may not be true—a logical order among the hash functions in the protocol does not necessarily translate into dependency between them. Hence, one may need to impose explicit dependency among the hash functions. A natural way to induce the dependency $H_1 \prec H_2$ is by making the input to H_2 a function of H_1 's output (like the binding we introduced for GG-IBS).

Next, we formally argue that the binding in GG-IBS does indeed translate into a dependency between the random oracles (see **Figure 4.3** for a detailed construction).

Claim 2.1 (Dependency for modified GG-IBS). *The random-oracle dependency of $H \prec G$ with $\eta := q(q - 1)/p$ applies to the modified GG-IBS.*

Argument. Consider an adversary A against the modified GG-IBS. Let's suppose that it produces a valid, non-trivial forgery $\hat{\sigma} = (\hat{b}, \hat{R}, \hat{A})$ on $(\hat{\text{id}}, \hat{m})$ at the end of one round of simulation. Let J (resp. I) denote the index of the target H-query $H(\hat{\text{id}}, \hat{R})$ (resp. G-query $G(\hat{m}, \hat{A}, c)$, where $c := H(\hat{\text{id}}, \hat{R})$) for the round. By using an argument *à la* that in **Footnote 2**, it can be established that $0 < I, J \leq q$. Now, there are two possibilities: $H < G$ and $G < H$. Let's look at the second of the cases. Since G takes c as an input, $G < H$ is plausible only if A correctly *anticipates* the value of c . Thus, for A to have a non-negligible advantage, it has to make the two target queries *and* in the order $H < G$. Simply put, the random oracles satisfy the first criterion.

Next, we simulate A again after forking at I . Let (I', J') be the target indices for this round. In

order to establish that the random oracles satisfy the second criterion, let's assume the contrary: $I' = I$ but $J' \neq J$. Since $I' = I$, A has to forge a signature $\hat{\sigma}' = (\hat{b}', \hat{R}', \hat{A})$ on (\hat{id}', \hat{m}) such that $H(\hat{id}', \hat{R}') = c$. This, in turn, is tantamount to a collision of the random function corresponding to H and, as per the birthday bound, the probability of this event is $q(q-1)/p$. Moreover, provided that $q \ll p$ the value is negligible². By implication, the second criterion is also satisfied. \square

More generally:

Claim 2.2 (Binding induces dependency). *Consider the hash functions (and the corresponding random oracles) described in **Definition 2.1**. Let q_1 denote the upper bound on the number of queries to the random oracle H_1 . In addition, let \mathbb{R}_1 denote the range of H_1 . Binding H_2 to H_1 (by making the input to H_2 a function of H_1 's output) induces a random-oracle dependency $H_1 \prec H_2$ with $\eta_b := q_1(q_1 - 1)/|\mathbb{R}_1|$.*

Argument. The line of argument is the same as in **Claim 2.1**. \square

Remark 17. Assuming q_1 to be polynomial and $|\mathbb{R}_1|$ to be exponential in n , the value of η_b is asymptotically negligible. Now, let's consider a concrete setting, say 80-bit security level. Assuming typical values of $q_1 := 2^{60}$ and $|\mathbb{R}_1| := 2^{80}$, the value can no longer be ignored. However, if we take into consideration the degradation in **Theorem 3**, we end up choosing $|\mathbb{R}_1| := 2^{260}$ rendering η_b negligible. Hence, we can safely assume full-dependency in subsequent discussions.

The consequences of (in)dependency. In a nutshell, our observations affect the security of GG-IBS in the following ways.

- (i) We saw that, once (in)dependency is taken into consideration, the condition for success of the multiple forkings is simplified to (4.4). Intuitively, this would bring down the degradation to $O(q^3)$.
- (ii) The incompleteness of B_2 stems from the fact that the simulator cannot really restrict the order in which the adversary makes the two target random oracle queries. In the previous chapter, the issue was fixed using the *two-reduction* strategy. By imposing the logical order $H < G$ through dependency, we can do away with the events F and $\neg F$ and hence the reduction R_2 . This leads to a simpler event structure.

²To be precise, q is polynomial in n whereas $1/p$ is a negligible function in n . Plus, the product of a polynomial and a non-negligible function is still non-negligible.

As we will see in the next section, we end up with a cleaner, tighter security argument for GG-IBS.

4.3 Galindo-Garcia IBS, Improved

We begin by describing, in detail, the GG-IBS scheme in which the binding required to induce the dependency $H \prec G$ has been set up.³ The construction is same as in **Figure 3.1** *save* for the structure of the hash function G , which is now a function of $c := H(\text{id}, R)$ (boxed in **Figure 4.3**). The binding that we have introduced is more refined⁴ than the one suggested in **Observation 6**.

Set-up, $G(1^n)$: Invoke the group generator G_{DL} (on 1^n) to obtain (\mathbb{G}, g, p) . Select $z \xleftarrow{U} \mathbb{Z}_p$ and set $Z = g^z$. Return z as the master secret key msk and $(\mathbb{G}, p, g, Z, H, G)$ as the master public key mpk , where H and G are hash functions

$$H : \{0, 1\}^* \times \mathbb{G} \mapsto \mathbb{Z}_p \quad \text{and} \quad G : \{0, 1\}^* \times \mathbb{G} \times \mathbb{Z}_p \mapsto \mathbb{Z}_p.$$

Key Extraction, $E(\text{id}, \text{msk})$: Select $r \xleftarrow{U} \mathbb{Z}_p$ and set $R := g^r$. Return $\text{usk} := (y, R) \in \mathbb{Z}_p \times \mathbb{G}$ as the user secret key, where

$$y := r + zc \quad \text{and} \quad c := H(\text{id}, R).$$

Signing, $S(\text{id}, m, \text{usk})$: Let $\text{usk} = (y, R)$ and $c = H(\text{id}, R)$. Select $a \xleftarrow{U} \mathbb{Z}_p$ and set $A := g^a$. Return $\sigma := (b, R, A) \in \mathbb{G} \times \mathbb{Z}_p \times \mathbb{G}$ as the signature, where

$$b := a + yd \quad \text{and} \quad \boxed{d := G(m, A, c)}.$$

Verification, $V(\sigma, \text{id}, m, \text{mpk})$: Let $\sigma = (b, R, A)$, $c := H(\text{id}, R)$ and $\boxed{d := G(m, A, c)}$. The signature is valid if

$$g^b = A(R \cdot Z^c)^d.$$

³A noteworthy observation is that, setting up the random-oracle dependency $G \prec H$ in GG-IBS allows more efficient reductions to DLP. However, this disturbs the “logical” order of the random oracles from the protocol’s perspective. In such a protocol, the PKG will have to issue private keys for each message to be signed, rendering it impractical.

⁴Recall that the suggestion in **Observation 6** was to set $d := G(\text{id}, A, m, c)$ where $c := H(R, \text{id})$. However, id is redundant here as it is anyways captured indirectly by c .

break

Figure 4.3: Galindo-Garcia IBS with binding.

4.3.1 Security Argument

Once the binding is in place, the adversary is bound to make the target queries in the logical order (except with a negligible probability, see **Claim 2.1**). This simplifies the event structure of the security argument to quite an extent and, accordingly, it consists of only two reductions R'_1 and R'_3 .

Theorem 3. *Let A be an $(\epsilon, q_\epsilon, q_H, q_G)$ -adversary against GG-IBS in the EU-ID-CMA model. If the hash functions H and G are modelled as random oracles, we can construct either*

(i) *Algorithm R'_1 that $\epsilon^2/(2 \exp(1)q_Gq_\epsilon)$ -breaks the DLP, or*

(ii) *Algorithm R'_3 that $\epsilon^4/64(q_H + q_G)^3$ -breaks the DLP.*

Here q_ϵ denotes the upper bound on the number of extract queries that A can make; q_H (resp. q_G) denotes the upper bound on the number of queries to the H -oracle (resp. G -oracle). \exp is the base of natural logarithm. (Plus, for simplicity, we have assumed both $1/p$ and $q(q-1)/p$ to be negligible).

Argument. Recall the events E and $\neg E$ that were defined in the security argument of GG-IBS. In the case of the event E (resp. $\neg E$) we give a reduction R'_1 (resp. R'_3) to the DLP. R'_1 can be regarded to be a simplified version of the reduction R_1 confined to handling the order $H < G$ (recall that R_1 is equipped to handle both the orders). As the gist of the reduction R'_1 is the same as in R_1 given in §3.3.3, we relegate it to **Appendix A.3**. As for R'_3 , the simulation of one round of the adversary, as well as the manner in which it is forked is the same as plotted in R_3 . However, we use the Rewinding Technique (see §2.3.2.2 and §2.5) instead of the MF Algorithm, as shown in **Figure 4.4**. The novelty lies in exploiting (in)dependency for the fruition of a better security bound. Hence, we focus on the probability analysis. \square

4.3.2 Analysis

The aim is to avoid the use of Extended Splitting Lemma in the analysis. This requires two iterations of the Splitting Lemma (**Lemma 1**) on different, but correlated, underlying sets. Consequently, we have two base notions of “good”: good and $\text{good}_{(1)}$. In addition, we have a higher notion of $\text{good}_{(3)}$ which we, ultimately, intend to bound.

Conventions. We stick to the conventions followed in the analysis of elementary oracle replay attack in §2.3.2.3. \mathbb{T} denotes the universe of random tapes participating in a single round of simulation of the adversary. This includes the internal coins of the adversary as well as the randomness from the random functions associated to the two random oracles, the signature oracle and the extract oracle. $\mathbb{T}_{(1)}$ denotes the universe of random tapes involved in two rounds of simulation resulting from the forking of the G-oracle. Thus, $\mathbb{T}_{(1)}$ is defined as

$$\mathbb{T}_{(1)} = \bigcup_{i=2}^q \mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}^2. \quad (4.5)$$

Here, $\mathbb{T}_{(i-)}$ denotes the set of all random tapes involved in the simulation before the adversary makes the query \mathbb{Q}_i^0 , whereas $\mathbb{T}_{(i+)}$, those after the query \mathbb{Q}_i^0 .⁵ Note that $\mathbb{T}_{(1)}$ consists of the partitions $(\mathbb{T}_{(2-)} \times \mathbb{T}_{(2+)}^2)$, $(\mathbb{T}_{(3-)} \times \mathbb{T}_{(3+)}^2)$ and $(\mathbb{T}_{(q-)} \times \mathbb{T}_{(q+)}^2)$. We use the shorthand $\mathbb{T}_{(1,j<)}$ (resp. $\mathbb{T}_{(1,j>)}$) to denote the union $\bigcup_{i=2}^j (\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}^2)$ (resp. $\bigcup_{i=j+1}^q (\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}^2)$). Thus, $\mathbb{T}_{(1)} = \mathbb{T}_{(1,j<)} \cup \mathbb{T}_{(1,j>)}$. Finally, $\mathbb{T}_{(3)}$ denotes the universe of random tapes for all four rounds of simulation⁶ with the adversary forked as per the illustration in **Figure 4.4**. We follow the same subscripting convention for individual tapes as well with the round is indicated in the superscript; e.g., \mathbb{T}_{i-}^0 denotes the random tape involved in the first round of simulation up till the adversary makes the query \mathbb{Q}_i^0 (see **Figure 4.4**).

Defining the notions of “good”. The notion of **good** is defined akin to the notion of “good” for the Schnorr signature scheme in §2.3.2.3: a tape in set \mathbb{T} is deemed to be **good** if it leads to the adversary successfully forging. Thus, by definition, at least ϵ fraction of the tapes in \mathbb{T} are **good**. However, the refined notion for **good** is a bit different as there are *two* random oracles in consideration. A tape in set \mathbb{T} is deemed to be **good**_(j,i) if it leads to the adversary forging

⁵Keep in mind that the set $\mathbb{T}_{(i-)}$ can be further represented by the Cartesian product of the two sets $\mathbb{T}_{(j-)}$ and $\mathbb{T}_{(ji)}$ as shown below.

$$\mathbb{T}_{(i-)} = \bigcup_{i=2}^q (\mathbb{T}_{(j-)} \times \mathbb{T}_{(ji)}) \times \mathbb{T}_{(i+)}^2$$

Here $\mathbb{T}_{(j-)}$ denotes the set of all random tapes involved in the simulation before the adversary makes the query \mathbb{Q}_j^0 , whereas $\mathbb{T}_{(ji)}$ denotes those from the query \mathbb{Q}_j^0 to the query \mathbb{Q}_i^0 . In fact, this representation is in agreement with the **Figure 4.4** more than the one in (4.5).

⁶It can be worked out that the universe of tapes for three forkings is

$$\mathbb{T}_{(3)} = \bigcup_{j=1}^{q-1} \left(\mathbb{T}_{(j-)} \times \left(\bigcup_{i=j+1}^q (\mathbb{T}_{(ji)} \times \mathbb{T}_{(i+)}^2) \right)^2 \right).$$

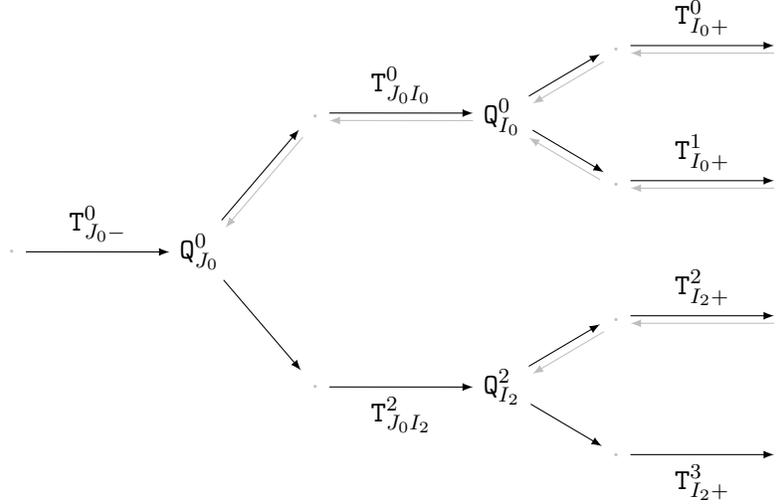


Figure 4.4: The tapes involved in the nested replay attack involving *three* forkings (carried out using the Rewinding Technique).

successfully and with a target H-index of j and a target G-index of i . In addition, $\mathbf{good}_{(j,*)}$ denotes a tape which leads to the adversary forging successfully with a target H-index i and *any* target G-index. In an analogous manner, we define the notion of $\mathbf{good}_{(*,i)}$. The key to an improved analysis requires circumventing the notion of $\mathbf{good}_{(j,i)}$ through the notion of $\mathbf{good}_{(1)}$ and, its refinements, $\mathbf{good}_{(1,j,*)}$ and $\mathbf{good}_{(1,*,i)}$. This, in turn, is made possible by our observations in §4.2—*i.e.*, through (in)dependency.

The notion of $\mathbf{good}_{(1)}$ is also related to the “higher” notion of $\mathbf{good}_{(1)}$ that was defined for Schnorr Signature (in **Remark 6**). Recall that, for Schnorr Signature, the notion of $\mathbf{good}_{(1)}$ applies to tapes participating in two rounds of simulations associated with (the only) forking. However, in the case of GG-IBS, there are two random oracles in consideration. Thus, the notion of $\mathbf{good}_{(1)}$ could be applied to tapes participating in replay of either of these oracles. But, as we are concerned with the replay of the G-oracle in the first two rounds of the nested replay attack (see **Figure 4.4**), the underlying set for which the notion of $\mathbf{good}_{(1)}$ applies is defined to be $\mathbb{T}_{(1)}$. A triplet of tape segments $(\mathbb{T}_{I-}, (\mathbb{T}_{I+}, \mathbb{T}'_{I+})) \in \mathbb{T}_{(1)}$ is considered to be $\mathbf{good}_{(1)}$ if the tapes $\mathbb{T} = (\mathbb{T}_{I-}, \mathbb{T}_{I+})$ and $\mathbb{T}' = (\mathbb{T}_{I-}, \mathbb{T}'_{I+})$ are both $\mathbf{good}_{(*,I)}$. To paraphrase, the tape is $\mathbf{good}_{(1)}$ if the replay of the G-oracle using tapes \mathbb{T} and \mathbb{T}' is successful. We define the refined notions of $\mathbf{good}_{(1,j,i)}$, $\mathbf{good}_{(1,*,i)}$ and $\mathbf{good}_{(1,j,*)}$ the same way we had defined it for the notion of \mathbf{good} . Finally, a tape in $\mathbb{T}_{(3)}$ is $\mathbf{good}_{(3)}$ if the nested replay attack launched as in **Figure 4.4** using the tape is successful. Keeping these notions in mind, we proceed to the probability analysis of the nested replay attack.

Probability Analysis. Our objective is to lower bound the fraction of $\text{good}_{(3)}$ tapes in $\mathbb{T}_{(3)}$ in terms of the fraction of good tapes in \mathbb{T} . This is accomplished through three claims: **Claim 3.1** through **Claim 3.3**, which, in turn, requires two iterations of the Splitting Lemma (**Lemma 1**). The objective of the first iteration of the Splitting Lemma (in **Claim 3.1**) is to establish a lower bound on the fraction of $\text{good}_{(1)}$ tapes in $\mathbb{T}_{(1)}$, in terms of the fraction of good tapes in \mathbb{T} . This bound is, in turn, to be used in the second iteration (in **Claim 3.3**) to lower bound the fraction of $\text{good}_{(3)}$ tapes in $\mathbb{T}_{(3)}$. However, the second iteration is far more involved than the first iteration—we need an intermediate **Claim 3.2**. Nevertheless, the usage of the Splitting Lemma, on a high level, is (self) similar to that in the first iteration; it’s the underlying set, the associated probabilities and the notion of “good” that differs. Let’s focus on the first two rounds of simulation of the adversary (see **Figure 4.4**).

Claim 3.1. *At least $\epsilon^2/4q$ fraction of the tapes in $\mathbb{T}_{(1)}$ are $\text{good}_{(1)}$.*

Argument. The analysis, because of dependency, turns out to be similar to analysing the two rounds of the elementary replay attack *bar* the definition of the notion of “good”—we use the notion of $\text{good}_{(*,i)}$ in place of $\text{good}_{(i)}$. Let’s presume that ϵ_i fraction of the tapes in \mathbb{T} are $\text{good}_{(*,i)}$. Thus, we have

$$\begin{aligned} \sum_{i=1}^q \epsilon_i &= \sum_{i=2}^q \mathbb{P}(I = i \wedge J > 0) \\ &= \mathbb{P}(1 \leq J < I \leq q) = \epsilon \text{ (by definition)}. \end{aligned} \tag{4.6}$$

The notion of $\text{good}_{(*,i)}$, which we defined for the set \mathbb{T} , is adapted for the set $\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ through the split and join functions (see **§2.3.2.3**). A tape $(\mathbb{T}_{i-}, \mathbb{T}_{i+}) \in \mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ is deemed to be $\text{good}_{(*,i)}$ if its counterpart in \mathbb{T} is $\text{good}_{(*,i)}$. We denote the (sub)set of all such tapes in $\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ by \mathbb{V}_i . Also, suppose that at least β_i fraction of tapes in $\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}$ are $\text{better}_{(*,i)}$. On these underlying assumptions, the analysis proceeds in a manner (almost) similar to that of the elementary replay attack (using the Rewinding Technique) in **§2.3.2.3**. Thus, we may conclude that the first two rounds of the nested replay attack are successful with a probability of at least $\epsilon^2/4q$. Moreover, it follows from the definition (of $\text{good}_{(1)}$) that at least $\epsilon^2/4q$ fraction of tape segments in $\mathbb{T}_{(1)}$ are $\text{good}_{(1)}$. \square

The second iteration of the Splitting Lemma is a bit trickier than the first one. We have to prime the set $\mathbb{T}_{(1)}$ before actually applying the Splitting Lemma. The partitions of the set $\mathbb{T}_{(1)}$ come

into play. Let's focus on a particular partition $(\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}^2)$. It is not difficult to infer that the set of $\text{good}_{(1,*,i)}$ tapes in $\mathbb{T}_{(1)}$, in fact, all belong⁷ to the partition $(\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)}^2)$. These $\text{good}_{(1,*,i)}$ elements can be further partitioned, depending on their target H-index, into subsets containing the $\text{good}_{(1,j,i)}$ elements. The set of all such $\text{good}_{(1,j,i)}$ elements, across the partitions (of $\mathbb{T}_{(1)}$), form the $\text{good}_{(1,j,*)}$ elements of $\mathbb{T}_{(1)}$. Let δ_j denote the fraction of all such tapes in $\mathbb{T}_{(1)}$. By implication,

$$\begin{aligned} \sum_{j=1}^{q-1} \delta_j &= \sum_{j=1}^{q-1} \mathbb{P}(J_1 = J_0 = j \wedge I_1 = I_0 \wedge 1 < I_0 \leq q) \\ &= \mathbb{P}(J_1 = J_0 \wedge I_1 = I_0 \wedge 1 \leq J_0 < I_0 \leq q) \\ &\geq \epsilon^2/4q \text{ (using \textbf{Claim 3.1})} \end{aligned} \tag{4.7}$$

That brings us to our second claim.

Claim 3.2. *At least δ_j fraction of tapes in $\mathbb{T}_{(1,j>)}$ are $\text{good}_{(1,j,*)}$.*

Argument. Let's presume that the tape involved in the first two rounds of simulation (the replay of G-oracle) was $\text{good}_{(1,j,i)}$. Our argument relies on two observations: *i*) the candidate tapes for the replay of the H-oracle belong to $\mathbb{T}_{(1,j>)}$; and *ii*) a $\text{good}_{(1,j,*)}$ element *cannot* belong to $\mathbb{T}_{(1,j<)}$. The first of the observations follows by definition, whereas the second—a consequence of dependency between the two random oracles—requires some explanation. Let's assume the contrary: a tape is $\text{good}_{(1,j,*)}$ and it belongs to $\mathbb{T}_{(1,j<)}$. Consequently, it is $\text{good}_{(1,j,i)}$ for some $i < j$. However, the dependency $H \prec G$ means that an adversary *has to* make the target queries in the order $H < G$ which, in turn, implies $j < i$ (leading to a contradiction). Thus, the $\text{good}_{(1,j,*)}$ elements all belong to $\mathbb{T}_{(1,j>)}$.

The two observations, in conjunction with the fact that $\mathbb{T}_{(1,j>)}$ is a subset of the set $\mathbb{T}_{(1)}$ (of which at least δ_j fraction are $\text{good}_{(1,j,*)}$ elements) completes the proof. \square

Claim 3.3. *The fraction of $\text{good}_{(3)}$ tapes in $\mathbb{T}_{(3)}$ is at least $\epsilon^4/64q^3$.*

Argument. Let's assume that the tape involved in the first two rounds of simulation $\mathbb{T}_1^0 := (\mathbb{T}_{I_0-}^0, (\mathbb{T}_{I_0+}^0, \mathbb{T}_{I_0+}^1))$ is $\text{good}_{(1,J_0,I_0)}$. We further split the tape segment $\mathbb{T}_{I_0-}^0$ into two: $\mathbb{T}_{J_0-}^0$ and $\mathbb{T}_{J_0I_0}^0$. Here, $\mathbb{T}_{J_0-}^0$ is the tape involved in the simulation before the adversary makes the target query $\mathbb{Q}_{J_0}^0$, whereas $\mathbb{T}_{J_0I_0}^0$,

⁷Strictly speaking, we have to take into consideration the non-core conditions.

that from the query $\mathbb{Q}_{J_0}^0$ up to the query $\mathbb{Q}_{I_0}^0$ (see **Figure 4.4**). Thus, an alternative representation of the tape \mathbb{T}_1^0 is

$$\begin{aligned}\mathbb{T}_1^0 &:= (\mathbb{T}_{J_0-}^0, \mathbb{T}_{J_0I_0}^0, (\mathbb{T}_{I_0+}^0, \mathbb{T}_{I_0+}^1)) \\ &:= (\mathbb{T}_{J_0-}^0, \mathbb{T}_{1,J_0I_0+}^0) \text{ (using shorthand notation)}\end{aligned}$$

The simulator, next, proceeds with the remaining two rounds of simulations by forking at $\mathbb{Q}_{J_0}^0$ and $\mathbb{Q}_{I_2}^2$ (in that order). Let $\mathbb{T}_{1,J_0I_2+}^2 := (\mathbb{T}_{J_0I_2}^2, (\mathbb{T}_{I_2+}^2, \mathbb{T}_{I_2+}^3)) \in \mathbb{T}_{(1,J_0I_2+)}$ denote the tapes involved in the aforementioned simulation. Thus, we end up with the tapes given in **Figure 4.4**. The nested replay attack, all four rounds of it, is successful if the tape

$$\mathbb{T}_1^2 := (\mathbb{T}_{J_0-}^0, (\mathbb{T}_{J_0I_2}^2, (\mathbb{T}_{I_2+}^2, \mathbb{T}_{I_2+}^3)))$$

turns out to be $\mathbf{good}_{(1,J_0,I_2)}$ (note that I_2 need not be the same as I_0). Our objective is to find the probability of this particular event using the Splitting Lemma. This requires $\mathbb{T}_{(1)}$ to be split into $\mathbb{T}_{(J_0-)}$ and $\cup_{i=J_0}^q (\mathbb{T}_{(J_0i)} \times \mathbb{T}_{(i+)})$. This, in turn, can be achieved by splitting the individual partitions of $\mathbb{T}_{(1)}$. But note that *not all* of the partitions in $\mathbb{T}_{(1)}$ can be split in such a way. In fact, by definition, it is viable only to the partitions $(\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)})$ for $i > J_0$, *i.e.* the set $\mathbb{T}_{(1,J_0>)}$. At this point **Claim 3.2** comes into play. It allows us to apply the Splitting Lemma to the “reduced” universe of $\mathbb{T}_{(1,J_0>)}$ (by ensuring that at least δ_{J_0} fraction of the tapes in $\mathbb{T}_{(1,J_0>)}$ too are $\mathbf{good}_{(1,J_0,*)}$). We are now ready to apply the Splitting Lemma.

Let’s assume that an arbitrary λ_j fraction of the tapes in $\mathbb{T}_{(1)}$ (and hence $\mathbb{T}_{(1,J_0>)}$) are $\mathbf{better}_{(1,j,*)}$. We denote the (sub)set of $\mathbf{good}_{(1,J_0,*)}$ elements and the $\mathbf{better}_{(1,J_0,*)}$ elements in $\mathbb{T}_{(1,J_0>)}$ by \mathbb{U}_{J_0} and $\mathbb{U}_{J_0}^*$ respectively. For ease of analysis, we use the sufficient condition that the tape \mathbb{T}_1^0 be $\mathbf{good}_{(1,J_0,*)}$ and the \mathbb{T}_1^2 be $\mathbf{better}_{(1,J_0,*)}$. Let’s denote the probability of this event (with the tapes sampled as described above) by δ'_{J_0} , *i.e.*,

$$\begin{aligned}\delta'_{J_0} &= \mathbb{P}((\mathbb{T}_{J_0-}^0, \mathbb{T}_{1,J_0I_0+}^0) \in \mathbb{U}_{J_0}^* \wedge (\mathbb{T}_{J_0-}^0, \mathbb{T}_{1,J_0I_2+}^2) \in \mathbb{U}_{J_0}) \\ &= \mathbb{P}((\mathbb{T}_{J_0-}^0, \mathbb{T}_{1,J_0I_2+}^2) \in \mathbb{U}_{J_0} \mid (\mathbb{T}_{J_0-}^0, \mathbb{T}_{1,J_0I_0+}^0) \in \mathbb{U}_{J_0}^*) \cdot \mathbb{P}((\mathbb{T}_{J_0-}^0, \mathbb{T}_{1,J_0I_0+}^0) \in \mathbb{U}_{J_0}^*) \\ &= (\delta_{J_0} - \lambda_{J_0})\lambda_{J_0} \text{ (using propositions 1 and 2 of **Lemma 1**)}\end{aligned}\tag{4.8}$$

The above expression attains a maxima of $\delta_{J_0}^2/4$ at the point $\lambda_{J_0} = \delta_{J_0}/2$. Now, the probability

that the nested oracle replay attack is successful for any J_0 is given by

$$\begin{aligned} \delta' &= \sum_{J_0=1}^{q-1} \delta'_{J_0} \geq \sum_{J_0=1}^{q-1} \frac{\delta_{J_0}^2}{4} \\ &\geq \frac{\epsilon^4}{64q^3} \text{ (using Hölder's inequality and (4.7))} \end{aligned} \tag{4.9}$$

□

Significance of (in)dependency. Dependency plays a crucial role in the first two claims. It is the first criterion of **Definition 14** which allows us to apply the Splitting Lemma to a reduced universe of tapes in **Claim 3.2**. The second criterion of dependency comes into play in **Claim 3.1**. It ensures: if, indeed, the two tapes $(\mathbb{T}_{I-}, \mathbb{T}_{I+})$ and $(\mathbb{T}_{I-}, \mathbb{T}'_{I+})$ are both $\text{good}_{(1,*,I)}$, then they have the same target H-index as well. In fact, this is the *sole* reason why the analysis proceeds as in **§2.3.2.3**. These two handles, together, reduce the degradation by $O(q^2)$. Independency, on the other hand, is used in **Claim 3.3**. Recall that a $\text{good}_{(1,*,i)}$ can only belong to the partition $(\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)})$. Without independency, the set of candidate tapes (for replay of the H-oracle at a particular index j) is restricted to $(\mathbb{T}_{(i-)} \times \mathbb{T}_{(i+)})$. But, relaxing this condition (*i.e.*, incorporating independency) “expands” the set of candidate tapes to $\mathbb{T}_{(1,j>)}$. On a high level, this reduces the degradation by $O(q)$. Apply the two observations together, the degradation is reduced by $O(q^3)$.

The characteristic expression. An interesting observation is that the second iteration of the Splitting Lemma, on a high level, seems to be similar to the first one. This is also captured by the fact that the analysis is abstracted out by the characteristic expressions⁸

$$\delta' \geq \sum_{j=1}^{q-1} \delta_j^2 \quad \text{and} \quad \epsilon' \geq \sum_{i=2}^q \epsilon_i^2$$

under the set of constraints

$$\sum_{j=1}^{q-1} \delta_j = \epsilon', \quad \sum_{i=2}^q \epsilon_i = \epsilon \quad \text{and} \quad (0 \leq \epsilon_i, \delta_j \leq 1)_{i,j \in \{1, \dots, q\}}.$$

Clearly, there is a degree of *self-similarity* in the two expressions and, in addition, the expressions themselves are quite disparate from that of the nested replay attack (and the MF algorithm for

⁸In fact, it was the characteristic expression that guided us to the analysis (and not the other way around).

$n = 3$) in (2.21).

4.3.3 Taking Stock

It is evident that the modified GG-IBS allows for tighter reductions. The effective degradation is reduced from $O((q_H + q_G)^6)$ to $O((q_H + q_G)^3)$ (see **Table 4.1**). The “gain”, which is substantial, is by the virtue of the multiple forkings being replaced, on a high level, by (two) general forkings. This, in turn, is endowed by our observations in §4.2.

Scheme	Security Argument			
GG-IBS (Chapter 3)	Reduction	R_1	R_2	R_3
	Forking	F_Y	M_{W1}	M_{W3}
	Degradation	$O(q_G q_\varepsilon)$	$O((q_H + q_G)^2)$	$O((q_H + q_G)^6)$
Modified GG-IBS (Figure 4.3)	Reduction	R'_1	R'_3	
	Forking	F_Y	Nested rewinding	
	Degradation	$O(q_G q_\varepsilon)$	$O((q_H + q_G)^3)$	

Table 4.1: Degradation for the original and modified GG-IBS. q_G (resp q_H) denotes the upper bound on the G-oracle (resp H-oracle) queries, whereas, q_ε denotes upper bound on the extract queries.

Chapter 5

From sID IBS to ID IBS without Random Oracles

5.1 Introduction

The selective-identity (sID) model for identity-based cryptographic schemes was introduced in [CHK03]. The distinguishing feature of this model is that the adversary has to commit, beforehand, to a “target” identity—*i.e.*, the identity which it eventually forges on (see **Definition 10**). Since its induction, the relationship of the sID notion with various other notions of security has been extensively studied [CS06, CFH⁺09, Gal06, GH05]. One of the interesting results is the separation between the sID models and ID models for IBE in the standard model [Gal06]. Therefore, it is a general consensus that the sID model is much weaker than the full-identity (ID) model. However, it is *easier* to design efficient schemes, based on weaker assumptions, that are secure in the sID model compared to the ID model. This is, in particular, highlighted by the disparity in the construction of IBE schemes given in [BB04a] and [Wat05]. The former is simple and efficient, whereas the latter, involved. Therefore, a generic transformation from an sID scheme to ID scheme would be a problem worth pursuing. We could design efficient sID-secure schemes and then just bootstrap it to ID-security using the transformation¹. In fact, this is a long-standing open problem.

The primary focus of this chapter is on the question of constructing an ID-secure IBS, given an sID-secure IBS, in the standard model, and with reasonable security degradation (say polynomial).

¹A well-known example of such an approach is the design of CCA-secure encryption schemes from CPA-secure ones using the Fujisaki-Okamoto transformation [FO99], as demonstrated in [BF01].

We accomplish this through a generic transformation which uses a CHF and a *weakly*-secure PKS as black-box [CK13b]. We go one step further by applying the same construction technique to a *relaxed* notion of IBS security which we call the *weak* selective-identity (wID) model. The distinguishing feature of the wID model is that the adversary, apart from committing to the target identity, has to commit to a set of “query” identities—the set of identities which it wishes to query the signature and extract oracle with (see §5.4 for the definition of the security model). Thus, we reduce the problem of constructing an ID-secure IBS to that of constructing wID-secure IBS, a EU-GCMA-secure PKS and a CHF. Our approach can be considered to be an alternative paradigm to the aforementioned folklore construction of IBS.

The security argument constitutes the main hurdle—the construction itself is quite straightforward. The line of argument, roughly, is: given an adversary that breaks the ID-IBS, we construct algorithms to break either the sID/wID-IBS, the PKS or the CHF. It leads to a tightness gap of $O(q_s)$, where q_s is the upper bound on the number of signature queries that the adversary is allowed to make.

5.2 Chameleon Hash Function

A chameleon hash function (CHF) is a randomised trapdoor hash function. Apart from the *collision resistance* property, it has an additional “chameleon” property which enables anyone with the trapdoor information to efficiently generate collisions.

Definition 15 (Chameleon Hash Function [BCC88, KR00, Moh11]). A family of CHF \mathcal{H} consists of three PPT algorithms $\{G, h, h^{-1}\}$ described below.

Key Generation, $G(1^n)$: It takes as input the security parameter n (in unary). It outputs the evaluation key \mathbf{ek} and the trapdoor key \mathbf{td} .

Hash Evaluation, $h(\mathbf{ek}, m, r)$: It takes as input the evaluation key \mathbf{ek} , a message m from the message-space \mathbb{M} and a randomiser r from the domain \mathbb{R} . It outputs the hash value y from the range \mathbb{Y} .

Collision Generation, $h^{-1}(\mathbf{td}, m, r, m')$: It takes as input the trapdoor key \mathbf{td} , two messages $m, m' \in \mathbb{M}$ and $r \in \mathbb{R}$. It outputs $r' \in \mathbb{R}$ such that $h(\mathbf{ek}, m, r) = h(\mathbf{ek}, m', r')$; in other words, (m, r) and (m', r') is a collision.

Any CHF should satisfy the following two properties.

- (i) *Uniformity.* The distribution induced by $h(\mathbf{ek}, m, r)$ for all messages m and a randomly chosen r should be the same. In other words, the distributions $(\mathbf{ek}, h(\mathbf{ek}, m, r))$ and (\mathbf{ek}, y) should be computationally indistinguishable, where $(\mathbf{ek}, \mathbf{td}) \xleftarrow{\$} \mathcal{G}(1^n), r \xleftarrow{\mathcal{U}} \mathbb{R}$ and $y \xleftarrow{\mathcal{U}} \mathbb{Y}$.
- (ii) *Collision Resistance.* Given the evaluation key \mathbf{ek} , it should be hard to compute a pair $(m, r) \neq (m', r')$ such that $h(\mathbf{ek}, m, r) = h(\mathbf{ek}, m', r')$, *i.e.* the probability given below should be negligible for all PPT adversaries \mathbf{A} .

$$\mathbb{P}(h(\mathbf{ek}, m, r) = h(\mathbf{ek}, m', r') \wedge (m, r) \neq (m', r') : (\mathbf{ek}, \mathbf{td}) \xleftarrow{\$} \mathcal{G}(1^n); (m, r, m', r') \xleftarrow{\$} \mathbf{A}(\mathbf{ek}))$$

5.3 The Generic Transformation

The transformation takes as input: *i*) an sID-secure IBS $\mathcal{I}_s := \{\mathcal{G}_s, \mathcal{E}_s, \mathcal{S}_s, \mathcal{V}_s\}$; *ii*) an EU-GCMA-secure PKS $\mathfrak{P} := \{\mathcal{K}, \mathcal{S}_p, \mathcal{V}_p\}$; and *iii*) a CHF $\mathfrak{H} := \{\mathcal{G}_h, h, h^{-1}\}$, to output an ID-secure IBS $\mathcal{I} := \{\mathcal{G}, \mathcal{E}, \mathcal{S}, \mathcal{V}\}$. The basic idea is to *map* an identity \mathbf{id} in \mathcal{I} to an identity \mathbf{id}_s in \mathcal{I}_s using the CHF. These two identities are then *bound* by using the PKS. A formal description follows.

Assumptions. We denote the identity-space of \mathcal{I}_s (and that of resulting \mathcal{I}) by \mathbb{I} and its message-space by \mathbb{M} . For simplicity, we assume that *i*) the message-space of \mathfrak{P} *ii*) the message-space of \mathfrak{H} (denoted by \mathbb{M}_h) and *iii*) the range of \mathfrak{H} (denoted by \mathbb{Y}) *are all* the same set \mathbb{I} , *i.e.*, $\mathbb{M}_h = \mathbb{Y} = \mathbb{I}$.² In addition, the randomness space of \mathfrak{H} is denoted by \mathbb{R} . Therefore, for a particular evaluation key \mathbf{ek} , the hash evaluation algorithm can be considered as a function $h : \mathbb{I} \times \mathbb{R} \mapsto \mathbb{I}$. The description of the transformation is given in **Figure 5.1** [CK13b]. It is followed by the argument that the resultant IBS \mathcal{I} is secure in the ID model.

$$\mathcal{I} \leftarrow \mathfrak{T}(\mathcal{I}_s, \mathfrak{P}, \mathfrak{H})$$

Set-up, $\mathcal{G}(1^n)$: Invoke the algorithms \mathcal{G}_s , \mathcal{K} and \mathcal{G}_h (all) on 1^n to obtain $(\mathbf{msk}_s, \mathbf{mpk}_s)$, $(\mathbf{sk}, \mathbf{pk})$ and $(\mathbf{ek}, \mathbf{td})$ respectively. Return $\mathbf{msk} := (\mathbf{msk}_s, \mathbf{sk})$ as the master secret key and $\mathbf{mpk} := (\mathbf{mpk}_s, \mathbf{pk}, \mathbf{ek})$ as the master public key.

²This assumption can be relaxed—to accommodate a CHF with $\mathbb{M}_h \neq \mathbb{Y} \neq \mathbb{I}$ —using two collision resistant hash functions H and G defined as follows:

$$H : \mathbb{I} \mapsto \mathbb{M}_h \quad \text{and} \quad G : \mathbb{Y} \mapsto \mathbb{I}$$

These hash functions can be used in the protocol, and also in the security argument, to couple the CHF with the IBS.

Key Extraction, $E(\text{id}, \text{msk})$: Select $r \xleftarrow{\text{U}} \mathbb{R}$ and compute $\text{id}_s \leftarrow h(\text{ek}, \text{id}, r)$. Next, run $E_s(\text{id}_s, \text{msk}_s)$ and $S_p(\text{id}_s, \text{sk})$ to obtain usk_s and σ_p respectively. Finally, return $\text{usk} := (\text{usk}_s, r, \sigma_p)$ as the user secret key.

Signing, $S(\text{id}, m, \text{usk})$: Parse the user secret key usk as $(\text{usk}_s, r, \sigma_p)$ and compute $\text{id}_s \leftarrow h(\text{ek}, \text{id}, r)$. Next, run $S_s(\text{id}_s, m, \text{usk}_s)$ to obtain σ_s . Finally, return $\sigma := (\sigma_s, r, \sigma_p)$ as the signature.

Verification, $V(\sigma, \text{id}, m, \text{mpk})$: Parse σ as (σ_s, r, σ_p) and compute $\text{id}_s \leftarrow h(\text{ek}, \text{id}, r)$. Return 1 only if σ_p is a valid signature on id_s and σ_s is a valid signature on (id_s, m) . In other words, if $\text{op}_p \leftarrow V_p(\sigma_p, \text{id}_s, \text{pk})$ and $\text{op}_s \leftarrow V_s(\sigma_s, \text{id}_s, m, \text{mpk})$, return $(\text{op}_p \wedge \text{op}_s)$.

Figure 5.1: Constructing ID-secure IBS from an sID-secure IBS.

break

The Hash Evaluation function h is used to maintain a map and identity id in \mathfrak{I} on to and identity id_s in \mathfrak{I}_s . The mapped identities are then bound using σ_p . This is reflected in the structure of the user secret key for id which is of the form $(\text{usk}_s, r, \sigma_p)$.

Remark 18. Note that we have omitted td from the master secret key. The trapdoor key td —hence the collision generation function h^{-1} —is not used *per se* in the transformation. However, it *does* play a crucial role in its security argument.

5.3.1 Security Argument

For simplicity, we consider the security of the specific case of EU-sID-CMA model; we argue that the resulting IBS is EU-ID-CMA-secure. The details of both the security models is given in §1.3.2. The line of argument can be easily extended to other models as well³.

Theorem 4. *Let A be an $(\epsilon, t, q_\epsilon, q_s)$ -adversary against the IBS \mathfrak{I} in the EU-ID-CMA model. We can construct either*

(i) *Algorithm B_s which $(\epsilon_s, t_s, q_\epsilon, q_s)$ -breaks \mathfrak{I}_s in the EU-sID-CMA model, where*

$$\epsilon_s \geq \frac{1}{3q_s} \epsilon \quad \text{and} \quad t_s \leq t + (q_\epsilon + q_s)\tau_1, \quad \text{or}$$

³*e.g.*, consider the sM-sID-CMA model - the *selective-message, selective-identity chosen-message attack* model. It is similar to the EU-sID-CMA model, except that the adversary—in addition to committing to the target identity—has to commit to the target message too. If we start from sM-sID-CMA-secure IBS, we end up with an sM-ID-CMA-secure IBS.

(ii) Algorithm B_p which $(\epsilon_p, t_p, q_\epsilon + q_s)$ -breaks \mathfrak{P} in the EU-GCMA model, where

$$\epsilon_p = \frac{1}{3}\epsilon \quad \text{and} \quad t_p \leq t + (q_\epsilon\tau_2 + q_s\tau_3), \quad \text{or}$$

(iii) Algorithm B_h which (ϵ_h, t_h) -breaks \mathfrak{H} , where

$$\epsilon_h = \frac{1}{3}\epsilon \quad \text{and} \quad t_h \leq t + (q_\epsilon + q_s)\tau_1 + (q_\epsilon\tau_2 + q_s\tau_3).$$

Here, q_ϵ (resp. q_s) denotes the upper bound on the number of extract (resp. signature) queries that A can make. τ_1 is the time taken for generating a signature in \mathfrak{P} ; τ_2 (resp. τ_3) denotes the time taken to generate a user secret key (resp. signature) in \mathfrak{I}_s .

Proof. We classify the forgeries (mutually-exclusively) into three: **type 1**, **type 2** and **type 3**. A forgery qualifies as **type 1** if the adversary makes *at least* one signature query on the “target” identity—the identity which A eventually forges on—and produces a forgery with the binding (provided by the simulator) *intact*. In both **type 2** and **type 3** forgeries, the binding is violated by the adversary by some means. The strategy adopted in each of the three cases is different; we give a reduction B_s for **type 1**, B_p for **type 2** and B_h for **type 3** adversary. The details follow.

Classifying the forgery. Consider an adversary A in the EU-ID-CMA model. At the beginning of the security game, A is given the challenge master public key mpk by (its challenger) C . A produces a forgery after making a series of queries—extract and signature—adaptively with C . Let id_i denote the i^{th} extract query made by A , which is responded to with $\text{usk}_i = (\text{usk}_{s,i}, \hat{r}_i, \hat{\sigma}_{p,i})$ by C . Similarly, (id_i, m_i) denotes the i^{th} signature query by A , which is responded to with $\sigma_i = (\sigma_{s,i}, r_i, \sigma_{p,i})$ by C . Note that the number of extract (resp. signature) queries is bounded by q_ϵ (resp. q_s). Finally, let $\hat{\sigma} = (\hat{\sigma}_s, \hat{r}, \hat{\sigma}_p)$ be the forgery produced by A on $(\hat{\text{id}}, \hat{m})$. The identity $\hat{\text{id}}$ is the so-called target identity. The forgeries can be partitioned into three types, *viz.*:

- (i) **type 1** forgery. A produces the forgery with $(\hat{\text{id}}, \hat{r}) = (\text{id}_i, r_i)$ for some $i \in \{1, \dots, q_s\}$.
- (ii) **type 2** forgery. A produces the forgery with $(\hat{\text{id}}, \hat{r}) \neq (\text{id}_i, r_i)$ for all $i \in \{1, \dots, q_s\}$, and with
 - (a) $h(\text{ek}, \hat{\text{id}}, \hat{r}) \neq h(\text{ek}, \text{id}_i, \hat{r}_i)$ for all $i \in \{1, \dots, q_\epsilon\}$, and
 - (b) $h(\text{ek}, \hat{\text{id}}, \hat{r}) \neq h(\text{ek}, \text{id}_i, r_i)$ for all $i \in \{1, \dots, q_s\}$.

(iii) **type 3 forgery**. A produces the forgery with $(\hat{\text{id}}, \hat{r}) \neq (\text{id}_i, r_i)$ for all $i \in \{1, \dots, q_s\}$, but with

- (a) $h(\text{ek}, \hat{\text{id}}, \hat{r}) = h(\text{ek}, \text{id}_i, r_i)$ for some $i \in \{1, \dots, q_\varepsilon\}$, or
- (b) $h(\text{ek}, \hat{\text{id}}, \hat{r}) = h(\text{ek}, \text{id}_i, r_i)$ for some $i \in \{1, \dots, q_s\}$.

If A produces a forgery of **type 1**, we construct an algorithm B_s which breaks the IBS scheme \mathcal{I}_s ; whereas, in case of **type 2** forgery, we construct an algorithm B_p which breaks the PKS scheme \mathfrak{P} ; and finally, in case of **type 3** forgery, we construct an algorithm B_h that breaks collision resistance property of the CHF \mathfrak{H} . We describe these reductions in the subsequent sections. \square

5.3.1.1 Reduction B_s .

Recall that in **type 1** forgeries, A makes at least one signature query on the target identity $\hat{\text{id}}$. The strategy is to guess the index of this identity and map it to the identity that B_s *commits* to (initially) in the EU-sID-CMA game. This leads to a degradation of $O(q_s)$.

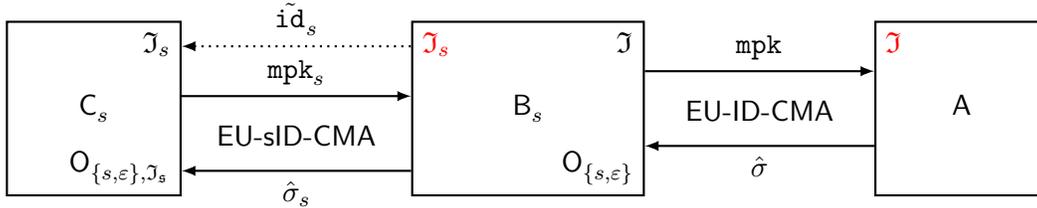


Figure 5.2: Reduction B_s

Let C_s be the challenger in the EU-sID-CMA game. B_s plays the role of the adversary in the EU-sID-CMA game and, at the same time, the role of the challenger to A in the EU-ID-CMA game. It starts by running the Key Generation algorithms K and G_h to obtain (pk, sk) and (ek, td) respectively. In order to initiate the EU-sID-CMA game, B_s has to commit to a target identity. It does so by selecting an identity $\hat{\text{id}} \xleftarrow{U} \mathbb{I}$ and a randomiser $\tilde{r} \xleftarrow{U} \mathbb{R}$, and committing $\hat{\text{id}}_s \leftarrow h(\text{ek}, \hat{\text{id}}, \tilde{r})$ to C_s . As a result, C_s releases the challenge master public key mpk_s to B_s . B_s is also allowed access to a signature oracle $O_{\mathcal{I}_s, \varepsilon}$. Now, B_s passes $\text{mpk} := (\text{mpk}_s, \text{pk}, \text{ek})$ as its own challenge master public key to A. Next, B_s guesses $1 \leq \tilde{\ell} \leq q_s$ as the index of the target identity.

Mapping the identities. In order to track the mapping between the identities in \mathcal{I} and \mathcal{I}_s , B_s maintains a table \mathcal{L} . It also maintains a counter ℓ (initially 1) to track the index of these identities.

\mathcal{L} contains tuples of the form $\langle \text{id}, \text{id}_s, \ell, \text{usk} \rangle$. Here, id and id_s are the related identities from \mathcal{I} and \mathcal{I}_s respectively; ℓ is the index of the identity id . The usk -field stores the user secret key for id and hence contains elements of the form $(\text{usk}_s, r, \sigma_p)$. If any component of the usk -field is yet to be generated, it is indicated by a ‘ \perp ’.

An identity id has already been mapped if there exists $\langle \text{id}_i, \text{id}_{s,i}, \ell_i, \text{usk}_i \rangle$ in \mathcal{L} such that $\text{id}_i = \text{id}$. For mapping a *fresh* identity id , B_s chooses $r \xleftarrow{\mathcal{U}} \mathbb{R}$ and sets $\text{id}_s \leftarrow h(\text{ek}, \text{id}, r)$.⁴ Finally, it adds $\langle \text{id}, \text{id}_s, \ell, (\perp, r, \perp) \rangle$ to \mathcal{L} and increments ℓ by one. A more formal description of the mapping function is given below.

```

Ms(id):
if  $\exists$  a tuple  $\langle \text{id}_i, \text{id}_{s,i}, \ell_i, \text{usk}_i \rangle \in \mathcal{L}$  such that  $(\text{id}_i = \text{id})$  then
    Set  $\tau := (\text{id}_{s,i}, \ell_i, \text{usk}_i)$ 
else
    if  $(\ell = \tilde{\ell})$  then set  $r \leftarrow h^{-1}(\text{td}, \tilde{\text{id}}, \tilde{r}, \text{id})$ 
    else choose  $r \xleftarrow{\mathcal{U}} \mathbb{R}$ 
    Compute  $\text{id}_s \leftarrow h(\text{ek}, \text{id}, r)$  and set  $\tau := (\text{id}_s, \ell, (\perp, r, \perp))$ 
    Add  $\langle \text{id}, \text{id}_s, \ell, (\perp, r, \perp) \rangle$  to  $\mathcal{L}$  and increment  $\ell$  by one
end if
return  $\tau$ 

```

Queries. The extract and signature queries by A are answered as per the following specifications.

Extract query, $O_{\varepsilon, \mathcal{I}}(\text{id})$: Invoke the function $M_s(\text{id})$ to obtain $(\text{id}_s, \ell, (\text{usk}_s, r, \sigma_p))$.

- (i) If $(\ell = \tilde{\ell})$ then B_s *aborts* (abort_1).
- (ii) Otherwise, if $(\text{usk}_s \neq \perp)$ then return $\text{usk} := (\text{usk}_s, r, \sigma_p)$ as the user secret key.
- (iii) Otherwise, B_s makes an extract query with $O_{\varepsilon, \mathcal{I}_s}$ on id_s to obtain usk_s . Next, it uses the knowledge of sk to compute $\sigma_p := S_p(\text{id}_s, \text{sk})$. Finally, it returns $\text{usk} := (\text{usk}_s, r, \sigma_p)$ as the user secret key and updates the usk -field of the tuple corresponding to id in \mathcal{L} .

Signature query, $O_{S, \mathcal{I}}(\text{id}, m)$: Invoke the function $M_s(\text{id})$ to obtain $(\text{id}_s, \ell, (\text{usk}_s, r, \sigma_p))$.

⁴If there already exists a tuple $\langle \text{id}_i, r_i, \text{id}_{s,i}, \ell_i, \text{usk}_i \rangle$ such that $\text{id}_{s,i} = \text{id}_s$, to maintain injection in the mapping, B_s repeats the process with a fresh r .

- (i) If $((\ell = \tilde{\ell}) \vee (\mathbf{usk}_s = \perp))$ then \mathcal{B}_s then makes a signature query with $\mathcal{O}_{S, \mathcal{J}_s}$ on (\mathbf{id}_s, m) to obtain σ_s . It uses the knowledge of \mathbf{sk} to compute $\sigma_p := S_p(\mathbf{id}_s, \mathbf{sk})$. Finally, it returns $\sigma := (\sigma_s, r, \sigma_p)$ as the signature.
- (ii) Otherwise, \mathcal{B}_s uses the knowledge of the user secret key \mathbf{usk} to generate the signature, *i.e.* it returns $\sigma := S(\mathbf{id}, m, \mathbf{usk})$.

Forgery. At the end of the simulation, \mathcal{A} produces a **type 1** forgery $\hat{\sigma} = (\hat{\sigma}_s, \hat{r}, \hat{\sigma}_p)$ on $(\hat{\mathbf{id}}, \hat{m})$. Let $\langle \mathbf{id}_{\hat{i}}, \mathbf{id}_{s, \hat{i}}, \ell_{\hat{i}}, \mathbf{usk}_{\hat{i}} \rangle$ be the tuple in \mathcal{L} such that $\mathbf{id}_{\hat{i}} = \hat{\mathbf{id}}$. If $\ell_{\hat{i}}$ matches \mathcal{B}_s 's initial guess for the target index (*i.e.* $\ell_{\hat{i}} = \tilde{\ell}$), it wins the EU-sID-CMA game with \mathcal{C}_s by passing $\hat{\sigma}_s$ as a forgery on $(\tilde{\mathbf{id}}_s, \hat{m})$ to \mathcal{C}_s ; otherwise it *aborts* (abort_2).

Analysis. The probability of success of the reduction \mathcal{B}_s is governed by the two events abort_1 and abort_2 . To be precise,

$$\begin{aligned} \epsilon_s &= \mathbb{P}(\neg \text{abort}_1 \wedge \neg \text{abort}_2) \epsilon \\ &= \mathbb{P}(\neg \text{abort}_1 \mid \neg \text{abort}_2) \mathbb{P}(\neg \text{abort}_2) \epsilon. \end{aligned}$$

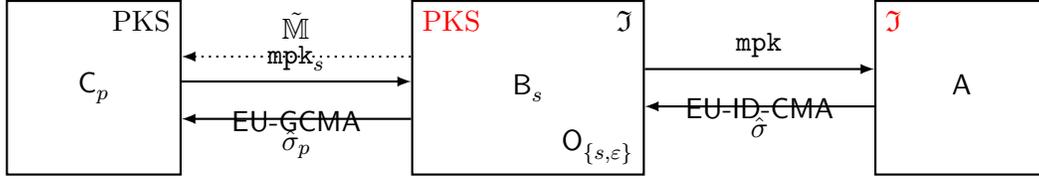
Since $\tilde{\ell}$ is hidden from the adversary, it is easy to see that

$$\mathbb{P}(\neg \text{abort}_2) = \mathbb{P}(\ell_{\hat{i}} = \tilde{\ell}) = 1/q_s.$$

On the other hand, $\mathbb{P}(\neg \text{abort}_1 \mid \neg \text{abort}_2) = 1$. This follows from the fact that if the simulator's guess of the target index was indeed correct ($\neg \text{abort}_2$), then the adversary *would not* have made an extract query on that identity (which causes abort_1). Thus, $\epsilon_s = \epsilon/q_s$. As for the time complexity, if τ_1 is the time taken for generating a signature in \mathfrak{P} , then the time taken by \mathcal{B}_s can be easily seen as $t_s \leq t + (q_\epsilon + q_s)\tau_1$.

5.3.1.2 Reduction \mathcal{B}_p .

The strategy adopted in \mathcal{B}_p is similar to that in security arguments of [HW09, ST01]. It is also, on a high level, related to the technique used in [BB04b] for proving the security of the EU-CMA-secure PKS scheme constructed from EU-GCMA-secure PKS scheme (using CHF implicitly). The details follow.

Figure 5.3: Reduction B_p

Let C_p be the challenger in the EU-GCMA game. B_p plays the role of the adversary in the EU-GCMA game and, at the same time, the role of the challenger to A in the EU-ID-CMA game. It starts by running the Key Generation algorithms G_s and G_h to obtain $(\text{mpk}_s, \text{msk}_s)$ and (ek, td) respectively. In order to initiate the EU-GCMA game, B_p has to commit to a set of q_s messages to C_p . On the other hand, it also has to answer the adaptive queries by A . Let's see how this is accomplished using the CHF. B_p first selects pairs $(\tilde{\text{id}}_1, \tilde{r}_1), \dots, (\tilde{\text{id}}_{q_s}, \tilde{r}_{q_s})$ independently and uniformly at random from $\mathbb{I} \times \mathbb{R}$. Next, it commits $\tilde{M} := \{\tilde{\text{id}}_{s,1}, \dots, \tilde{\text{id}}_{s,q_s}\}$ to C_s , where $\tilde{\text{id}}_{s,i} \leftarrow h(\text{ek}, \tilde{\text{id}}_i, \tilde{r}_i)$. As a result, C_p releases the challenge public key pk to B_p along with the set of signatures $\{\sigma_{p,1}, \dots, \sigma_{p,q_s}\}$ on the (respective) committed messages. All this information is stored in a table \mathfrak{C} as tuples $\langle \tilde{\text{id}}_i, \tilde{r}_i, \tilde{\text{id}}_{s,i}, \sigma_{p,i} \rangle$. Now, B_p initiates the EU-ID-CMA game by passing $\text{mpk} := (\text{mpk}_s, \text{pk}, \text{ek})$ as the challenge master public key to A .

Mapping the identities. B_p too maintains the table \mathfrak{L} ; but, its structure is slightly different from that in B_s . \mathfrak{L} contains tuples of the form $\langle \text{id}, \tilde{\text{id}}_s, \text{usk} \rangle$. Here, id and id_s are the related identities from \mathfrak{I} and \mathfrak{I}_s respectively. The usk -field stores the user secret key for id and hence contains elements of the form $(\text{usk}_s, r, \sigma_p)$. If any component of the usk -field is yet to be generated, it is indicated by a ' \perp '.

The way in which the mapping is maintained between the identities is somewhat different from that in B_s . For mapping a *fresh* identity id , B_p first picks a tuple $\mathfrak{t} = \langle \tilde{\text{id}}_s, \tilde{\text{id}}, \tilde{r}, \sigma_p \rangle$ randomly from \mathfrak{C} . It then computes $r \leftarrow h^{-1}(\text{td}, \tilde{\text{id}}, \tilde{r}, \text{id})$, and adds the tuple $\langle \text{id}, \tilde{\text{id}}_s, (\perp, r, \sigma_p) \rangle$ to \mathfrak{L} . Finally it removes the tuple \mathfrak{t} from \mathfrak{C} . As a result of these actions, id is effectively mapped to $\tilde{\text{id}}_s$ since $h(\text{ek}, \text{id}, r) = h(\text{ek}, \tilde{\text{id}}, \tilde{r}) = \tilde{\text{id}}_s$. A more formal description follows.

```

Mp(id):
if ∃ a tuple ⟨idi, ids,i, uski⟩ ∈ ℒ such that (idi = id) then
    Set τ := (ids,i, uski)
else

```

<p style="margin: 0;">Pick $\mathbf{t} \xleftarrow{\\$} \mathfrak{C}$ and parse it as $\langle \tilde{\mathbf{id}}, \tilde{r}, \tilde{\mathbf{id}}_s, \sigma_p \rangle$</p> <p style="margin: 0;">Compute $r \leftarrow h^{-1}(\mathbf{td}, \tilde{\mathbf{id}}, \tilde{r}, \mathbf{id})$ and set $\tau := (\mathbf{id}_s, (\perp, r, \sigma_p))$</p> <p style="margin: 0;">Add $\langle \mathbf{id}, \mathbf{id}_s, (\perp, r, \sigma_p) \rangle$ to \mathfrak{L} and remove \mathbf{t} from \mathfrak{C}</p> <p style="margin: 0;">end if</p> <p style="margin: 0;">break return τ</p>
--

Queries: The extract and signature queries by \mathbf{A} are answered as follows.

Extract query, $\mathcal{O}_{\varepsilon, \mathcal{J}}(\mathbf{id})$: Invoke the function $M_p(\mathbf{id})$ to obtain $(\tilde{\mathbf{id}}_s, \ell, (\mathbf{usk}_s, r, \sigma_p))$.

- (i) If $(\mathbf{usk}_s \neq \perp)$ then return $\mathbf{usk} := (\mathbf{usk}_s, r, \sigma_p)$ as the user secret key.
- (ii) Otherwise, \mathbf{B}_s uses the knowledge of the master secret key \mathbf{msk}_s to generate the user secret key $\mathbf{usk}_s := \mathbf{E}_s(\tilde{\mathbf{id}}_s, \mathbf{msk}_s)$ for $\tilde{\mathbf{id}}_s$. It returns $\mathbf{usk} := (\mathbf{usk}_s, r, \sigma_p)$ as the user secret key for \mathbf{id} and updates the \mathbf{usk}_s -field of the tuple corresponding to \mathbf{id} in \mathfrak{L} .

Signature query, $\mathcal{O}_{\mathcal{S}, \mathcal{J}}(\mathbf{id}, m)$: Invoke the function $M_p(\mathbf{id})$ to obtain $(\tilde{\mathbf{id}}_s, \ell, (\mathbf{usk}_s, r, \sigma_p))$.

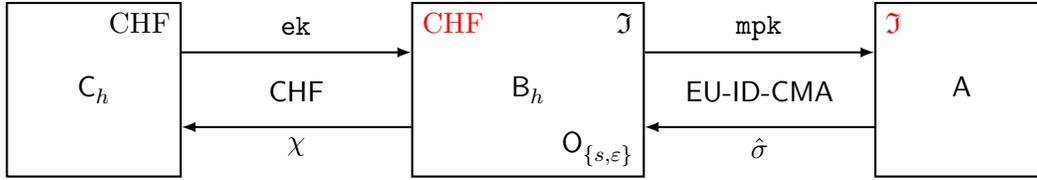
- (i) If $(\mathbf{usk}_s \neq \perp)$ then \mathbf{B}_p uses the knowledge of \mathbf{usk} to return the signature $\sigma := \mathbf{S}(\mathbf{id}, m, \mathbf{usk})$
- (ii) Otherwise, \mathbf{B}_s uses step (ii) of **Extract query** to generate a user secret key \mathbf{usk} for \mathbf{id} and then use this \mathbf{usk} to return a signature $\sigma := \mathbf{S}(\mathbf{id}, m, \mathbf{usk})$.

Forgery. Finally, \mathbf{A} produces a forgery $\hat{\sigma} = (\hat{\sigma}_s, \hat{r}, \hat{\sigma}_p)$ on $(\hat{\mathbf{id}}, \hat{m})$. As the forgery is of **type 2**, it implies $\hat{\mathbf{id}}_s := h(\mathbf{ek}, \hat{\mathbf{id}}, \hat{r}) \notin \tilde{\mathcal{M}}$. Therefore $\hat{\sigma}_p$ is a valid forgery in the EU-GCMA game and \mathbf{B}_p passes it to \mathbf{C}_s to win the game.

Analysis. Since no abort is involved in \mathcal{S}_p , there is no degradation involved either. Thus, its advantage in attacking \mathfrak{P} is $\epsilon_p = \epsilon$. If τ_2 and τ_3 denote the time taken for generating a secret key and a signature respectively in \mathcal{J}_s , then the time taken by \mathbf{B}_p is $t_p \leq t + (q_\varepsilon \tau_2 + q_s \tau_3)$.

5.3.1.3 Reduction \mathbf{B}_h .

\mathbf{B}_h first obtains the challenge evaluation key \mathbf{ek} for \mathfrak{F} from its challenger \mathbf{C}_h . Then it invokes the algorithms \mathbf{G}_s and \mathbf{K} to generate $(\mathbf{msk}_s, \mathbf{mpk}_s)$ and $(\mathbf{sk}, \mathbf{pk})$ respectively. Finally, it passes $(\mathbf{mpk}_s, \mathbf{pk}, \mathbf{ek})$ as the challenge master public key to \mathbf{A} .

Figure 5.4: Reduction B_h

Mapping the identities. The table used for maintaining mapping has the same structure as in B_p . However, the actual method used for mapping identities is far simpler than in B_p as shown below.

```

hM(id):
if  $\exists$  a tuple  $\langle \text{id}_i, \text{id}_{s,i}, \text{usk}_i \rangle \in \mathcal{L}$  such that  $(\text{id}_i = \text{id})$  then
  Set  $\tau := (\text{id}_{s,i}, \text{usk}_i)$ 
else
  Pick  $r \xleftarrow{\mathbb{U}} \mathbb{R}$  and compute  $\text{id}_s := h(\text{ek}, \text{id}, r)$ 
  Set  $\tau := (\text{id}_s, (\perp, r, \sigma_p))$  add  $\langle \text{id}, \text{id}_s, (\perp, r, \perp) \rangle$  to  $\mathcal{L}$ 
end if
return  $\tau$ 
  
```

Queries: The extract and signature queries by A are answered as follows.

Extract query, $O_{\varepsilon, \mathcal{J}}(\text{id})$: Invoke the function $hM(\text{id})$ to obtain $(\text{id}_s, (\text{usk}_s, r, \sigma_p))$.

- (i) If $(\text{usk}_s \neq \perp)$ then return $\text{usk} := (\text{usk}_s, r, \sigma_p)$ as the user secret key.
- (ii) Otherwise, B_s uses the knowledge of the master secret key msk_s to generate the user secret key $\text{usk}_s := E_s(\text{id}_s, \text{msk}_s)$ for id_s . It also uses sk to generate $\sigma_p := S_p(\text{id}_s)$. Finally, B_h returns $\text{usk} := (\text{usk}_s, r, \sigma_p)$ as the user secret key for id and updates the usk_s -field and σ_p -field of the tuple corresponding to id in \mathcal{L} .

Signature query, $O_{S, \mathcal{J}}(\text{id}, m)$: Invoke the function $hM(\text{id})$ to obtain $(\text{id}_s, (\text{usk}_s, r, \sigma_p))$.

- (i) If $(\text{usk}_s \neq \perp)$ then B_p uses the knowledge of usk to return the signature $\sigma := S(\text{id}, m, \text{usk})$
- (ii) Otherwise, B_s uses step (ii) of **Extract query** to generate a user secret key usk for id and then use this usk to return a signature $\sigma := S(\text{id}, m, \text{usk})$.

Queries. The extract and signature queries by A are answered as follows.

Extract query, $O_{\varepsilon, \mathcal{J}}(\text{id})$: Let $(\text{id}, \text{id}_s, r)$ be the mapping for id . B_s uses the knowledge of the master secret key msk_s to generate the user secret key for id_s in \mathcal{J}_s , *i.e.* $\text{usk}_s := E_s(\text{id}_s, \text{msk}_s, \text{mpk}_s)$. It then uses the knowledge of secret key sk to generate a signature σ_p on id_s . It returns the user secret key for id

$$\text{usk} := (\text{usk}_s, r, \sigma_p).$$

Signature query, $O_{S, \mathcal{J}}(\text{id}, m)$: First, B_p generates the user secret key for id by running E , *i.e.* $\text{usk} = E(\text{id}, \text{msk}, \text{mpk})$. Then, B_p uses usk to return the signature

$$\sigma := S(\text{id}, m, \text{usk}).$$

Forgery. Finally, A produces a **type 3** forgery $\hat{\sigma} = (\hat{\sigma}_s, \hat{r}, \hat{\sigma}_p)$ on $(\hat{\text{id}}, \hat{m})$. Recall that this implies A produces the forgery with $(\hat{\text{id}}, \hat{r}) \neq (\text{id}_i, r_i)$ for all $i \in \{1, \dots, q_s\}$, but with

- (a) $h(\text{ek}, \hat{\text{id}}, \hat{r}) = h(\text{ek}, \text{id}_i, r_i)$ for some $i \in \{1, \dots, q_s\}$, or
- (b) $h(\text{ek}, \hat{\text{id}}, \hat{r}) = h(\text{ek}, \text{id}_i, r_i)$ for some $i \in \{1, \dots, q_s\}$.

Both the cases tantamount to breaking the collision resistance property of \mathfrak{H} . In case (a) (resp. case (b)) B_h passes $\chi := ((\hat{\text{id}}, \hat{r}), (\text{id}_i, r_i))$ (resp. $\chi := ((\hat{\text{id}}, \hat{r}), (\text{id}_i, r_i))$) as a collision to the challenger C_h to win the game.

Analysis. As in B_p , there is no abort involved in B_h . Therefore, its advantage in attacking \mathfrak{H} is $\epsilon_p = \epsilon$. Again, if τ_2 and τ_3 denote the time taken for generating a secret key and a signature respectively in \mathcal{J}_s , then the time taken by B_h is $t_p \leq t + (q_\varepsilon \tau_2 + q_s \tau_3)$.

5.4 Transforming from the EU-wID-CMA model

The construction technique described in the previous section can as well be used with a *relaxed* version of the selective-identity model which we call the *weak selective-identity* (wID) model. In this model the adversary, apart from committing to the “target” identity $\tilde{\text{id}}$, has to commit a set of “query” identities $\tilde{\mathbb{I}}$. The adversary is allowed to query the extract oracle only on identities

belonging to $\tilde{\mathbb{I}}$; whereas, it is allowed to query the signature oracle with identities from $\tilde{\mathbb{I}}$ *as well as* the target identity. Finally, as in the sID model, the adversary has to produce a forgery on $\tilde{\text{id}}$. One may see the analogy between the EU-GCMA model for PKS and the wID model—both involve the adversary committing, beforehand, to the identities/messages that it wants to queries to. The only change involved is in the security argument—the way in which mapping is handled by the simulator. We elaborate on this later. But first, let’s formally define the EU-wID-CMA model for IBS.

Definition 16 (EU-wID-CMA Game). The security of an IBS scheme in the EU-wID-CMA model is argued in terms of the following game between a challenger C and an adversary A.

Commitment: A commits to a target identity $\tilde{\text{id}}$ and a set of query identities $\tilde{\mathbb{I}} := \{\tilde{\text{id}}_1, \dots, \tilde{\text{id}}_{\tilde{q}}\} \subset \mathbb{I} \setminus \{\tilde{\text{id}}\}$.

Set-up: C runs the set-up algorithm G to obtain the master keys (mpk, msk) . It passes mpk as the challenge master public key to A.

Queries: A can adaptively make extract queries on identities from $\tilde{\mathbb{I}}$ to an oracle \mathcal{O}_E and signature queries involving identities from $\tilde{\mathbb{I}} \cup \{\tilde{\text{id}}\}$ to an oracle \mathcal{O}_s . These queries are handled as follows.

Extract query, $\mathcal{O}_E(\text{id})$: A asks for the secret key of a user with identity $\text{id} \in \tilde{\mathbb{I}}$. C computes $\text{usk} := E(\text{id})$ and passes it to A.

Signature query, $\mathcal{O}_s(\text{id}, m)$: A asks for the signature of a user with identity $\text{id} \in \tilde{\mathbb{I}} \cup \{\tilde{\text{id}}\}$ on a message m . C first runs E on id to obtain the user secret key usk . Next, it computes $\sigma := S(\text{id}, m, \text{usk})$ and forwards it to A.

Forgery: A outputs a signature $\hat{\sigma}$ on a message \hat{m} and the target identity $\tilde{\text{id}}$ wins the game if:

1. $\hat{\sigma}$ is a valid signature on \hat{m} by $\tilde{\text{id}}$.
2. A has not made a signature query on $(\tilde{\text{id}}, \hat{m})$.

The advantage that A has in the above game, denoted by $\text{Adv}_A^{\text{EU-wID-CMA}}(n)$, is defined as the probability with which it wins the game, *i.e.*

$$\mathbb{P}(1 \leftarrow \mathcal{V}(\hat{\sigma}, \hat{\text{id}}, \hat{m}, \text{mpk}) : (\tilde{\text{id}}, \tilde{\mathbb{I}}) \xleftarrow{\$} A; (\text{msk}, \text{mpk}) \xleftarrow{\$} G(1^n); (\hat{\sigma}, \hat{\text{id}}, \hat{m}) \xleftarrow{\$} A^{\mathcal{O}_E, \mathcal{O}_s}(\text{mpk}))$$

where the oracles \mathcal{O}_E and \mathcal{O}_s are restricted to answering queries involving identities from $\tilde{\mathbb{I}}$ and $\tilde{\mathbb{I}} \cup \{\tilde{\text{id}}\}$ respectively. An adversary A is said to be an $(\epsilon, t, q_\epsilon, q_s, \tilde{q})$ -forger of an IBS scheme in the EU-wID-CMA model if it has advantage of at least ϵ in the above game, runs in time at most t and makes at most q_ϵ and q_s extract and signature queries respectively, provided the number of identities involved in the signature and extract queries, excluding the target identity, is at most \tilde{q} . It is easy to see that $\tilde{q} \leq q_\epsilon + q_s$. As we pointed out, the same transformation technique applies; the only change is in the security argument.

Theorem 5. *Let A be an $(\epsilon, t, q_\epsilon, q_s)$ -adversary against the IBS \mathfrak{I} in the EU-ID-CMA model. We can construct either*

(i) *Algorithm B_w which $(\epsilon_w, t_w, q_\epsilon, q_s, q_\epsilon + q_s)$ -breaks \mathfrak{I}_w in the EU-wID-CMA model, where*

$$\epsilon_w \geq \frac{1}{3q_s}\epsilon \quad \text{and} \quad t_w \leq t + (q_\epsilon + q_s)\tau_1, \quad \text{or}$$

(ii) *Algorithm B_p which $(\epsilon_p, t_p, q_\epsilon + q_s)$ -breaks \mathfrak{P} in the EU-GCMA model, where*

$$\epsilon_p = \frac{1}{3}\epsilon \quad \text{and} \quad t_p \leq t + (q_\epsilon\tau_2 + q_s\tau_3), \quad \text{or}$$

(iii) *Algorithm B_h which (ϵ_h, t_h) -breaks \mathfrak{H} , where*

$$\epsilon_h = \frac{1}{3}\epsilon \quad \text{and} \quad t_h \leq t + (q_\epsilon + q_s)\tau_1 + (q_\epsilon\tau_2 + q_s\tau_3).$$

Here, q_ϵ (resp. q_s) denotes the upper bound on the number of extract (resp. signature) queries that A can make. τ_1 is the time taken for generating a signature in \mathfrak{P} ; τ_2 (resp. τ_3) denotes the time taken to generate a user secret key (resp. signature) in \mathfrak{I}_s .

Proof. The security argument is similar to the one discussed in §5.3.1. The only difference lies in the way in which the mapping of identities is handled in B_w .

Let C_w be the challenger in the EU-wID-CMA game. B_w plays the role of the adversary in the EU-wID-CMA game and, at the same time, the role of the challenger to A in the EU-ID-CMA game. In order to initiate the EU-wID-CMA game, B_w has to commit to a target identity and a target set. It selects an identity $\tilde{\text{id}} \xleftarrow{\mathbb{U}} \mathbb{I}$ and a randomiser $\tilde{r} \xleftarrow{\mathbb{U}} \mathbb{R}$, and commits $\tilde{\text{id}}_w \leftarrow h(\text{ek}, \tilde{\text{id}}, \tilde{r})$ as the target identity to C_w . Similarly, it selects $\{\tilde{\text{id}}_i, \dots, \tilde{\text{id}}_{\tilde{q}}\} \xleftarrow{\mathbb{S}} \mathbb{I}$, $\{\tilde{r}_1, \dots, \tilde{r}_{\tilde{q}}\} \xleftarrow{\mathbb{S}} \mathbb{R}$ and commits

$\hat{\mathbb{I}} := \{\tilde{\text{id}}_{1,w}, \dots, \tilde{\text{id}}_{q,w}\}$, where $\tilde{\text{id}}_{i,w} \leftarrow h(\mathbf{ek}, \tilde{\text{id}}_i, \tilde{r}_i)$, as the target set to C_w . As a result, C_w releases the challenge master public key mpk_w to B_s . All this information is stored in a table, denoted by \mathfrak{D} , as tuples $\langle \tilde{\text{id}}_i, \tilde{r}_i, \tilde{\text{id}}_{w,i} \rangle$.

Mapping. B_w maintains a table \mathfrak{L} with structure the same as that in reduction B_p . For mapping a *fresh* identity id , B_s chooses a tuple $\mathfrak{t} = \langle \tilde{\text{id}}, \tilde{r}, \tilde{\text{id}}_w \rangle$ randomly from \mathfrak{D} . Next, it computes $r := h^{-1}(\text{td}, \tilde{\text{id}}, \tilde{r}, \text{id})$ and adds $\langle \text{id}, \tilde{\text{id}}_w, (\perp, r, \perp) \rangle$ to \mathfrak{L} . Finally, it removes the tuple \mathfrak{t} from \mathfrak{D} . As a result of these actions, id is effectively mapped to $\tilde{\text{id}}_w$ as $h(\mathbf{ek}, \text{id}, r) = h(\mathbf{ek}, \tilde{\text{id}}, \tilde{r}) = \tilde{\text{id}}_w$. A more formal description follows.

```

Mw(id):
  if ∃ a tuple ⟨idi, idw,i, uski⟩ ∈ ℒ such that (idi = id) then
    Set τ := (idw,i, uski)
  else
    Pick t  $\stackrel{\$}{\leftarrow}$  ℄ and parse it as ⟨id̃, r̃, id̃w⟩
    Compute r ← h-1(td, id̃, r̃, id) and set τ := (idw, (⊥, r, ⊥))
    Add ⟨id, idw, (⊥, r, ⊥)⟩ to ℒ and remove t from ℄
  end if
  return τ

```

□

Remark 19 (Comparison with the folklore paradigm). The (identities-based) signature of an IBS scheme constructed using the folklore technique consists of *two* (public-key) signatures and *one* public key of the underlying (fully-secure) PKS. In contrast, the signature of an IBS scheme using our approach consists of *one* signature each of the underlying (wID-secure) IBS and (weakly-secure) PKS and *one* randomiser from the CHF. The time taken for signing and verification is comparable, bar the time taken to compute the hash value.

Chapter 6

Conclusions

In this thesis we have identified certain shortcomings in the original security argument of GG-IBS in [GG09]. Based on our observations, we provide a new elaborate security argument for the same scheme. Although, the reductions are tighter than their counterparts in the original security argument, they are still quite loose due to the usage of the MF Algorithm. This motivated us to explore means to launch the nested replay attack in a more effective manner than in the MF Algorithm and our (re)search culminated with the notion of (in)dependency. The result was a cleaner, tighter security argument for GG-IBS with the effective degradation down from $O(q^6)$ to $O(q^3)$.

The bound on security reductions for the Schnorr signature has been well-studied [PV05, GBL08, Seu12]. In the same vein, giving a bound on the reductions for GG-IBS would truly complete the jigsaw (of GG-IBS). In addition, studying the effect of (in)dependency on the MF Algorithm should also be worthwhile [CK13a].

On a separate note, we described a generic transformation from sID/wID IBS to full-identity IBS using a chameleon hash function and a GCMA-secure PKS scheme. It was also supported by an argument, without using random oracles, that the resulting IBS is secure in the full-identity model with only linear degradation incurred. An interesting problem would be to replace the EU-GCMA PKS with a more primitive construct. Extending the transformation for Hierarchical IBS could be yet another challenging task.

Bibliography

- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009. (Cited on page [iii](#).)
- [ARP03] Sattam Al-Riyami and Kenneth Paterson. Certificateless public key cryptography. In Chi-Sung Lai, editor, *Advances in Cryptology - ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, pages 452–473. Springer Berlin / Heidelberg, 2003. (Cited on page [2](#).)
- [BB04a] Dan Boneh and Xavier Boyen. Efficient selective-id secure identity-based encryption without random oracles. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 223–238. Springer Berlin / Heidelberg, 2004. (Cited on pages [13](#), [45](#), [47](#), [59](#), [84](#) and [115](#).)
- [BB04b] Dan Boneh and Xavier Boyen. Short signatures without random oracles. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 56–73. Springer Berlin / Heidelberg, 2004. (Cited on page [91](#).)
- [BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2):156 – 189, 1988. (Cited on page [85](#).)
- [BF01] Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In Joe Kilian, editor, *Advances in Cryptology — CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 213–229. Springer Berlin / Heidelberg, 2001. (Cited on pages [2](#), [6](#), [8](#), [13](#) and [84](#).)

- [BN06] Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In *Proceedings of the 13th ACM conference on Computer and communications security*, CCS '06, pages 390–399, New York, NY, USA, 2006. ACM. (Cited on pages 14, 16, 32, 33 and 34.)
- [BNN04] Mihir Bellare, Chanathip Namprempre, and Gregory Neven. Security proofs for identity-based identification and signature schemes. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 268–286. Springer Berlin / Heidelberg, 2004. (Cited on pages 10, 12, 14 and 44.)
- [BPW12] Alexandra Boldyreva, Adriana Palacio, and Bogdan Warinschi. Secure proxy signature schemes for delegation of signing rights. *Journal of Cryptology*, 25:57–115, 2012. (Cited on pages 14, 16, 27, 33, 35 and 42.)
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *Proceedings of the 1st ACM conference on Computer and communications security*, CCS '93, pages 62–73, New York, NY, USA, 1993. ACM. (Cited on pages 6 and 13.)
- [CC02] Jae Choon and Jung-Hee Cheon. An identity-based signature from gap Diffie-Hellman groups. In Yvo Desmedt, editor, *Public Key Cryptography — PKC 2003*, volume 2567 of *Lecture Notes in Computer Science*, pages 18–30. Springer Berlin / Heidelberg, 2002. (Cited on page 1.)
- [CFH⁺09] Yang Cui, Eiichiro Fujisaki, Goichiro Hanaoka, Hideki Imai, and Rui Zhang. Formal security treatments for ibe-to-signature transformation: Relations among security notions. *IEICE Transactions*, 92-A(1):53–66, 2009. (Cited on page 84.)
- [CHK03] Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In Eli Biham, editor, *Advances in Cryptology — EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 646–646. Springer Berlin / Heidelberg, 2003. (Cited on pages 11, 15 and 84.)

- [CK13a] Sanjit Chatterjee and Chethan Kamath. A closer look at multiple-forking: Leveraging (in)dependence for a tighter bound. Cryptology ePrint Archive, Report 2013/651, 2013. <http://eprint.iacr.org/>. (Cited on page 99.)
- [CK13b] Sanjit Chatterjee and Chethan Kamath. From selective-id to full-id IBS without random oracles. In Benedikt Gierlichs, Sylvain Guilley, and Debdeep Mukhopadhyay, editors, *Security, Privacy, and Applied Cryptography Engineering*, volume 8204 of *Lecture Notes in Computer Science*, pages 172–190. Springer Berlin Heidelberg, 2013. (Cited on pages 15, 85 and 86.)
- [CKK13] Sanjit Chatterjee, Chethan Kamath, and Vikas Kumar. Galindo-Garcia identity-based signature revisited. In Taekyoung Kwon, Mun-Kyu Lee, and Daesung Kwon, editors, *Information Security and Cryptology - ICISC 2012*, volume 7839 of *Lecture Notes in Computer Science*, pages 456–471. Springer Berlin / Heidelberg, 2013. Full version available in Cryptology ePrint Archive, Report 2012/646, <http://eprint.iacr.org/2012/646>. (Cited on pages 14 and 35.)
- [CMW12] Sherman S. M. Chow, Changshe Ma, and Jian Weng. Zero-knowledge argument for simultaneous discrete logarithms. *Algorithmica*, 64(2):246–266, 2012. (Cited on page 35.)
- [Cor00] Jean-Sébastien Coron. On the exact security of full domain hash. In Mihir Bellare, editor, *Advances in Cryptology — CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 229–235. Springer Berlin / Heidelberg, 2000. (Cited on pages 44, 45, 57 and 112.)
- [CS06] Sanjit Chatterjee and Palash Sarkar. Generalization of the selective-id security model for HIBE protocols. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *Public Key Cryptography - PKC 2006*, volume 3958 of *Lecture Notes in Computer Science*, pages 241–256. Springer Berlin / Heidelberg, 2006. (Cited on page 84.)
- [ElG85] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In George Robert Blakley and David Chaum, editors, *Advances in Cryptology*, volume 196 of *Lecture Notes in Computer Science*, pages 10–18. Springer Berlin Heidelberg, 1985. (Cited on page 16.)

- [FLR⁺10] Marc Fischlin, Anja Lehmann, Thomas Ristenpart, Thomas Shrimpton, Martijn Stam, and Stefano Tessaro. Random oracles with(out) programmability. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 303–320. Springer Berlin / Heidelberg, 2010. (Cited on pages 6 and 7.)
- [FO99] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Michael Wiener, editor, *Advances in Cryptology — CRYPTO’ 99*, volume 1666 of *Lecture Notes in Computer Science*, pages 79–79. Springer Berlin / Heidelberg, 1999. (Cited on page 84.)
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew Odlyzko, editor, *Advances in Cryptology — CRYPTO’ 86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer Berlin / Heidelberg, 1987. (Cited on pages 1, 6 and 16.)
- [Gal05] David Galindo. Boneh-Franklin identity based encryption revisited. In Luís Caires, Giuseppe Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 102–102. Springer Berlin / Heidelberg, 2005. (Cited on page 44.)
- [Gal06] David Galindo. A separation between selective and full-identity security notions for identity-based encryption. In Marina Gavrilova, Osvaldo Gervasi, Vipin Kumar, C. Tan, David Taniar, Antonio Laganá, Youngsong Mun, and Hyunseung Choo, editors, *Computational Science and Its Applications - ICCSA 2006*, volume 3982 of *Lecture Notes in Computer Science*, pages 318–326. Springer Berlin / Heidelberg, 2006. (Cited on page 84.)
- [GBL08] Sanjam Garg, Raghav Bhaskar, and Satyanarayana V. Lokam. Improved bounds on security reductions for discrete log based signatures. In *Proceedings of the 28th Annual conference on Cryptology: Advances in Cryptology, CRYPTO 2008*, pages 93–107, Berlin, Heidelberg, 2008. Springer-Verlag. (Cited on page 99.)
- [GG09] David Galindo and Flavio Garcia. A Schnorr-like lightweight identity-based signature scheme. In Bart Preneel, editor, *Progress in Cryptology – AFRICACRYPT 2009*, volume

- 5580 of *Lecture Notes in Computer Science*, pages 135–148. Springer Berlin / Heidelberg, 2009. (Cited on pages [12](#), [14](#), [35](#), [44](#), [45](#), [48](#), [49](#), [55](#), [63](#), [66](#), [68](#), [99](#), [110](#), [120](#) and [124](#).)
- [GH05] David Galindo and Ichiro Hasuo. Security notions for identity based encryption. Cryptology ePrint Archive, Report 2005/253, 2005. (Cited on page [84](#).)
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ron Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, 1988. (Cited on pages [7](#) and [8](#).)
- [GQ90] Louis Guillou and Jean-Jacques Quisquater. A “paradoxical” identity-based signature scheme resulting from zero-knowledge. In Shafi Goldwasser, editor, *Advances in Cryptology — CRYPTO’ 88*, volume 403 of *Lecture Notes in Computer Science*, pages 216–231. Springer Berlin / Heidelberg, 1990. (Cited on page [1](#).)
- [Her05] Javier Herranz. Deterministic identity-based signatures for partial aggregation. *The Computer Journal*, 49(3):322–330, 2005. (Cited on page [1](#).)
- [Hes03] Florian Hess. Efficient identity based signature schemes based on pairings. In Kaisa Nyberg and Howard Heys, editors, *Selected Areas in Cryptography*, volume 2595 of *Lecture Notes in Computer Science*, pages 310–324. Springer Berlin / Heidelberg, 2003. (Cited on page [1](#).)
- [HW09] Susan Hohenberger and Brent Waters. Realizing hash-and-sign signatures under standard assumptions. In Antoine Joux, editor, *Advances in Cryptology - EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 333–350. Springer Berlin Heidelberg, 2009. (Cited on pages [12](#) and [91](#).)
- [HWS⁺06] Anne Marie Hegland, Eli Winjum, Pal Spilling, Chunming Rong, and Oivind Kure. Analysis of IBS for MANET security in emergency and rescue operations. In *Advanced Information Networking and Applications, 2006. AINA 2006. 20th International Conference on*, volume 2, pages 155–159, 2006. (Cited on page [1](#).)
- [Kia07] Aggelos Kiayias. *Cryptography: Primitives and protocols*, 2007. (Cited on page [22](#).)
- [KLS00] Stephen Kent, Charles Lynn, and Karen Seo. Secure border gateway protocol (S-BGP). *Selected Areas in Communications, IEEE Journal on*, 18(4):582–592, 2000. (Cited on page [1](#).)

- [KM07] Neal Koblitz and Alfred J. Menezes. Another look at "provable security". *Journal of Cryptology*, 20:3–37, 2007. (Cited on pages 3 and 23.)
- [KR00] Hugo Krawczyk and Tal Rabin. Chameleon signatures. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2000, San Diego, California, USA*. The Internet Society, 2000. (Cited on pages 12 and 85.)
- [LBZ⁺10] Joseph K. Liu, Joonsang Baek, Jianying Zhou, Yanjiang Yang, and JunWen Wong. Efficient online/offline identity-based signature for wireless sensor network. *International Journal of Information Security*, 9(4):287–296, 2010. (Cited on page 1.)
- [Moh11] Payman Mohassel. One-time signatures and chameleon hash functions. In Alex Biryukov, Guang Gong, and Douglas Stinson, editors, *Selected Areas in Cryptography*, volume 6544 of *Lecture Notes in Computer Science*, pages 302–319. Springer Berlin / Heidelberg, 2011. (Cited on page 85.)
- [Oka93] Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In ErnestF. Brickell, editor, *Advances in Cryptology — CRYPTO' 92*, volume 740 of *Lecture Notes in Computer Science*, pages 31–53. Springer Berlin Heidelberg, 1993. (Cited on page 16.)
- [PS00] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13:361–396, 2000. (Cited on pages 14, 16, 21 and 22.)
- [PV05] Pascal Paillier and Damien Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In Bimal Roy, editor, *Advances in Cryptology - ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, pages 1–20. Springer Berlin Heidelberg, 2005. (Cited on pages 4 and 99.)
- [RS11] V. Radhakishan and S. Selvakumar. Prevention of man-in-the-middle attacks using id-based signatures. In *Second International Conference on Networking and Distributed Computing - ICNDC, 2011*, pages 165–169, 2011. (Cited on page 45.)
- [Sch90] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, volume 435 of

- Lecture Notes in Computer Science*, pages 239–252. Springer Berlin / Heidelberg, 1990. (Cited on page 16.)
- [Sch91] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4:161–174, 1991. 10.1007/BF00196725. (Cited on pages 7, 12, 14 and 45.)
- [Seu12] Yannick Seurin. On the exact security of Schnorr-type signatures in the random oracle model. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 554–571. Springer Berlin / Heidelberg, 2012. (Cited on page 99.)
- [Sha85] Adi Shamir. Identity-based cryptosystems and signature schemes. In George Blakley and David Chaum, editors, *Advances in Cryptology*, volume 196 of *Lecture Notes in Computer Science*, pages 47–53. Springer Berlin / Heidelberg, 1985. (Cited on page 1.)
- [Sho01] Victor Shoup. OAEP reconsidered. In Joe Kilian, editor, *Advances in Cryptology — CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 239–259. Springer Berlin / Heidelberg, 2001. (Cited on page 44.)
- [Sho04] Victor Shoup. Sequences of games: a tool for taming complexity in security proofs, 2004. (Cited on page 3.)
- [ST01] Adi Shamir and Yael Tauman. Improved online/offline signature schemes. In Joe Kilian, editor, *Advances in Cryptology — CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 355–367. Springer Berlin / Heidelberg, 2001. (Cited on pages 12 and 91.)
- [Teg03] Max Tegmark. Parallel universes. In John Barrow, Paul Davies, and Charles Harper Jr., editors, *Science and Ultimate Reality: Quantum Theory, Cosmology, and Complexity*, pages 459–491. Cambridge University Press, 2003. (Cited on page 25.)
- [TWZL08] Chiu C. Tan, Haodong Wang, Sheng Zhong, and Qun Li. Body sensor network security: an identity-based cryptography approach. In *Proceedings of the first ACM conference on Wireless network security*, WiSec '08, pages 148–153, New York, NY, USA, 2008. ACM. (Cited on page 1.)
- [Wat05] Brent Waters. Efficient identity-based encryption without random oracles. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture*

Notes in Computer Science, pages 557–557. Springer Berlin / Heidelberg, 2005. (Cited on pages 6 and 84.)

- [XW12] Min Xie and Libin Wang. One-round identity-based key exchange with perfect forward security. *Information Processing Letters*, 112(14–15):587 – 591, 2012. (Cited on page 45.)

Appendices

Appendix A

Galindo-Garcia IBS

A.1 The Fixed Security Argument

Let A be an adversary against GG-IBS in EU-ID-CMA model. Eventually, A outputs an attempted forgery of the form $\sigma = (A, b, R)$. Let E be the event that σ is a valid signature and R was contained in an answer of the signature oracle O_s . Let NE be the event that σ is a valid signature and R was never part of an answer of O_s . Galindo and Garcia construct algorithms B_1 (resp. B_2) that break the DLP in case of event E (resp. NE). We describe the modified reductions below.

A.1.1 Reduction B_1

B_1 takes as argument the description of a group (\mathbb{G}, p, g) and a challenge g^α with $\alpha \xleftarrow{U} \mathbb{Z}_p$ and tries to extract the discrete logarithm α . The environment is simulated as shown below.

$B_{1.1}$ B_1 picks $\hat{i} \xleftarrow{U} \{1, \dots, q_G\}$,¹ where q_G is the maximum number of queries that the adversary A makes to the G-oracle. Let \hat{id} (the target identity) be the \hat{i} th distinct identity queried to the G-oracle. Next, B_1 chooses $z \xleftarrow{U} \mathbb{Z}_p$ and sets $(\mathbf{mpk}, \mathbf{msk}) := ((\mathbb{G}, g, p, G, H, g^z), z)$, where G, H are descriptions of hash functions modelled as random oracles. As usual, B_1 simulates these oracles with the help of two tables \mathcal{L}_G and \mathcal{L}_H containing the queried values along with the answers given to A .

$B_{1.2}$ Every time A queries the key extraction oracle O_E , for user id , B_1 chooses $c, y \xleftarrow{U} \mathbb{Z}_p$, sets $R := g^{-zc}g^y$ and adds $\langle R, id, c \rangle$ to the table \mathcal{L}_H . Then it returns the key (y, R) to A .

¹The number of different identities involved in the G-oracle query, *i.e.* n , can be at most q_G . Hence, B_1 has to choose one index from this set.

$\mathcal{B}_1.3$ When A queries the signature oracle \mathcal{O}_s with (id, m) where $\text{id} \neq \hat{\text{id}}$, \mathcal{B}_1 simply computes id 's secret key as described in the previous bullet. Then it invokes the signing algorithm S and returns the produced signature to A.

$\mathcal{B}_1.4$ When A queries the signature oracle \mathcal{O}_s with (id, m) where $\text{id} = \hat{\text{id}}$, \mathcal{B}_1 chooses $b, d \xleftarrow{\text{U}} \mathbb{Z}_p$, sets $B := g^b, R := g^\alpha, c := \text{H}(\text{id}, R), A := B(g^\alpha g^{zc})^{-d}$ and programs the random oracle in such a way that $d := \text{G}(\text{id}, A, m)$. Then it returns the signature (A, b, R) to A.

$\mathcal{B}_1.5$ \mathcal{B}_1 invokes the algorithm $\text{M}_{\text{W}1}(\text{mpk})$ as described in Lemma 1 (§4 in [GG09]). Here algorithm W is simply a wrapper that takes as explicit input, the answers from the random oracles. Then it calls A and returns its output together with two integers I, J . These integers are the indices of A's queries to the random oracles G, H with the target identity $\hat{\text{id}}$.

$\mathcal{B}_1.6$ In this way we get two forgeries of the form $\sigma_0 = (\text{id}, m, (A, b_0, R))$ and $\sigma_1 = (\text{id}, m, (A, b_1, R))$. Let d_0 be the answer from the G-oracle given to A in the first simulation, $s_{I_0}^0$ in $\text{M}_{\text{W}1}$ and let d_1 be the second answer $s_{I_0}^1$. If the identity id is not equal to the target identity $\hat{\text{id}}$ then \mathcal{B}_1 aborts. Otherwise it terminates and outputs the attempted discrete logarithm

$$\alpha = \frac{b_0 - b_1}{d_0 - d_1} - zc.$$

A.1.2 Reduction \mathcal{B}_2

It takes as argument, the description of a group (\mathbb{G}, p, g) and a challenge g^α with $\alpha \xleftarrow{\text{U}} \mathbb{Z}_p$ and outputs the discrete logarithm α . To do so, it will invoke A simulating the environment as shown below.

$\mathcal{B}_2.1$ At the beginning of the experiment, \mathcal{B}_2 sets the master public key $\text{mpk} := (\mathbb{G}, p, g, \text{G}, \text{H})$ and $\text{msk} := (g^\alpha)$, where G, H are description of hash functions modelled as random oracles. As usual, \mathcal{B}_2 simulates these oracles with the help of two tables \mathfrak{L}_{G} and \mathfrak{L}_{H} containing the queried values together with the answers given to A.

$\mathcal{B}_2.2$ Every time A queries the key extraction oracle \mathcal{O}_{E} , for user id , \mathcal{B}_2 chooses $c, y \xleftarrow{\text{U}} \mathbb{Z}_p$, sets $R := g^{-\alpha c} g^y$ and adds $\langle R, \text{id}, c \rangle$ to the table \mathfrak{L}_{H} . Then it returns the key (y, R) to A.

$\mathcal{B}_2.3$ When A queries the signature oracle \mathcal{O}_s with (id, m) , \mathcal{B}_2 simply computes id 's secret key as described in the previous step. Then it computes a signature by calling S, adding the

respective call to the G-oracle, $((\text{id}, g^a, m), d)$ to the table \mathfrak{L}_G and gives the resulting signature to the adversary.

$\mathcal{B}_{2.4}$ B_2 invokes the algorithm $M_{W3}(\text{mpk})$. In this way either B_2 aborts prematurely or we get, for some identity id , some message m and some R , four forgeries $(\text{id}, m, (A_k, b_k, R_k))$, $k := 0, \dots, 3$. Now, two situations may arise

(a) If $R_3 = R_2 = R_1 = R_0$ ($H < G$) then, the signatures will be of the form

$$\begin{aligned} \{b_0 = \log A_0 + (\log R + c_0\alpha)d_0, b_1 = \log A_0 + (\log R + c_0\alpha)d_1, \\ b_2 = \log A_2 + (\log R + c_2\alpha)d_2, b_3 = \log A_2 + (\log R + c_2\alpha)d_3\} \end{aligned} \quad (\text{A.1})$$

B_2 solves for α using the equation

$$\alpha = \frac{(b_0 - b_1)(d_2 - d_3) - (b_2 - b_3)(d_0 - d_1)}{(c_0 - c_1)(d_0 - d_1)(d_2 - d_3)}. \quad (\text{A.2})$$

(b) Else, if $A_3 = A_2 = A_1 = A_0$ ($G < H$) then, the signatures will be of the form

$$\begin{aligned} \{b_0 = \log A + (\log R_0 + c_0\alpha)d_0, b_1 = \log A + (\log R_0 + c_1\alpha)d_0, \\ b_2 = \log A + (\log R_2 + c_2\alpha)d_2, b_3 = \log A + (\log R_2 + c_3\alpha)d_2\}. \end{aligned} \quad (\text{A.3})$$

B_2 solves for α using the equation

$$\alpha = \frac{b_0 - b_1}{d_0(c_0 - c_1)}. \quad (\text{A.4})$$

A.2 A Security Argument without Wrappers

The reductions in the subsequent sections are described in two steps. In the first step, called “Handling the queries”, we describe the protocol set-up and the methods in which adversarial queries–signature, extract and random oracle queries, are to be handled. This gives us sufficient know-how to simulate the protocol environment. In the second step, called “Solving the DLP”, we describe how the reductions use forking algorithms to solve the underlying hard problem, which in this case is the DLP. The reductions use the forking algorithm simply as a black-box to get hold of the forgeries. The actual simulation is handled by the wrapper algorithm Y in the forking algorithms, according to the plan laid down by the reductions in the first step. For details on how

the forking algorithms work, see §2.4.

Simulating the random oracles. A random oracle query is defined to be *fresh* if it is the first query involving that particular input. If a query is not fresh for an input, in order to maintain consistency, the random oracle has to respond with the same output as in the previous query on that input. We say that a fresh query does not require *programming* if the simulator can simply return a random value as the response. The crux of most security arguments involving random oracles, including ours, is the way the simulator answers the queries that require programming. In our case, random oracle programming is used to resolve the circularity involved while dealing with the *implicit* random oracle queries. A random oracle query is said to be implicit if it is not an explicit query from the adversary or the simulator. As usual, to simplify the book-keeping, all implicit random oracle queries involved in answering the extract and signature queries are put into the account of A .

A.2.1 Reduction R_1

R_1 uses the so-called “partitioning strategy”, first used by Coron in the security argument of FDH [Cor00]. The basic idea is to divide the identity-space \mathbb{I} into two disjoint sets, \mathbb{I}_ε and \mathbb{I}_s , depending upon the outcome of a biased coin. The simulator is equipped to respond to both extract and signature queries on identities from \mathbb{I}_ε . But it fails if the adversary does an extract query on any identity from \mathbb{I}_s ; it can answer only to signature queries on identities from \mathbb{I}_s . Finally, the simulator hopes that the adversary produces a forgery on an identity from \mathbb{I}_s . The optimal size of the sets is determined on analysis.

In R_1 the problem instance is embedded in the randomiser R , depending on the outcome of a biased coin. As R_1 maintains a unique R for each identity, the structure of R decides whether that identity belongs to \mathbb{I}_ε or to \mathbb{I}_s . The details follow.

Let $\Delta := (\mathbb{G}, p, g, g^\alpha)$ be the given DLP instance. R_1 sets $z \xleftarrow{U} \mathbb{Z}_p$ as the master secret key. The master public key $\text{mpk} := (\mathbb{G}, p, g, g^z, H, G)$ are released to the adversary. The hash functions H and G are modelled as random oracles. This is done with the aid of two tables, \mathfrak{L}_H and \mathfrak{L}_G .

A.2.1.1 Handling the Queries

H-oracle query. \mathfrak{L}_H contains tuples of the form

$$\langle R, r, \text{id}, c, \beta \rangle \in \mathbb{G} \times \mathbb{Z}_p \cup \{\perp\} \times \{0, 1\}^* \times \mathbb{Z}_p \times \{0, 1, \phi\}.$$

Here, (R, id) is the query to the H-oracle and c is the corresponding output. Therefore, an oracle query $H(R, \text{id})$ is fresh if there exists no tuple $\langle R_i, r_i, \text{id}_i, c_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (R_i = R)$. If such a tuple exists, then the oracle has to return c_i as the output.

The r -field is used to store additional information related to the R -field. The tuples corresponding to the explicit H-oracle queries, made by A , are tracked by storing ‘ \perp ’ in the r -field. This indicates that R_1 does not have any additional information regarding R . In these tuples, the β -field is irrelevant and this is indicated by storing ‘ ϕ ’. In tuples with $r \neq \perp$, the field β indicates whether the DLP instance is embedded in R or not. If $\beta = 0$ then $R = (g^\alpha)^r$ for some known $r \in \mathbb{Z}_p$, which is stored in the r -field. On the other hand, $\beta = 1$ implies $R = g^r$ for some known $r \in \mathbb{Z}_p$, which is, again, stored in the r -field. Therefore, the r -field stores one of the values: *i*) ‘ \perp ’, or *ii*) $\log gR$, or *iii*) r , if $R = (g^\alpha)^r$. As a result, \mathfrak{L}_H can contain three types of tuples determined by the content of r -field and β -field, viz.

1. $r = \perp$: These tuples correspond to the explicit H-oracle queries made by A .
2. $r \neq \perp \wedge \beta = 0$: These tuples correspond to identities in \mathbb{I}_s . They are added by R_1 while answering the signature queries. As the DLP instance is embedded in R , extract query fails on these identities.
3. $r \neq \perp \wedge \beta = 1$: These tuples correspond to identities in \mathbb{I}_ε . They are added by R_1 while answering signature or extract queries.

We now explain how the fresh H-oracle queries are handled.

$H(R, \text{id})$: The query may be

- (i) H_1 , Explicit query made by A : In this case R_1 returns $c \stackrel{U}{\leftarrow} \mathbb{Z}_p$ as the output. $\langle R, \perp, \text{id}, c, \phi \rangle$ is added to \mathfrak{L}_H .
- (ii) H_2 , Explicit query made by R_1 : As in the previous case, R_1 returns $c \stackrel{U}{\leftarrow} \mathbb{Z}_p$ as the output. As R_1 knows $r = \log gR$, $\langle R, r, \text{id}, c, 1 \rangle$ is added to \mathfrak{L}_H .

- (iii) H_3 , Implicit query by R_1 in order to answer a signature query made by A : See **Sign** (iii) on how to program the random oracle in this situation.

G-oracle query. \mathfrak{L}_G contains tuples of the form

$$\langle \mathbf{id}, A, m, d \rangle \in \{0, 1\}^* \times \mathbb{G} \times \{0, 1\}^* \times \mathbb{Z}_p.$$

Here, (\mathbf{id}, A, m) is the query to the G-oracle and d is the corresponding output. Therefore, a random oracle query $G(\mathbf{id}, A, m)$ is fresh if there exists no tuple $\langle \mathbf{id}_i, A_i, m_i, d_i \rangle$, in \mathfrak{L}_G such that $(\mathbf{id}_i = \mathbf{id}) \wedge (A_i = A) \wedge (m_i = m)$. If such a tuple exists, then the oracle has to return d_i as the output.

We now explain how the fresh G-oracle queries are handled.

$G(\mathbf{id}, A, m)$: The query may be

- (i) G_1 , Explicit query made by the either A or R_1 : In this case R_1 returns $d \xleftarrow{U} \mathbb{Z}_p$ as the output. $\langle \mathbf{id}, A, m, d \rangle$ is added to \mathfrak{L}_G .
- (ii) G_2 , Implicit query by R_1 in order to answer a signature query made by A : See **Sign** (i),
- (iii) on how to program the random oracle in this situation.

Remark 20. In the case of the implicit queries H_3 and G_2 , R_1 has to program the respective random oracles in an appropriate way to deal with the circularity involved. For ease of understanding, they are dealt with in their respective sections.

Now that R_1 can handle the random oracle queries, the extract and signature queries are answered as follows.

Extract query. R_1 first checks if \mathbf{id} has an associated R . This is done by searching for tuples $\langle R_i, r_i, \mathbf{id}_i, c_i, \beta_i \rangle$ in \mathfrak{L}_H with $(\mathbf{id}_i = \mathbf{id}) \wedge (r_i \neq \perp)$. If such a tuple exists, R_1 checks for the value of β_i in the tuple. $\beta_i = 0$ implies the identity belongs to \mathbb{I}_s and consequently the extract query fails, leading to an abort, $\text{abort}_{1,1}$. On the other hand, $\beta_i = 1$ implies that there was a prior extract query on \mathbf{id} and also that the identity belongs to \mathbb{I}_e . R_1 generates the secret key (same as in prior extract query) using the information available in the tuple. On the other hand, if such a tuple does not exist, R_1 selects a fresh r and assigns \mathbf{id} to \mathbb{I}_e . R_1 has this freedom since the adversary cannot forge on this identity. This is captured by the oracle O_E shown below.

$O_E(\text{id})$: If there exists a tuple $\langle R_i, r_i, \text{id}_i, c_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$

- (i) If $\beta_i = 0$, R_1 *aborts* ($\text{abort}_{1,1}$).
- (ii) Otherwise, $\beta_i = 1$ and R_1 returns $\text{usk} := (r_i + zc_i, R_i)$ as the secret key for id .

Otherwise

- (iii) R_1 chooses $r \xleftarrow{U} \mathbb{Z}_p$, sets $R := g^r$ and asks the H-oracle for $c := H(R, \text{id})$. It returns $\text{usk} := (r + zc, R)$ as the secret key.

Signature query. As in the extract query, R_1 checks the identity for an associated R by searching tuples $\langle R_i, r_i, \text{id}_i, c_i, \beta_i \rangle$ in \mathfrak{L}_H with $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$. If such a tuple exists, the identity has been assigned to either of \mathbb{I}_ε or \mathbb{I}_s , determined by the value of β_i . If such a tuple does not exist, then the identity is unassigned and R_1 assigns the identity to either \mathbb{I}_ε or \mathbb{I}_s by tossing a biased coin β . If the outcome is 0, id is assigned to \mathbb{I}_s ; else it is assigned to \mathbb{I}_ε . Identities assigned to \mathbb{I}_s have the problem instance g^α embedded in the randomiser R . Although the private key cannot be calculated, an algebraic technique, similar to one adopted by Boneh-Boyen in [BB04a], coupled with random oracle programming enables us to give the signature. On the other hand, signature queries involving identities from \mathbb{I}_ε are answered by first generating usk as in the extract query and then invoking S . This is captured by the oracle O_s described below.

$O_s(\text{id}, m)$: If there exists a tuple $\langle R_i, r_i, \text{id}_i, c_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$

- (i) If $\beta_i = 0$, R_1 selects $s, d \xleftarrow{U} \mathbb{Z}_p$ and sets $A := g^s(g^\alpha)^{-r_i d}$. Then (id, A, m, d) is added to \mathfrak{L}_G (*Deferred case G_2*)². The signature returned is

$$\sigma := (A, s + zcd, R_i).$$

- (ii) Otherwise, $\beta_i = 1$ and the secret key for id is $\text{usk} = (y, R_i)$, where $y = r_i + zc_i$ and $R_i = g^{r_i}$. R_1 selects $a \xleftarrow{U} \mathbb{Z}_p$, sets $A := g^a$ and asks G-oracle for $d := G(\text{id}, A, m)$. The signature returned is

$$\sigma := (A, a + yd, R_i).$$

²If there exists a tuple $\langle \text{id}_i, A_i, m_i, d_i \rangle$ in \mathfrak{L}_G with $\text{id}_i = \text{id} \wedge A_i = A \wedge m_i = m$ but $d_i \neq d$ then $G(\text{id}, A, m)$ cannot be set to d . In that case R_1 can simply choose a fresh set of randomisers s, d and repeat the process.

Otherwise, R_1 tosses a coin β with a bias δ (i.e., $\Pr[\beta = 0] = \delta$). The value of δ will be quantified on analysis.

- (iii) If $\beta = 0$, R_1 selects $c, d, s, r \xleftarrow{U} \mathbb{Z}_p$ and sets $R := (g^\alpha)^r$, $A := g^s(g^\alpha)^{-rd}$. Next, it adds $\langle (g^\alpha)^r, r, \text{id}, c, 0 \rangle$ to \mathcal{L}_H (*Deferred case H_3*) and $\langle \text{id}, A, m, d \rangle$ to \mathcal{L}_G (*Deferred case G_2*).³

The signature returned is

$$\sigma := (A, s + zc_d, R).$$

- (iv) Otherwise, $\beta = 1$ and R_1 selects $a, r \xleftarrow{U} \mathbb{Z}_p$ and sets $A := g^a$, $R := g^r$. It then asks the respective oracles for $c := H(R, \text{id})$ and $d := G(\text{id}, A, m)$. The signature returned is

$$\sigma := (A, a + (r + zc)d, R).$$

Correctness. For $\beta = 0$, the signature given by R_1 is of the form (A, b, R) , where $A = g^s(g^\alpha)^{-rd}$, $b = s + zcd$ and $R = (g^\alpha)^r$. R_1 also sets $c := H(R, \text{id})$ and $d := G(\text{id}, A, m)$. The signature verifies as shown below.

$$\begin{aligned} g^b &= g^{s+zcd} \\ &= g^{s-\alpha rd + \alpha rd + zcd} \\ &= g^{(s-\alpha rd)} g^{(\alpha r + zc)d} \\ &= g^s (g^\alpha)^{-rd} ((g^\alpha)^r (g^z)^c)^d \\ &= A (R (g^z)^c)^d. \end{aligned}$$

For $\beta = 1$, the signatures are generated as in the protocol. Therefore they fundamentally verify.

To conclude the queries section, we calculate the probability of the event $\neg \text{abort}_{1,1}$. R_1 aborts only when A does an extract query on an identity from \mathbb{I}_s , i.e. an identity with $\beta = 0$. Therefore, R_1 does not abort if all the extract queries are from \mathbb{I}_ε and we have

$$\mathbb{P}(\neg \text{abort}_{1,1}) = (1 - \delta)^{q_\varepsilon}. \tag{A.5}$$

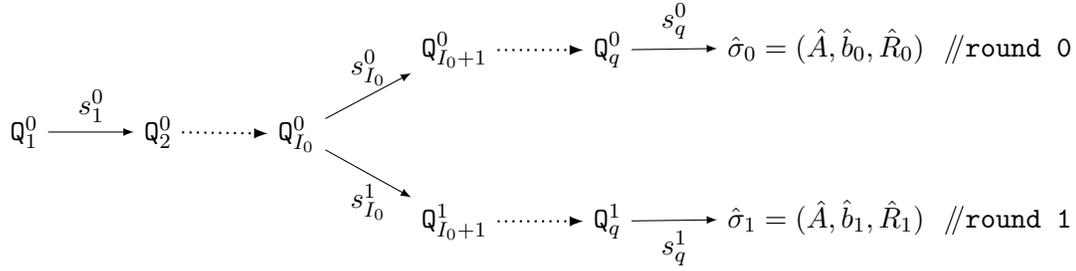


Figure A.1: Successful forking by R_1 . $Q_{I_0}^0$ denotes the target G-query $G(\hat{\text{id}}, \hat{A}, \hat{m})$.

A.2.1.2 Solving the DLP

R_1 now uses the general forking algorithm F_W (see §2.4 for details on the working of F_W) to solve the DLP challenge. It invokes F_W on the given DLP instance Δ ,⁴ with the G-oracle involved in the replay attack. If F_W fails, so does R_1 and it *aborts* ($\text{abort}_{1,2}$). On the other hand, if F_W is successful, it gets two valid forgeries

$$\{\hat{\sigma}_0 = (\hat{A}, \hat{b}_0, \hat{R}_0), \hat{\sigma}_1 = (\hat{A}, \hat{b}_1, \hat{R}_1)\}$$

on $(\hat{\text{id}}, \hat{m})$ (see **Figure A.1**). R_1 now retrieves two tuples

$$t_i := \langle R_i, r_i, \text{id}_i, c_i, \beta_i \rangle \mid (\text{id}_i = \hat{\text{id}}) \wedge (R_i = \hat{R}_0) \text{ and}$$

$$t_j := \langle R_j, r_j, \text{id}_j, c_j, \beta_j \rangle \mid (\text{id}_j = \hat{\text{id}}) \wedge (R_j = \hat{R}_1)$$

from \mathfrak{L}_H . R_1 *aborts* ($\text{abort}_{1,3}$) if both β_i and β_j are equal to 1. Otherwise it solves for α as shown below. Note that d_0 and d_1 represent the value of $G(\hat{\text{id}}, \hat{A}, \hat{m})$ in the two rounds, *i.e.*, $d_0 = s_{I_0}^0$ and $d_1 = s_{I_0}^1$. Let $\hat{a} := \log g \hat{A}$.

- (i) $(\beta_i = 1) \wedge (\beta_j = 0)$: In this case, $\hat{R}_0 = g^{r_i}$ and $\hat{R}_1 = g^{r_j \alpha}$. Thus we have $\hat{b}_0 = \hat{a} + (r_i + zc_i)d_0$

³ R_1 chooses different randomisers if there is a collision as explained in **Footnote 2**.

⁴In reductions involving the forking algorithm, the problem instance is usually embedded in mpk . Therefore the forking algorithm F_W is invoked on mpk . But in case of R_1 , the master secret key is chosen by R_1 itself. The problem instance (g^α) is embedded as part of answers to signature queries. Therefore, R_1 invokes F_W on Δ while (msk, mpk) is considered to be part of ρ .

and $\hat{b}_1 = \hat{a} + (r_j\alpha + zc_j)d_1$.

$$\begin{aligned}\hat{b}_0 - \hat{b}_1 &= (r_id_0 - r_j\alpha d_1) + z(c_id_0 - c_jd_1), \\ \alpha &= \frac{z(c_id_0 - c_jd_1) + r_id_0 - (\hat{b}_0 - \hat{b}_1)}{r_jd_1}.\end{aligned}\tag{A.6}$$

(ii) $(\beta_i = 0) \wedge (\beta_j = 1)$: In this case, $\hat{R}_0 = g^{r_i\alpha}$ and $\hat{R}_1 = g^{r_j}$. Thus we have $\hat{b}_0 = \hat{a} + (r_i\alpha + zc_i)d_0$ and $\hat{b}_1 = \hat{a} + (r_j + zc_j)d_1$.

$$\begin{aligned}\hat{b}_1 - \hat{b}_0 &= (r_jd_1 - r_i\alpha d_0) + z(c_jd_1 - c_id_0), \\ \alpha &= \frac{z(c_jd_1 - c_id_0) + r_jd_1 - (\hat{b}_1 - \hat{b}_0)}{r_id_0}.\end{aligned}\tag{A.7}$$

(iii) $(\beta_i = 0) \wedge (\beta_j = 0)$: In this case, $\hat{R}_0 = g^{r_i\alpha}$ and $\hat{R}_1 = g^{r_j\alpha}$. Thus we have $\hat{b}_0 = \hat{a} + (r_i\alpha + zc_i)d_0$ and $\hat{b}_1 = \hat{a} + (r_j\alpha + zc_j)d_1$.

$$\begin{aligned}\hat{b}_0 - \hat{b}_1 &= \alpha(r_id_0 - r_jd_1) + z(c_id_0 - c_jd_1), \\ \alpha &= \frac{(\hat{b}_0 - \hat{b}_1) - z(c_id_0 - c_jd_1)}{(r_id_0 - r_jd_1)}.\end{aligned}\tag{A.8}$$

Remark 21. The equations (A.6), (A.7) and (A.8) hold even if $\hat{R}_1 = \hat{R}_0$ (and consequently $r_j = r_i$ and $c_j = c_i$). Note that this can happen if the adversary makes the random oracle query $H(\hat{R}_0, \hat{id})$ before the query $G(\hat{id}, \hat{A}, \hat{m})$ in round 0. Hence, the order in which A makes the aforementioned random oracle queries is not relevant.

Structure of R . The event E guarantees the existence of the tuple \mathbf{t}_i in \mathfrak{L}_H . As A cannot make an extract query on \hat{id} , the choice of \hat{R}_i must have been made during a signature query, where the structure of R is determined by the coin β . Therefore

$$\hat{R}_i = \begin{cases} g^{\hat{r}_i\alpha} & \text{If } \beta_i = 0 \\ g^{\hat{r}_i} & \text{Otherwise} \end{cases}$$

This, in turn, is determined by the bias in β , *i.e.* δ . A similar argument holds for \mathbf{t}_j . R_1 is successful in solving the DLP if either of the forgeries have $\beta = 0$. If $(\beta_i = 1) \wedge (\beta_j = 1)$, R_1 fails and does $\text{abort}_{1,3}$. We conclude by calculating the probability of $\text{abort}_{1,3}$ provided $\text{abort}_{1,2}$ has not

occurred. It is same as the probability with which $(\beta_i = 1) \wedge (\beta_j = 1)$, *i.e.*

$$\mathbb{P}(\text{abort}_{1,3} \mid \neg\text{abort}_{1,2}) = (1 - \delta)^2. \quad (\text{A.9})$$

Let $gfrk$ be the probability with which F_W is successful. Since $\text{abort}_{1,2}$ occurs if F_W fails, we have

$$\mathbb{P}(\neg\text{abort}_{1,2}) = gfrk. \quad (\text{A.10})$$

A.2.1.3 Analysis

The probability analysis is done in terms of the aborts $\text{abort}_{1,1}$, $\text{abort}_{1,2}$ and $\text{abort}_{1,3}$. From (A.5), (A.10) and (A.9), we have $\mathbb{P}(\neg\text{abort}_{1,1}) = (1 - \delta)^{q_\epsilon}$, $\mathbb{P}(\neg\text{abort}_{1,2}) = gfrk$ and $\mathbb{P}(\text{abort}_{1,3} \mid \neg\text{abort}_{1,2}) = (1 - \delta)^2$. F_W is successful during **round 0** if there is no abort during the query phase ($\neg\text{abort}_{1,1}$) and A produces a valid forgery. We denote this probability by acc_1 . Thus

$$\begin{aligned} acc_1 &\geq \mathbb{P}(\neg\text{abort}_1) \cdot \epsilon \\ &\geq (1 - \delta)^{q_\epsilon} \cdot \epsilon. \end{aligned}$$

Applying the GF lemma (**Lemma 2.4**) with $|\mathbb{S}| = p$ and $q = q_G$, we get

$$\begin{aligned} gfrk &\geq acc_1 \cdot \left(\frac{acc_1}{q_G} - \frac{1}{p} \right) \\ &\geq (1 - \delta)^{q_\epsilon} \epsilon \cdot \left(\frac{(1 - \delta)^{q_\epsilon} \epsilon}{q_G} - \frac{1}{p} \right). \end{aligned}$$

If we assume $p \gg 1$, the above expression approximates to

$$gfrk \geq \frac{(1 - \delta)^{2q_\epsilon} \epsilon^2}{q_G}.$$

Now, R_1 is successful in solving DLP if neither of the aborts, $\text{abort}_{1,2}$ and $\text{abort}_{1,3}$, occur. Thus the advantage it has is

$$\begin{aligned}
\epsilon_1 &= \mathbb{P}(\neg \text{abort}_{1,3} \wedge \neg \text{abort}_{1,2}) \\
&= \mathbb{P}(\neg \text{abort}_{1,3} \mid \neg \text{abort}_{1,2}) \cdot \mathbb{P}(\neg \text{abort}_{1,2}) \\
&\geq (1 - (1 - \delta)^2) \cdot \text{gfrk} \\
&\geq (2\delta - \delta^2) \frac{(1 - \delta)^{2q_\epsilon} \epsilon^2}{q_G}.
\end{aligned} \tag{A.11}$$

Assuming $p \gg 1$, (3.9) attains maximum value at the point $\delta = \left(1 - \sqrt{q_\epsilon / (q_\epsilon + 1)}\right)$, at which

$$\epsilon_1 \geq \frac{\epsilon^2}{\exp(1) q_G q_\epsilon}.$$

Here, \exp is the base of natural logarithm.

Remark 22. The above reduction is tighter than the reduction B_1 given by Galindo and Garcia [GG09]. This can be attributed to two reasons: *i*) R_1 using the GF Algorithm F_W instead of the MF Algorithm M_W1 ; and *ii*) B_1 in [GG09] randomly chooses one of the identities involved in the G-oracle query as the target identity (refer to §3.2.2.1) which contributes a factor of q_G to the degradation in B_1 . In contrast, we apply Coron’s technique in R_1 to partition the identity space in some optimal way.

Time complexity. If τ is the time taken for an exponentiation in \mathbb{G} then the time taken by R_1 is $t_1 \leq t + 2(q_\epsilon + 3q_s)\tau$. It takes at most one exponentiation for answering the extract query and three exponentiations for answering the signature query. This contributes the $(q_\epsilon + 3q_s)\tau$ factor in the running time. The factor of two comes from the forking algorithm, since it involves simulating the adversary twice.

A.2.2 Reduction R_2

The reduction R_2 is similar in some aspects to the (incomplete) reduction argument B_2 in [GG09]. However, a *major* difference is that R_2 uses the multiple-forking algorithm M_W1 instead of M_W3 to solve the DLP challenge. Therefore, only one forking is involved leading to a much tighter reduction than B_2 . The details follow.

Let $\Delta := (\mathbb{G}, p, g, g^\alpha)$ be the given DLP instance. R_2 sets $\text{mpk} := (\mathbb{G}, p, g, g^\alpha, H, G)$ as the master public key and releases it to A . Note that R_2 does not know the master secret key msk , which is α , the solution to the DLP challenge. The hash functions H and G are modelled as random oracles. This is done with the aid of two tables, \mathcal{L}_H and \mathcal{L}_G .

A.2.2.1 Handling the Queries

H-oracle query. \mathcal{L}_H contains tuples of the form

$$\langle R, \text{id}, c, y \rangle \in \mathbb{G} \times \{0, 1\}^* \times \mathbb{Z}_p \times \mathbb{Z}_p \cup \{\perp\}.$$

Here, (R, id) is the query to the H-oracle and c the corresponding output. The y -field stores either the corresponding component of the secret key for id or ' \perp ' if the field is invalid. A random oracle query $H(R, \text{id})$ is fresh if there exists no tuple $\langle R_i, \text{id}_i, c_i, y_i \rangle$ in \mathcal{L}_H such that $(\text{id}_i = \text{id}) \wedge (R_i = R)$. If such a tuple exists, then the oracle has to return c_i as the output. We now explain how the H-oracle queries are answered.

$H(R, \text{id})$: The query may be

- (i) H_1 , Explicit query made by A : In this case R_2 returns $c \xleftarrow{U} \mathbb{Z}_p$ as the output. $\langle R, \text{id}, c, \perp \rangle$ is added to \mathcal{L}_H .
- (ii) H_2 , Implicit query by R_2 in order to answer an extract query made by A : See **Extract** (ii) on how to program the random oracle in this situation.

G-oracle query. \mathcal{L}_G has the same structure as in R_1 (See §A.2.1.1). The queries to G-oracle are handled as shown below.

$G(\text{id}, A, m)$: R_2 returns $d \xleftarrow{U} \mathbb{Z}_p$ as the output. $\langle \text{id}, A, m, d \rangle$ is added to \mathcal{L}_G .

Signature and Extract queries. Since R_2 does not know the master secret key α , it has to use the algebraic technique used in R_1 to come up with the secret key corresponding to an identity. The choice of R and c enables it to give the secret key. The circularity involved in this choice is resolved by programming the H-oracle appropriately. Signature queries are answered by generating usk as in the extract query, followed by calling S .

Extract query, $O_E(\text{id})$:

- (i) If there exists a tuple $\langle R_i, \text{id}_i, c_i, y_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (y_i \neq \perp)$, R_2 returns $\text{usk} := (y_i, R_i)$ as the secret key.
- (ii) Otherwise, R_2 chooses $c, y \xleftarrow{U} \mathbb{Z}_p$, sets $R := (g^\alpha)^{-c} g^y$ and adds $\langle R, \text{id}, c, y \rangle$ to \mathfrak{L}_H (*Deferred case H_2*). It returns $\text{usk} := (y, R)$ as the secret key.

Signature query, $O_s(\text{id}, m)$:

- (i) If there exists a tuple $\langle R_i, \text{id}_i, c_i, y_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (y_i \neq \perp)$, then $\text{usk} = (y_i, R_i)$. R_2 now uses the knowledge of usk to invoke S and returns the signature.
- (ii) Otherwise, R_2 generates usk as in **Extract**(ii) and invokes S to return the signature.

We conclude the queries section with the remark that R_2 never aborts during the query stage.

A.2.2.2 Solving the DLP

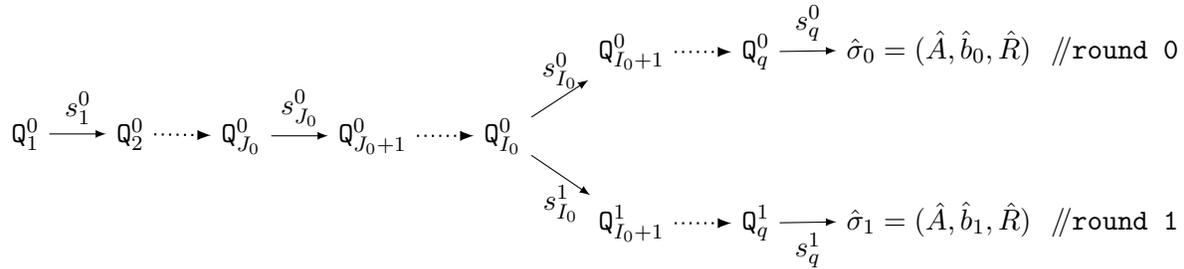


Figure A.2: Successful forking by R_2 . $Q_{J_0}^0$ denotes the target G-query $G(\hat{\text{id}}, \hat{A}, \hat{m})$ while $Q_{I_0}^0$ denotes the target H-query $H(\hat{R}, \hat{\text{id}})$.

R_2 now uses the multiple-forking algorithm M_{W1} (see §2.5 for details on the working of M_{Wn}) to solve the DLP challenge. It invokes M_{W1} on mpk , with both H and G-oracle involved in the replay attack. If M_{W1} fails, so does R_2 and it *aborts* ($\text{abort}_{2,1}$). On the other hand, if M_{W1} is successful, R_2 gets two valid, non-trivial forgeries $\{\hat{\sigma}_0, \hat{\sigma}_1\}$ on $(\hat{\text{id}}, \hat{m})$ with $\hat{\sigma}_i := (\hat{A}, \hat{b}_i, \hat{R})$ for each $i = 0, 1$ and

$$\{\hat{b}_0 = \hat{a} + (\hat{r} + \alpha c_0)\hat{d}, \hat{b}_1 = \hat{a} + (\hat{r} + \alpha c_1)\hat{d}\} \quad (\text{A.12})$$

where $\hat{a} := \log g\hat{A}$ and $\hat{r} := \log g\hat{R}$. Note that c_0 and c_1 represent the value of $H(\hat{R}, \hat{\text{id}})$ in the two rounds, *i.e.*, $c_0 = s_{I_0}^0$ and $c_1 = s_{I_0}^1$. The event F guarantees that A makes the G-oracle query $Q_{J_0}^0 : G(\hat{\text{id}}, \hat{A}, \hat{m})$, before the H-oracle query $Q_{I_0}^0 : H(\hat{R}, \hat{\text{id}})$. Finally it outputs the solution to the

DLP instance,

$$\alpha = \frac{\hat{b}_0 - \hat{b}_1}{\hat{d}(c_0 - c_1)}. \quad (\text{A.13})$$

Structure of the forgeries. Now we justify the structure of the component b of the forgeries given in (A.12). Recall that the signature queries are answered by doing an extract query on the identity followed by calling S . Therefore, the resultant secret keys are of the form $\mathbf{usk} = (y, R)$, where $R = (g^\alpha)^{-c}g^y$ and we have $r = -\alpha c + y$. If a forgery is produced using the same R as given by R_2 as part of the signature query on id , then b will be of the form $b = a + (-\alpha c + y + \alpha c)d = a + yd$. Therefore, it will not contain the solution to the DLP challenge α , and such forgeries are of no use to R_2 . But the event $\neg E$ guarantees that A does not forge using an R which was given as part of the signature query on id and hence, for the forgery to be valid b will necessarily be of the form:

$$b = a + (r + \alpha c)d. \quad (\text{A.14})$$

We conclude with the remark that the event $\text{abort}_{2,1}$ does not occur if the multiple-forking algorithm is successful (let this probability be $mfrk$). Therefore

$$\mathbb{P}(\neg \text{abort}_{2,1}) = mfrk. \quad (\text{A.15})$$

A.2.2.3 Analysis

The only abort involved in R_2 is $\text{abort}_{2,1}$, which occurs when M_W1 fails. Therefore R_2 is successful if M_W1 is and from (A.15) we have

$$\epsilon_2 = \mathbb{P}(\neg \text{abort}_{2,1}) = mfrk.$$

We denote the probability with which M_W1 is successful during round 0 as acc_2 . Since there is no abort involved during query phase, M_W1 is successful during round 0 if A produces a valid forgery, *i.e.* $acc_2 = \epsilon$. Applying the MF lemma (**Lemma 6**) with $n := 1$, $q := q_H + q_G$ and $|\mathbb{S}| = p$, we get

$$\begin{aligned} \epsilon_2 = mfrk &\geq acc_2 \cdot \left(\frac{acc_2}{(q_H + q_G)^2} - \frac{1}{p} \right) \\ &\geq \epsilon \left(\frac{\epsilon}{(q_H + q_G)^2} - \frac{1}{p} \right). \end{aligned}$$

Time complexity. Drawing analogy from the analysis of time complexity of R_1 , the time taken by R_2 is easily seen to be bounded by $t_2 \leq t + 2(2q_\varepsilon + 3q_s)\tau$.

A.2.3 Reduction R_3

The approach used in R_3 is the same as in the reduction B_2 in [GG09]. Let $\Delta := (\mathbb{G}, p, g, g^\alpha)$ be the given DLP instance. R_3 sets $\text{mpk} := (\mathbb{G}, p, g, g^\alpha, H, G)$ as the master public key and releases it to A . As in R_2 , R_3 does not know the master secret msk , which is α . The hash functions H and G are modelled as random oracles. This is done with the aid of two tables, \mathfrak{L}_H and \mathfrak{L}_G .

A.2.3.1 Handling the Queries

The queries are handled in the same way as in R_2 . So we refer to §A.2.2.1 for details.

A.2.3.2 Solving the DLP

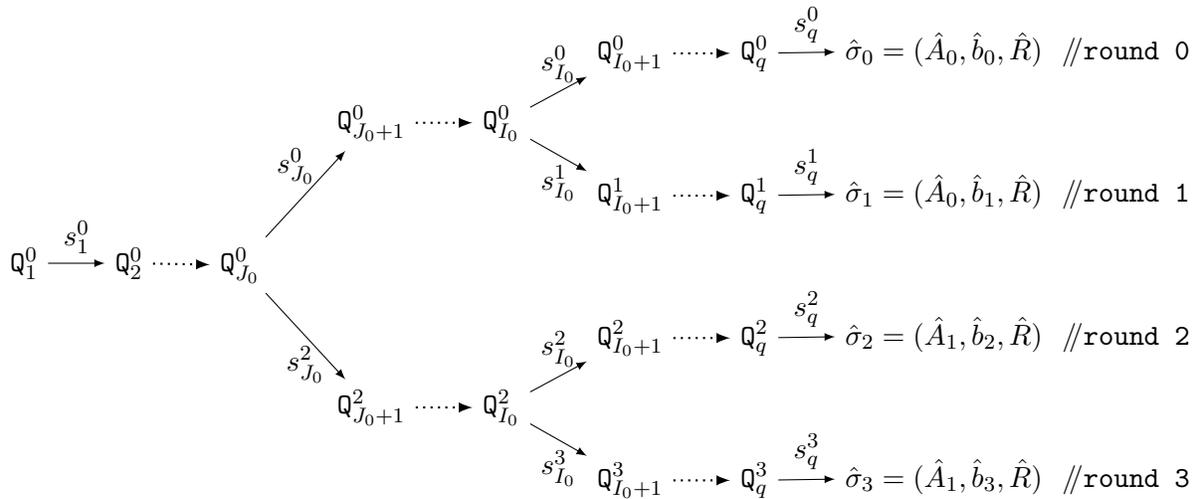


Figure A.3: Successful multiple forkings by R_3 . $Q_{J_0}^0$ denotes the target H-query $H(\hat{R}, \hat{\text{id}})$; $Q_{I_0}^0$ (resp $Q_{I_0}^2$) denotes the target G-query $G(\hat{\text{id}}, \hat{A}_0, \hat{m})$ (resp $G(\hat{\text{id}}, \hat{A}_1, \hat{m})$).

R_3 now uses the multiple-forking algorithm M_{W3} , to solve the DLP challenge. It invokes M_{W3} on mpk , with both H and G -oracle involved in the replay attack. If M_{W3} fails, so does R_3 and it *aborts* ($\text{abort}_{3,1}$). On the other hand, if M_{W3} is successful (see **Figure A.3**), R_3 gets four valid

forgeries

$$\begin{aligned}\hat{\sigma}_0 &= (\hat{A}_0, \hat{b}_0, \hat{R}), \quad \hat{\sigma}_1 = (\hat{A}_0, \hat{b}_1, \hat{R}), \\ \hat{\sigma}_2 &= (\hat{A}_1, \hat{b}_2, \hat{R}) \quad \text{and} \quad \hat{\sigma}_3 = (\hat{A}_1, \hat{b}_3, \hat{R})\end{aligned}\tag{A.16}$$

with $\hat{\sigma}_0$ and $\hat{\sigma}_1$ on $(\hat{\text{id}}, \hat{m}_0)$ and $\hat{\sigma}_2$ and $\hat{\sigma}_3$ on $(\hat{\text{id}}, \hat{m}_1)$, where

$$\begin{aligned}\hat{b}_0 &= \hat{a}_0 + (\hat{r} + \alpha c_0)d_0, \quad \hat{b}_1 = \hat{a}_0 + (\hat{r} + \alpha c_0)d_1, \\ \hat{b}_2 &= \hat{a}_1 + (\hat{r} + \alpha c_1)d_2 \quad \text{and} \quad \hat{b}_3 = \hat{a}_1 + (\hat{r} + \alpha c_1)d_3,\end{aligned}\tag{A.17}$$

where $\hat{r} := \log g\hat{R}$, $\hat{a}_i := \log g\hat{A}_i$. Note that c_0 and c_1 represent the value of $H(\hat{R}, \hat{\text{id}})$ in the H-oracle forks; d_0 and d_1 represent the value of $G(\hat{\text{id}}, \hat{A}_0, \hat{m})$ in the first two rounds; d_2 and d_3 represent the value of $G(\hat{\text{id}}, \hat{A}_1, \hat{m})$ in the last two runs. Finally it outputs the solution to the DLP challenge,

$$\alpha = \frac{(\hat{b}_0 - \hat{b}_1)(d_2 - d_3) - (\hat{b}_2 - \hat{b}_3)(d_0 - d_1)}{(c_0 - c_1)(d_0 - d_1)(d_2 - d_3)}.\tag{A.18}$$

We conclude with the remark that the event $\text{abort}_{3,1}$ does not occur if the multiple-forking algorithm is successful (let this probability be $mfrk$). Therefore

$$\mathbb{P}(\neg \text{abort}_{3,1}) = mfrk.\tag{A.19}$$

A.2.3.3 Analysis

As in R_2 , the only abort involved in R_3 is $\text{abort}_{3,1}$, which occurs when $M_{\mathbb{W}3}$ fails. Therefore R_3 is successful if $M_{\mathbb{W}3}$ is and from (A.19) we have

$$\epsilon_3 = \mathbb{P}(\neg \text{abort}_{3,1}) = mfrk.$$

We denote the probability with which $M_{\mathbb{W}3}$ is successful during the first round as acc_3 . Since there is no abort involved during query phase, $M_{\mathbb{W}3}$ is successful during the first round if A produces a valid forgery, *i.e.* $acc_3 = \epsilon$. Applying the MF lemma (**Lemma 6**) with $n = 3$, $q = q_H + q_G$ and

$|\mathbb{S}| = p$, we have

$$\begin{aligned} \epsilon_3 = mfrk &\geq acc_3 \cdot \left(\frac{acc_3^3}{(q_H + q_G)^6} - \frac{3}{p} \right) \\ &\geq \epsilon \left(\frac{\epsilon^4}{(q_H + q_G)^6} - \frac{3}{p} \right). \end{aligned}$$

Time complexity. The time taken by R_3 is easily seen to be bounded by $t_3 \leq t + 4(2q_\epsilon + 3q_s)\tau$.

A.3 Reduction R'_1

Let $\Delta := (\mathbb{G}, p, g, g^\alpha)$ be the given DLP instance. The reduction involves invoking the GF Algorithm on the wrapper Y as shown in **Algorithm 7**. As a result, it obtains a set of two congruences in two unknowns and solves for α .

Algorithm 7 Reduction $R'_1(\Delta)$

```

Select  $z \xleftarrow{U} \mathbb{Z}_p^*$  as the msk and set mpk :=  $(\mathbb{G}, g, p, g^z)$ .
 $(op, \sigma_0, \sigma_1) \xleftarrow{\$} F_Y((\mathbf{mpk}, \mathbf{msk}), g^\alpha)$ 
if  $(op = 0)$  then  $\perp$  //abort1,2
Parse  $\sigma_i$  as  $(\hat{b}_i, c_i, r_i, \beta_i, d_i)$ .
Let  $\beta := \beta_0$ ,  $c := c_0$  and  $r := r_0$ 
if  $\beta = 0$  then return  $((\hat{b}_0 - \hat{b}_1) - zc(d_0 - d_1))/r(d_0 - d_1)$ 
else return  $\perp$  //abort1,3
end if

```

The Wrapper

Suppose that $q := q_G$ and $\mathbb{S} := \mathbb{Z}_p$. Y takes as input the master keys $(\mathbf{mpk}, \mathbf{msk})$, the problem instance g^α and s_1, \dots, s_q . It returns a pair (I, σ) where I is an integer that refers to the target G-query and σ is the side-output. In order to track the index of the current G-oracle query, Y maintains a counter ℓ , initially set to 1. It also maintains a table \mathfrak{L}_H (resp. \mathfrak{L}_G) to manage the random oracle H (resp. G). Y initiates the EU-ID-CMA game by passing **mpk** as the challenge master public key to the adversary A. The queries by A are handled as per the following specifications.

(a) **Random oracle query**, $H(\text{id}, R)$: \mathfrak{L}_H contains tuples of the form

$$\langle \text{id}, R, c, r, \beta \rangle \in \{0, 1\}^* \times \mathbb{G} \times \mathbb{Z}_p \times \mathbb{Z}_p \cup \{\perp\} \times \{0, 1, \phi\}.$$

Here, (id, R) is the query to the H-oracle and c is the corresponding output. Therefore, a query $H(\text{id}, R)$ is fresh if there exists no tuple $\langle \text{id}_i, R_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (R_i = R)$. If such a tuple exists, then the oracle has to return the corresponding c_i as the output. The role of the r and β -field is the same as in R_1 . We now explain how the fresh H-oracle queries are handled. The query may be

- (i) H_1 , Explicit query made by A: In this case Y returns $c \xleftarrow{U} \mathbb{Z}_p$ as the output. $\langle \text{id}, R, c, \perp, \phi \rangle$ is added to \mathfrak{L}_H .
- (ii) H_2 , Explicit query made by Y: As in the previous case, Y returns $c \xleftarrow{U} \mathbb{Z}_p$ as the output. As Y knows $r = \log gR$, $\langle \text{id}, R, c, r, 1 \rangle$ is added to \mathfrak{L}_H .
- (iii) H_3 , Implicit query by Y in order to answer a signature query made by A: See step (iii) of **Signature query** on how to program the random oracle in this situation.

(b) **Random oracle query**, $G(m, A, c)$: \mathfrak{L}_G contains tuples of the form

$$\langle m, A, c, d, \ell \rangle \in \{0, 1\}^* \times \mathbb{G} \times \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}^+.$$

Here, (m, A, c) is the query to the G-oracle and d is the corresponding output. The index of the query is stored in the ℓ -field. Therefore, a random oracle query $G(m, A, c)$ is fresh if there exists no tuple $\langle m_i, A_i, c_i, d_i, \ell_i \rangle$, in \mathfrak{L}_G such that $(m_i = m) \wedge (A_i = A) \wedge (c_i = c)$. If such a tuple exists, then the oracle has to return the corresponding d_i as the output. We now explain how the fresh G-oracle queries are handled. The query may be

- (i) G_1 , Explicit query made by either A or Y: In this case Y returns $d := s_\ell$ as the output. $\langle m, A, c, d, \ell \rangle$ is added to \mathfrak{L}_G and ℓ is incremented by one.
- (ii) G_2 , Implicit query by Y in order to answer a signature query made by A: See steps (i) and (iii) of **Signature query** on how to program the random oracle in this situation.

(c) **Extract query**, $O_E(\text{id})$:

If there exists a tuple $\langle \text{id}_i, R_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$

- (i) If $\beta_i = 0$, Y *aborts* ($\text{abort}_{1,1}$).
- (ii) Otherwise, $\beta_i = 1$ and Y returns $\text{usk} := (r_i + zc_i, R_i)$ as the user secret key.

Otherwise

- (iii) Υ chooses $r \xleftarrow{\text{U}} \mathbb{Z}_p$, sets $R := g^r$ and queries the H-oracle for $c := \text{H}(\text{id}, R)$. It returns $\text{usk} := (r + zc, R)$ as the secret key.

(d) **Signature query**, $\text{O}_s(\text{id}, m)$:

If there exists a tuple $\langle \text{id}_i, R_i, c_i, r_i, \beta_i \rangle$ in \mathfrak{L}_H such that $(\text{id}_i = \text{id}) \wedge (r_i \neq \perp)$

- (i) If $\beta_i = 0$, Υ selects $s, \ell \xleftarrow{\text{U}} \mathbb{Z}_p$ and sets $d := s_\ell$, $A := g^s(g^\alpha)^{-r_i d}$. It then adds $\langle m, A, c, d, \ell \rangle$ to \mathfrak{L}_G (deferred case G_2)⁵ and increments ℓ by one. The signature returned is $\sigma := (A, s + zcd, R_i)$.
- (ii) Otherwise, $\beta_i = 1$ and the user secret key is $\text{usk} := (y, R_i)$, where $y = r_i + zc_i$ and $R_i = g^{r_i}$. Υ then selects $a \xleftarrow{\text{U}} \mathbb{Z}_p$, sets $A := g^a$ and queries the G-oracle with $d := \text{G}(m, A, c_i)$. The signature returned is $\sigma := (a + yd, R_i, A_i)$.

Otherwise, Υ tosses a coin β with a bias δ (i.e., $\text{Pr}[\beta = 0] = \delta$). The value of δ will be quantified on analysis.

- (iii) If $\beta = 0$, Υ selects $c, s, r \xleftarrow{\text{U}} \mathbb{Z}_p$ and sets $d := s_\ell$, $R := (g^\alpha)^r$, $A := g^s(g^\alpha)^{-rd}$. Next, it adds $\langle \text{id}, (g^\alpha)^r, c, r, 0 \rangle$ to \mathfrak{L}_H (deferred case H_3), $\langle m, A, c, d, \ell \rangle$ to \mathfrak{L}_G (deferred case G_2) and increments ℓ by one.⁶ The signature returned is $\sigma := (s + zc_i d, R, A)$.
- (iv) Otherwise, $\beta = 1$ and Υ selects $a, r \xleftarrow{\text{U}} \mathbb{Z}_p$ and sets $A := g^a, R := g^r$. It then queries the respective oracles with $c := \text{H}(\text{id}, R)$ and $d := \text{G}(m, A, c)$. The signature returned is $\sigma := (a + (r + zc)d, R, A)$.

At the end of the simulation, a successful adversary outputs a valid forgery $\hat{\sigma} := (\hat{b}, \hat{R}, \hat{A})$ on a $(\hat{\text{id}}, \hat{m})$. Let $\langle \text{id}_j, R_j, c_j, r_j, \beta_j \rangle$ be the tuple in \mathfrak{L}_H that corresponds to the target H-query. Similarly, let $\langle m_i, A_i, c_i, d_i, \ell_i \rangle$ be the tuple in \mathfrak{L}_G that corresponds to the target G-query. Υ returns $(\ell_i, (\hat{b}, c_j, r_j, \beta_j, d_i))$ as its own output. Note that the side-output σ consists of $(\hat{b}, c_j, r_j, \beta_j, d_i)$. That concludes the description of the wrapper.

Correctness of the discrete-log. In the event of successful forking, \mathcal{R}'_1 obtains two (related) sets of side-outputs σ_0 and σ_1 , where σ_i (for $i = 1, 2$) is of the form $(\hat{b}_i, c_i, r_i, \beta_i, d_i)$. Due to dependency, the adversary is forced to follow the order $\text{H} < \text{G}$. As a result, $\beta_0 = \beta_1$ (denoted by β), $c_0 = c_1$ (denoted by c) and $r_0 = r_1$ (denoted by r). In addition, let \hat{a} denote $\log g \hat{A}_1 = \log g \hat{A}_0$.

⁵ In the unlikely event of there already existing a tuple $\langle m_i, A_i, c_i, d_i, \ell_i \rangle$ in \mathfrak{L}_G with $(m_i = m) \wedge (A_i = A) \wedge (c_i = c)$ but $(d_i \neq d)$ then $\text{G}(m, A, c)$ cannot be set to d . In that case Υ can simply increment ℓ and repeat step (i).

⁶ Υ chooses different randomisers if there is a collision as explained in **Footnote 5**.

Y aborts in the event that $\beta = 1$ ($\text{abort}_{1,3}$). In the case $\beta = 0$, R'_1 ends up with a system of two congruences $\{\hat{b}_0 = \hat{a} + (r\alpha + zc)d_0, \hat{b}_1 = \hat{a} + (r\alpha + zc)d_1\}$ in two unknowns $\{\hat{a}, \alpha\}$. α can be solved for as shown below.

$$\alpha = \frac{(\hat{b}_0 - \hat{b}_1) - zc(d_0 - d_1)}{r(d_0 - d_1)} \quad (\text{A.20})$$

Notice that (A.20) is precisely what R'_1 outputs in **Algorithm 7**.

A.3.1 Analysis

The probability analysis is governed by the three events $\text{abort}_{1,1}$, $\text{abort}_{1,2}$ and $\text{abort}_{1,3}$. The accepting probability of Y is the same as for R_1 and, as a result, so is the probability of $\text{abort}_{1,2}$. The probability of event $\text{abort}_{1,3}$, on the other hand, is the same as that with which ($\beta = 1$), *i.e.* $\mathbb{P}(\text{abort}_{1,3} \mid \neg \text{abort}_{1,2}) = (1 - \delta)$. On putting it all together, we get

$$\begin{aligned} \epsilon'_1 &= \mathbb{P}(\neg \text{abort}_{1,3} \wedge \neg \text{abort}_{1,2}) \\ &\geq (1 - (1 - \delta)) \cdot (1 - \delta)^{q_\epsilon} \cdot \left(\frac{(1 - \delta)^{q_\epsilon} \epsilon}{q_G} - \frac{1}{p} \right) \\ &= \delta \cdot (1 - \delta)^{q_\epsilon} \cdot \left(\frac{(1 - \delta)^{q_\epsilon} \epsilon}{q_G} - \frac{1}{p} \right) \end{aligned} \quad (\text{A.21})$$

Assuming $p \gg 1$, (A.21) attains maximum value at the point $\delta = 1/(1 + 2q_\epsilon)$, at which

$$\epsilon'_1 \geq \frac{\epsilon^2}{2 \exp(1) q_G q_\epsilon}.$$

Here, \exp is the base of natural logarithm.