On the Complexity of Cryptographic Groups and Generic Group Models

Keyu Ji[1,2], Cong Zhang[1,2(✉)], Taiyu Wang[1,2], Bingsheng Zhang[1,2(✉)], Hong-Sheng Zhou[3(✉)], Xin Wang[4], and Kui Ren[1,2]

[1] the State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, China
{jikeyu,congresearch,taiyuwang,bingsheng,kuiren}@zju.edu.cn
[2] Hangzhou High-Tech Zone (Binjiang) Blockchain and Data Security Research Institute, Hangzhou, China
[3] Virginia Commonwealth University, Richmond, USA
hszhou@vcu.edu
[4] Digital Technologies, Ant Group, Hangzhou, China
wx352699@antgroup.com

Abstract. Ever since the seminal work of Diffie and Hellman, cryptographic (cyclic) groups have served as a fundamental building block for constructing cryptographic schemes and protocols. The security of these constructions can often be based on the hardness of (cyclic) group-based computational assumptions. Then, the generic group model (GGM) has been studied as an idealized model (Shoup, EuroCrypt 1997), which justifies the hardness of many (cyclic) group-based assumptions and shows the limits of some group-based cryptosystems. We stress that, the importance of the *length* of group encoding, either in a concrete group-based construction or assumption, or in the GGM, has not been studied.

In this work, we initiate a systematic study on the complexity of cryptographic groups and generic group models, varying in different lengths of group encodings, and demonstrate evidences that "the length matters". More concretely, we have the following results:

- We show that there is no black-box/relativizing reduction from the CDH-secure groups (i.e., over such groups, the computational Diffie-Hellman assumption holds) with shorter encodings, to the CDH-secure groups with longer encodings, within the same security parameter. More specifically, given any arbitrary longer CDH-secure group, it is impossible to generically shorten the group encoding and obtain a shorter CDH-secure group within the same group order.
- We show that there is a strict hierarchy of the GGMs with different lengths of encodings. That is, in the framework of indifferentiability, the shorter GGM is *strictly stronger* than the longer ones, even in the presence of computationally *bounded* adversaries.

The work was mainly supported by National Key Research and Development Program of China, Grant No. 2023YFB3106000. Cong Zhang is the co-first author.

K.-M. Chung and Y. Sasaki (Eds.): ASIACRYPT 2024, LNCS 15490, pp. 3–35, 2025.
https://doi.org/10.1007/978-981-96-0941-3_1

1 Introduction

Provable Security and Black-Box Reduction. In the past decades, *provable security* becomes one of the cornerstones of modern cryptography. As the main technique of provable security, reductions are involved to justify the security of a scheme based on a cryptographic primitive. Essentially, given an allegedly successful adversary that breaks the scheme, one can convert it into another successful adversary against the underlying primitive. To a large extent, we study the reductions that are in a black-box manner, in the sense that reductions consider the primitive and/or the adversary against the scheme only via the input-output behavior, without exploring the internal code of the primitive or of the adversary.

In the realm of group-based cryptography (initiated by Diffie and Hellman in their seminal work [DH76]), reductions are established based on the security of cryptographic groups. Serving as the foundation, the community is motivated to study cryptographic groups from various perspectives.

From an Efficiency Perspective. In the literature, with few exceptions, group-based cryptosystems are often built on cryptographic groups in an abstract and black-box manner, which means the underlying groups can be instantiated by any concrete ones as long as the desired security properties are fulfilled. For instance, the well-known public key encryption (PKE) scheme, the ElGamal encryption [ElG85], is chosen-plaintext attack secure (IND-CPA) w.r.t. any concrete prime-order cyclic group in which the decisional Diffie-Hellman (DDH) assumption holds.

In practice, when it comes to instantiating cryptosystems for better efficiency, we typically prefer concrete groups with shorter descriptions. Specifically, the ElGamal encryption utilizes the prime-order subgroup of $\mathbb{Z}_p^*$, for prime p, where the typical bit-length of a group element is 3072 (for 128-bit security) [Bar20]; an alternative approach involves elliptic curves, an increasingly popular choice, and NIST SP 800-186 [CMR+23] provides a list of recommended curves for 128-bit security, such as Curve25519 (with a 255-bit prime modulus). With classic point compression technique, each group element of Curve25519 can be encoded in 256 bits.

This highlights a critical yet subtle issue that has long been overlooked by the community. That is, *the bit-length of a group element is not explicitly taken into account* when the group is utilized in a black-box manner. Note, in real-world applications, groups with shorter descriptions are often preferred to minimize communication and computation overhead. Hereby, we ask the following questions: Does the length of the group description matter when using it in a black-box manner? Is it possible to construct a group with a shorter description generically from groups with longer descriptions? For notation simplicity, throughout this work, we will use *shorter groups* and *longer groups*, to denote "groups with shorter descriptions" and "groups with longer descriptions," respectively.

From a Security Perspective. Unfortunately, despite the advancement of modern cryptography, to the best of our knowledge, there is a fundamental limitation

in provable security—the inability of establishing *unconditional hardness* with respect to a *concrete* group. In the past decades, researchers have made significant efforts to explore various ways to demonstrate the hardness of those group-based problems, and one approach is through the class of generic algorithms.

In essence, generic algorithms do not explore the specific encoding of group elements, but instead treat them in a generic manner. Studying this class of algorithms is highly motivated, since several well-known algorithms such as the baby-step/giant-step algorithm [PH78] and Pollard's rho algorithm [Pol78] fall within this classification. To formally describe generic algorithms, ever since the initial work by Nechaev [Nec94], variants of generic group models (GGMs) have been proposed. In Shoup's GGM [Sho97], the group is conceptualized as a random injective encoding from the additive group $\mathbb{Z}_N$ into bit strings uniformly sampled from a set S, where algorithms are allowed to retrieve group encodings and perform group operations, through oracle access. In Maurer's GGM [Mau05], the group is modeled as pointers with respect to a stateful register, where group encodings are the handles (or the indexes) of the register. Within both models, we can establish the unconditional hardness, affirming the justification of the security of cryptographic groups.

When it comes to the study of the lengths of group encodings in the GGMs, Maurer, Portmann, and Zhu [MPZ20] initiate the models varying in the length of the group encoding, and illustrate a *partial hierarchy* of the GGMs, wherein any adversary within the GGM with a longer group encoding (below we denote it as "the longer GGM" for simplicity) can be converted into an adversary within the GGM with a shorter group encoding (below we denote it as "the shorter GGM" for simplicity). Despite the partial hierarchy, the connection and distinction between the longer GGM and the shorter GGM remains unexplored, which hinders a comprehensive interpretation and comparison of the numerous positive and negative results in the GGM.

To deepen our understanding on cryptographic groups, we ask the following question:

> *Will the longer group/GGM and the shorter group/GGM yield the same complexity?*

1.1 Our Results

In this work, we initiate a fine-grained study of cryptographic groups and generic group models with different lengths of group encodings. Specifically, we give evidences that:

- There is a black-box separation between the shorter CDH-secure groups and the longer CDH-secure groups with the same security parameter; in other words, given longer CDH-secure groups, one cannot build a shorter CDH-secure group with the same group order from any standard techniques;
- The shorter GGMs are strictly stronger than the longer GGMs, even in the presence of *computationally bounded* adversaries.

To illustrate our findings, we first formalize the notions of groups/GGMs, parameterized by (N, m)[1], where N and m denote the order of the group and the length of the group encodings[2], respectively. More concretely, we respectively denote the (parameterized) groups and GGMs as $\mathcal{P}^{\mathsf{CDH}}_{N,m}$ and $\mathcal{G}_{N,m}$.

To establish the black-box separation between $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ and $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$ where m_2 is much larger than m_1, we apply the common technique, namely, the relativizing separation. Concretely, we identify an idealized oracle $\mathcal{O}$ and prove that the longer CDH-secure groups exist relative to $\mathcal{O}$, but the shorter one does not exist. In our strategy, we set this oracle to be the GGM with longer group encodings, namely $\mathcal{G}_{N,m_2}$. At the first glance, this seems impossible, because the GGM is designed as the idealized model for cryptographic groups, and the GGM justifies the CDH by having the unconditional lower bound of the hardness. Fortunately, we observe that the analysis becomes subtle when the "length" is involved.

Theorem 1 (Main Theorem, informal). *Consider $m_1 < m_2$. The shorter CDH-secure groups $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ are black-box separated from the longer CDH-secure groups $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$. Concretely,*

- *$\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ does not exist in the generic group model $\mathcal{G}_{N,m_2}$;*
- *$\mathcal{G}_{N,m_2}$ implies $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$.*

Remark 1. Careful readers might wonder what is the relationship between the longer groups and the shorter groups in which *discrete logarithm* is assumed to be hard. We emphasize that, due to technical challenges[3], the relationship between the longer and shorter groups remains unknown—neither positively nor negatively established. We leave it as an open problem.

Next, we turn our attention to understanding the relationship between the GGMs with different lengths of encoding. Based on the trivial observation that "$\mathcal{G}_{N,m_1}$ implies $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$", we immediately note that the shorter GGM, $\mathcal{G}_{N,m_1}$, and the longer GGM, $\mathcal{G}_{N,m_2}$ do not yield the same black-box complexity. However, when attempting to grasp the relationship between two idealized models, solely relying on black-box complexity might not provide us a comprehensive understanding. Essentially, the black-box complexity of a model only demonstrates the limit of standard-model cryptographic systems it implies and considers the computationally unbounded adversary.

To supplement this, Zhang and Zhandry [ZZ23] propose an orthogonal perspective to the black-box complexity, namely the heuristic complexity. It considers computationally bounded adversaries, thereby excluding the impact of all standard-model cryptosystems on the complexity. We investigate "the length matters" of GGMs within this new perspective, showing that:

[1] Typically N is sufficiently large, and $2^m \geq N$.

[2] By the length of group encoding, we mean that the binary length of the longest canonical representation for all group elements. For instance, let $\mathbb{G}$'s be a group such that the order is 3 and the canonical representation of the group elements is $\{00, 111, 0101\}$, then the length of $\mathbb{G}$, denoted as $\mathsf{len}_{\mathbb{G}}$, is 4.

[3] It is still unclear that whether discrete logarithm implies key agreement or not yet.

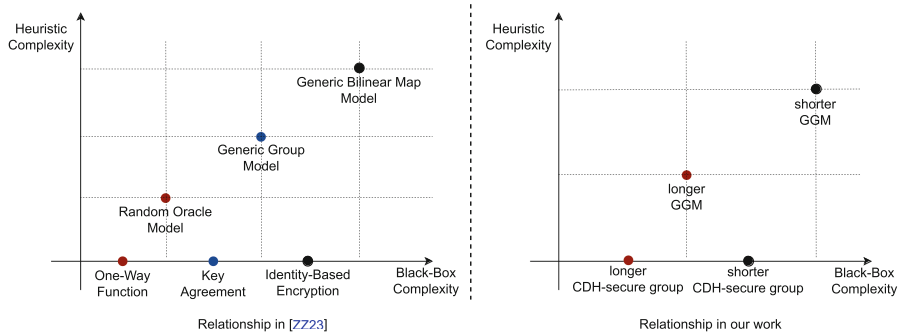

Fig. 1. Relationship between idealized models.

Theorem 2 (Hierarchy of GGMs, Informal). *In the framework of indifferentiability, the GGM with shorter encodings is strictly stronger than the GGM with longer ones, even against computational bounded adversaries.*

To make it clearer, we show our results in Fig. 1. Following the notions in [ZZ23], we give evidence that the shorter GGM statistically implies the longer ones, but the existence of longer GGM's does not computationally imply the existence of a shorter one. More concretely, there exists an indifferentiable construction of a longer generic group with oracle access to shorter generic group without any computational assumption; whereas, as long as the difference in encoding lengths is sufficiently large, there does not exist an indifferentiable construction of a shorter generic group from a long generic group, even with any additional computational assumption.

1.2 Interpretation

Below, we offer interpretations of our findings.

From the Perspective of Black-Box Separation. Our results will bring the research community a better understanding of the cryptographic groups and the generic group models, from the perspective of the black-box reduction/separation[4]. In literature, generic group models have been frequently used to show the impossibility of constructing advanced group-based cryptosystems. Examples include identity-based encryption (IBE) [PRV12,SGS21, Zha22], indistinguishable obfuscation (iO) [MMN16], registration-based encryption (RBE) [HMQS23], accumulators [SGS20], order revealing encryption (ORE) [ZZ18], verifiable delay functions (VDF) [RSS20], and digital signature [DHH+21]. Most of the separation results (e.g., [MMN16,ZZ18,RSS20, SGS20]) are established in Maurer's GGM. Meanwhile, Zhandry [Zha22] illustrates the limits of Maurer's GGM by proving that there are many natural

[4] In this work, when talking about the black-box reduction/separation, we mean that the *fully* black-box reduction/separation that is explicitly defined in [RTV04].

group-based cryptographic schemes (e.g., efficient IND-CPA secure PKE) cannot be modeled by Maurer's GGM, and motivates the line of research, i.e., separations in Shoup's GGM (e.g., IBE in [Zha22] and RBE in [HMQS23]).

Our results demonstrate the *first* evidence that the generic group model can also be used to show the impossibility of constructing plain cryptographic groups, varying in distinct length of group encodings. Speaking of the "lengths" in cryptographic primitives, prior to our work, Garg et.al. [GMM17] prove that there is no iO construction from the single-key functional encryption (FE), if the output length of the functions is much shorter than the length of the ciphertexts[5]. Therefore, we believe that, our result would motivate the community to study the "lengths" in fundamental primitives (e.g., PKE).

However, when delving deeper into our analysis, we stress that our separation results have a limitation. That is, we only establish the separations between the shorter CDH-secure groups and the longer CDH-secure groups under the condition that those groups yield the same security-parameter, which indicates that our separations are somehow security-parameter *dependent*.

For readability, we now explain the limitation through a concrete example. Let λ and λ' be two security parameters. Let p and p' be two primes where $\lfloor \log p \rfloor = \lambda$ and $\lfloor \log p' \rfloor = \lambda'$. Let $\mathbb{G}_1$ be a CDH-secure cryptographic group where the group order is p and the length is 2λ. Let $\mathbb{G}_2$ be another CDH-secure cryptographic group where the group order is p' and the length is $4\lambda'$. Apparently, $\mathbb{G}_1$ is the shorter group and $\mathbb{G}_2$ is the longer one. According to our findings, if $\lambda' \geq \lambda$, then one cannot generically build $\mathbb{G}_1$ from $\mathbb{G}_2$. Unfortunately, if $\lambda' < \lambda$ (say, $\lambda' = \frac{1}{3}\lambda$, indicating that $4\lambda' = \frac{4}{3}\lambda < 2\lambda$), then the relationship between $\mathbb{G}_1$ and $\mathbb{G}_2$ becomes unclear.

In contrast, most known separations are security-parameter *independent*. Take the separation of IBE in Shoup's GGM [Zha22] for instance; according to Zhandry's analysis, we have that for any sufficiently large λ and λ', one cannot generically build an IBE along with security-parameter λ', in Shoup's GGM with security-parameter λ. In order to establish a complete black-box separation (i.e., in the sense of security-parameter independent) between shorter groups and longer groups, novel techniques must be developed to resolve the limitation; we leave this as an important open problem.

Next, we justify that despite of the limitation, our results are interesting and important. First, when it comes to the problem that building a cryptographic group (say, $\mathbb{G}_1$) from another one (say, $\mathbb{G}_2$), it is natural to study the cases that: (1) $\mathbb{G}_1$ and $\mathbb{G}_2$ are with the same group order; (2) the order of $\mathbb{G}_1$ is a factor of $\mathbb{G}_2$[6]. Second, to the best of our knowledge, we are aware of no technique that can be used to generically build $\mathbb{G}_1$ from $\mathbb{G}_2$ if the group orders of $\mathbb{G}_1$ and $\mathbb{G}_2$ are distinct and co-prime. Therefore, we stress that our separations do capture the *natural* settings.

[5] The separation is established under the condition that one-way functions (OWFs) exist and $\mathbf{NP} \not\subseteq \mathbf{coAM}$.

[6] This case indicates that the security parameter of $\mathbb{G}_2$ is bigger than $\mathbb{G}_1$'s, and fortunately our analysis does capture such a case.

Moreover, our results serve as the first attempt to pin down the "lengths" problem for a fundamental primitive (i.e., cryptographic groups), which might open up new research directions (say, the "lengths" problem for other fundamental primitives). Below, for the ease of exposition, when we say the black-box separation between groups, we always mean the one with the same security parameter.

From a Heuristic Perspective. Our results will deepen our understanding of the generic group models from the perspective of heuristic complexity. Inspired by [MRH04, ZZ23], an idealized model can be interpreted through two orthogonal perspectives: the black-box complexity and the heuristic complexity, as depicted in Fig. 1.

For the heuristic aspect, initiated by Maurer et.al. [MRH04] and explicitly studied by Zhang and Zhandry [ZZ23], we consider the framework of indifferentiability against computationally bounded adversaries, where all cryptosystems that exist in the standard model are incorporated. Therefore, the perspective of heuristic is orthogonal to the one of black-box reduction/separation, and understanding the heuristic aspect of various idealized models is important for the relative security of cryptosystems based on idealized models. We establish a strict hierarchy of GGMs from this perspective and prove that the shorter GGM is strictly stronger than the longer one.

In the following, we will give an overview of our approach to comparing the various primitives/models, varying in different lengths of encodings, and our solutions for separating them.

1.3 Technical Overview

Separation Between Cryptographic Groups. Given two cryptographic primitives $\mathcal{P}$ and $\mathcal{Q}$, the typical technique to establish the black-box separation is "relativizing separation" [IR89]. That is, we find a proper oracle $\mathcal{O}$ and prove that the primitive $\mathcal{P}$ exists relative to $\mathcal{O}$ but $\mathcal{Q}$ does not. In our setting, we consider the primitives $\mathcal{P}$ and $\mathcal{Q}$ to be the longer CDH-secure group and shorter one, respectively.

Compared to prior works, the main technical challenge is that, we need to show the gap between two primitives within the same security game (i.e., the CDH game), rather than within different games[7]. The first obstacle is to find a proper oracle. Apparently, the random oracle does not serve our purpose, because the random oracle is weak and there is no construction for CDH-secure groups in the random oracle model.

Our idea is to use a *stronger* oracle, the generic group model. At the first glance, this is impossible, because GGM implies CDH trivially! Fortunately, the GGMs varying in length of group encodings might also yield different levels of complexity, and thus we set this oracle to be the longer GGM within the same security parameter. Concretely, we denote the shorter groups, the longer groups

[7] Games in [IR89] are one-wayness and key recovery attack, respectively.

and the longer GGM as $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}, \mathcal{P}^{\mathsf{CDH}}_{N,m_2}, \mathcal{G}_{N,m_2}$, respectively; recall that $m_2 > m_1$; and we prove that:

- $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$ exists relative to $\mathcal{G}_{N,m_2}$;
- $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ does not exist relative to $\mathcal{G}_{N,m_2}$.

As the former statement is trivial, below we only explain the latter one. To show that $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ does not exist in $\mathcal{G}_{N,m_2}$, it suffices to construct an adversary $\mathcal{A}$ that breaks the CDH game for any construction of shorter group relative to $\mathcal{G}_{N,m_2}$. Due to technical difficulties, we switch to an alternative path. First we pin down a new primitive—non-interactive key exchange (NIKE) with shorter public key, denoted as $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$ [8]. Then we prove that:

1. $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ implies $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$;
2. $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$ does not exist relative to $\mathcal{G}_{N,m_2}$.

As the first statement is straightforward, we will prove the second one. Essentially, $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$, in and of itself, is a key agreement scheme. Next, we give a brief explanation of the separation between NIKE and the random oracle [BKSY11] and then demonstrate how to incorporate the ideas into our analysis. Let $\mathcal{H}$ be a random oracle and $\Pi^{\mathcal{H}} := (\mathsf{KGen}^{\mathcal{H}}, \mathsf{SHK}^{\mathcal{H}})$ be an NIKE scheme with *perfect* correctness. Assuming that the algorithms KGen and SHK make at most q queries, we then construct the adversary $\mathcal{A}$ as follows[9]. Let Alice and Bob be two honest parties. Given the transcript of an execution between Alice and Bob, i.e., pk_A and pk_B, in the present of $\mathcal{H}$, the adversary $\mathcal{A}$ maintains a set $S_{\mathsf{que\text{-}res}}$ of query/response pairs of $\mathcal{H}$, and a multiset S_{key} of candidate keys, both initialized to be $\emptyset$. The adversary $\mathcal{A}$ then runs $4q+1$ iterations of the following attack:

- *Simulation Phase.* The adversary $\mathcal{A}$ searches a proper view of Alice that is consistent with pk_A and $S_{\mathsf{que\text{-}res}}$. Specifically, this view contains the randomness r^*_A used by KGen and SHK, as well as a set of oracle queries/responses $\hat{S}_{\mathcal{A}}$ made by KGen and SHK. The set $\hat{S}_{\mathcal{A}}$ is chosen to be consistent with $S_{\mathsf{que\text{-}res}}$, but it is unnecessary to be consistent with the true oracle $\mathcal{H}$. Let key be the value computed by $\mathsf{SHK}(r^*_A, \mathsf{pk}_B)$. Now, $\mathcal{A}$ adds key into S_{key}.
- *Update Phase.* The adversary $\mathcal{A}$ makes all queries in $\hat{S}_{\mathcal{A}} \setminus S_{\mathsf{que\text{-}res}}$ to the oracle $\mathcal{H}$, and adds the corresponding pairs into $S_{\mathsf{que\text{-}res}}$.

Finally, $\mathcal{A}$ outputs the majority of the keys in S_{key}. Next, we explain why $\mathcal{A}$ recovers the key. Let S_B denote the queries made by Bob in the real execution of the key exchange protocol. In a given iteration, there are two events:

- Event 1 (**Bad event**): $\exists(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}, (\mathsf{que}_B, \mathsf{res}_B) \in S_B$ s.t. $\mathsf{que}_A = \mathsf{que}_B$ but $\mathsf{res}_A \neq \mathsf{res}_B$.

[8] Here, m_1 means the length of the public key; please find the formal definition in Sect. 2.1.

[9] The adversary here is computational unbounded but query-efficient.

– Event 2 (**Good event**): $\forall(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}, (\mathsf{que}_B, \mathsf{res}_B) \in S_B$, we have that if $\mathsf{que}_A = \mathsf{que}_B$, then $\mathsf{res}_A = \mathsf{res}_B$.

Note that event 1 only occurs in at most $2q$ iterations because $|S_B| \leq 2q$ and once it happens, the update phase would absorb at least one pair $(\mathsf{que}_B, \mathsf{res}_B) \in S_B$ into $S_{\mathsf{que\text{-}res}}$. For event 2, we observe that, when it occurs, there is another oracle $\tilde{\mathcal{H}}$ that is consistent with both $\hat{S}_{\mathcal{A}}$ and S_B. Based on the perfect correctness, we have that the shared key computed in that iteration is valid. Moreover, event 2 occurs in at least 2q+1 iterations, indicating that the majority in S_{key} is valid.

However, when it comes to the GGM, the attack fails. Comparing to ROM, there are two kinds of queries in GGM, namely the labeling query $(x, \mathcal{G}^{\mathsf{label}}_{N,m_2}(x))$ and the addition query $(\mathcal{G}^{\mathsf{label}}_{N,m_2}(x), \mathcal{G}^{\mathsf{label}}_{N,m_2}(y), \mathcal{G}^{\mathsf{label}}_{N,m_2}(x+y))$. Therefore, we should define S_B that covers all the group encodings that appear in the queries (both labeling and addition) with the discrete logarithms (Bob might not know the value). Then, in a given iteration, there are three events:

– Event 1 (**Bad event**): $\exists(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}, (\mathsf{que}_B, \mathsf{res}_B) \in S_B$ s.t. $\mathsf{que}_A = \mathsf{que}_B$ but $\mathsf{res}_A \neq \mathsf{res}_B$.
– Event 2 (**Bad event**): $\exists(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}, (\mathsf{que}_B, \mathsf{res}_B) \in S_B$ s.t. $\mathsf{que}_A \neq \mathsf{que}_B$ but $\mathsf{res}_A = \mathsf{res}_B$.
– Event 3 (**Good event**): $\forall(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}, (\mathsf{que}_B, \mathsf{res}_B) \in S_B$, we have that if $\mathsf{que}_A = \mathsf{que}_B$, then $\mathsf{res}_A = \mathsf{res}_B$.

Note that event 1 and event 3 can be handled similarly as above. However, the fatal problem is that event 2 might always happen. In other words, we cannot find a GGM that is consistent with both $\hat{S}_{\mathcal{A}}$ and S_B, indicating that the above attack fails immediately.

More specifically, we note that the reason why event 2 happens is that, given pk_A and pk_B, algorithms can obtain valid group encoding without making labeling query[10]. Moreover, if algorithms *cannot* obtain valid group encodings without making labeling queries, then the GGM can be simulated by a stateful oracle that only provides labeling queries, as the addition queries can be easily converted into labeling queries. Such an oracle is close to the random oracle model and thus our goal is to design a mechanism that prevent extracting valid group elements without knowing the corresponding discrete logarithms.

Here we introduce our *length tool*, intuitively, if the length of the public key is much shorter than the group encoding (say, the length gap is at least $\omega(\log \lambda)$), then the public key would not carry enough information to recover the group encodings. This also explains why we choose NIKE other than general key agreement (say, multi-round KA), because the adversary only obtains two public keys in the setting of NIKE.

[10] This in fact is natural in group-based cryptosystem, take the ElGamal encryption scheme [ElG85] for instance, the public key itself is a valid group element.

Concretely, let Q_{sk_A} and Q_{sk_B} be the set of the query/response pairs (only labeling queries[11]) made when running $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}_A)$ and $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}_B)$, respectively. Let h be the valid group encoding that an algorithm outputs, by having pk_A and pk_B, we then consider the following four cases:

- Case 1: (**Independent**) $h \notin Q_{\mathsf{sk}_A} \cup Q_{\mathsf{sk}_B}$.
- Case 2: (**Frequent**) $h \in Q_{\mathsf{sk}_A} \cap Q_{\mathsf{sk}_B}$.
- Case 3: (**Dependent but hard to extract**) $h \in Q_{\mathsf{sk}_A} \setminus Q_{\mathsf{sk}_B}$.
- Case 4: (**Dependent but hard to extract**) $h \in Q_{\mathsf{sk}_B} \setminus Q_{\mathsf{sk}_A}$.

For case 1, h is independent of pk_A and pk_B. Due to the sparseness of the group encodings in $\mathcal{G}_{N,m_2}$, no algorithm can output h except for negligible probability.

For case 2, it is apparent that pk_A and pk_B together might carry enough information for h. Fortunately, with high probability h is a frequent query, therefore the discrete logarithm of h can be easily obtained by repeatedly running $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\cdot)$ on sufficiently many random inputs.

For case 3 (or case 4), note that h is independent of pk_B, which means that only pk_A carries the information of h. Note that the length of pk_A is m_1 but length of h is m_2. Moreover, h is uniformly distributed over the probability of GGM. Therefore, conditioned on that $m_2 - m_1$ is sufficiently large, no algorithm can extract such an h except for negligible probability.

The above sketch is not precise; please find low-level details, in Sect. 3.

Hierarchy of GGMs. To establish the hierarchy of the generic group models against computational bounded adversaries, we formalize our goal in the framework of indifferentiability. Specifically, we prove that the shorter GGM (denoted as $\mathcal{G}_{N,m_1}$) statistically implies the longer one (denoted as $\mathcal{G}_{N,m_2}$), but the longer GGM does not computationally imply the shorter one.

$\mathcal{G}_{N,m_1}$ statistically implies $\mathcal{G}_{N,m_2}$. We first explain how $\mathcal{G}_{N,m_1}$ implies $\mathcal{G}_{N,m_2}$ against computationally unbounded adversaries. Let $\mathcal{H}$ be a random oracle that maps $\{0,1\}^* \to \{0,1\}^{m_2-m_1}$; as the first attempt, it is natural to design the labeling function as:

$$L^{\mathcal{G}_{N,m_1},\mathcal{H}}(x) := \mathcal{G}^{\mathsf{label}}_{N,m_1}(x) || \mathcal{H}(\mathcal{G}^{\mathsf{label}}_{N,m_1}(x)).$$

However, there always exists an efficient distinguisher that breaks the indifferentiability w.r.t. the above scheme. Specifically, in the ideal world, the distinguisher uniformly samples $x \in \mathbb{Z}_N$, makes a labeling query with x, and obtains $\mathcal{G}^{\mathsf{label}}_{N,m_2}(x)$. Let str and str' be the first m_1 bits and the last $m_2 - m_1$ bits of $\mathcal{G}^{\mathsf{label}}_{N,m_2}(x)$, respectively. Then the distinguisher makes a query to the simulator with input str and checks whether the response matches str'. Note that, without knowing x, the

[11] We stress that, for the algorithm $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}$, without loss of generality, it only makes labeling queries. Essentially, the group encodings of $\mathcal{G}_{N,m_2}$ are sparse, which means any algorithm with inputs that are independent of $\mathcal{G}_{N,m_2}$ cannot obtain a valid group encoding without making labeling query, indicating that any addition query can be absorbed by the corresponding labeling query.

simulator cannot answer this query properly except for a negligible probability. To prevent the attack above, we enhance the power of the simulator. We involve an additional oracle, the random permutation oracle $\mathcal{E}$, that permutes $\{0,1\}^{m_2}$ with its inverse[12] $\mathcal{E}^{-1}$, and design the labeling function as:

$$L^{\mathcal{G}_{N,m_1},\mathcal{H},\mathcal{E}}(x) := \mathcal{E}(\mathcal{G}^{\mathsf{label}}_{N,m_1}(x)||\mathcal{H}(\mathcal{G}^{\mathsf{label}}_{N,m_1}(x))).$$

Careful readers may wonder why it works. Note that both $\mathcal{E}$ and $\mathcal{E}^{-1}$ are under full control of the simulator, which means that the distinguisher is independent of the value $\mathcal{H}(\mathcal{G}^{\mathsf{label}}_{N,m_1}(x))$ without making queries to the simulator. This extra information gained from these queries is exactly what the simulator requires for the proof to go through. In fact, with the aid of $\mathcal{E}$, we can even simplify the construction by replacing $\mathcal{H}(\mathcal{G}^{\mathsf{label}}_{N,m_1}(x))$ with a fixed string, say $0\cdots0$, concretely:

$$L^{\mathcal{G}_{N,m_1},\mathcal{E}}(x) := \mathcal{E}(\mathcal{G}^{\mathsf{label}}_{N,m_1}(x)||\underbrace{0\cdots0}_{m_2-m_1}).$$

The addition algorithm can be easily constructed by applying the inverse oracle $\mathcal{E}^{-1}$. While the additional oracle $\mathcal{E}$ and its inverse $\mathcal{E}^{-1}$ have protected against certain natural attacks, we need to argue indifferentiability against all possible attacks. To do so, we use a careful simulation strategy for $\mathcal{G}^{\mathsf{label}}_{N,m_1}, \mathcal{G}^{\mathsf{add}}_{N,m_1}, \mathcal{E}$, and $\mathcal{E}^{-1}$, and prove indifferentiability through a careful sequence of hybrids. Due to the space limit, we omit the formal descriptions of our simulation, and we refer interesting readers to see it in the full version of this paper [ZJW+24].

Remark 2. Careful readers might note that the building blocks of construction above contain both the shorter GGM $\mathcal{G}_{N,m_1}$ and an *additional* independent randome oracle, rather than the shorter GGM solely. Although we have that GGM implies ROM statistically [ZZ23], it is unclear to us that how to build an indifferentiable GGM plus an independent ROM from a single GGM. Therefore, we stress that our hierarchy of the GGM is established with the aid of an additional independent random oracle.

Moreover, this even motivates an interesting research question that whether one single GGM implies multiple independent GGMs, comparing to the fact that the random oracle does.

$\mathcal{G}_{N,m_2}$ does not computationally imply $\mathcal{G}_{N,m_1}$. Suppose we have a purported construction $\Pi^{\mathcal{G}_{N,m_2}} := (L^{\mathcal{G}_{N,m_2}}, A^{\mathcal{G}_{N,m_2}})$ of a shorter group from a longer GGM. How could we prove that $\Pi^{\mathcal{G}_{N,m_2}}$ can be differentiated from $\mathcal{G}_{N,m_1}$ by a computationally bounded distinguisher?

Following the strategy in [ZZ23], we should find some security property P that holds for $\mathcal{G}_{N,m_1}$ but fails for *any* $\Pi^{\mathcal{G}_{N,m_2}}$. As explained in [ZZ23], any standard model assumption cannot serve as the property, and thus, this property P is set to be a variant of discrete logarithm problem, called *discrete log identification*

[12] According to [HKT11], the random oracle and random permutation oracle with inverse are equivalent, therefore we take $\mathcal{E}$ and $\mathcal{E}^{-1}$ for granted.

(DLI). Intuitively, DLI is defined as: given $h := L(x)$, construct a (probabilistic, efficient, and query-free) circuit C such that $C(x)$ accepts with a high probability, but $C(x')$ rejects with a overwhelming probability on all $x' \neq x$. Apparently, the DLI problem is easy on any standard-model group: for any y, set $C(y)$ to be 1 if and only if $L(y) = h$, where $L(y) := g^y$ is computed as part of the circuit[13]. To establish the separation between GGM and ROM, Zhandry and Zhang prove that the DLI problem is also easy on any group built within the random oracle model. Intuitively, they "compile out" the random oracle $\mathcal{H}$ and design an attacker that can easily construct an oracle-aided circuit $C^{\mathcal{H}}(\cdot)$, breaking the DLI problem by computing $L^{\mathcal{H}}(\cdot)$. The subtlety is to anticipate the oracle queries that C will make to the random oracle model and have the attacker make the corresponding queries for itself. Concretely, given input $L^{\mathcal{H}}(x)$, the attacker runs the addition algorithm $A^{\mathcal{H}}(L^{\mathcal{H}}(y), L^{\mathcal{H}}(x-y))$ and $L^{\mathcal{H}}(\cdot)$ on several random inputs, records all queries/responses that were made, and hardcodes the queries/responses into the C to obtain an oracle-free circuit, which C outputs.

Below, we outline our method for integrating the aforementioned technique into the analysis within the longer GGM. The difficulty is that, our goal seems to conflict with the results in [ZZ23], as they have proven that the DLI problem is hard with respect to the generic group model. To bypass the obstacle, we here leverage the *length tool* again.

Consider computing $L^{\mathcal{G}_{N,m_2}}(x)$ from x, which in turn makes queries to the longer GGM, $\mathcal{G}_{N,m_2}$. Let Q_x be the set of query/response pairs made during the procedure of computing $L^{\mathcal{G}_{N,m_2}}(x)$. Similarly as above, we assume that, without loss of generality, each query/response pair $(\mathsf{que}, \mathsf{res}) \in Q_x$ is a labeling query. Consider computing $A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(y), L^{\mathcal{G}_{N,m_2}}(z))$ where y and z are random, conditioned on $y + z = x$. The output of this addition is $L^{\mathcal{G}_{N,m_2}}(y+z) = L^{\mathcal{G}_{N,m_2}}(x)$. For each query/response pair $(\mathsf{que}, \mathsf{res}) \in Q_x$, there are roughly four possible cases:

- Case 1: The label $L^{\mathcal{G}_{N,m_2}}(x)$ does *not* depend on the response res at all;
- Case 2: The label $L^{\mathcal{G}_{N,m_2}}(x)$ depends on the response res, but with a overwhelming probability over the choice of y and z, res does not appear when computing $A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(y), L^{\mathcal{G}_{N,m_2}}(z))$;
- Case 3: The label $L^{\mathcal{G}_{N,m_2}}(x)$ depends on the response res, and with a non-negligible probability over the choice of y and z, $A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(y), L^{\mathcal{G}_{N,m_2}}(z))$ makes a "labeling" query to $\mathcal{G}_{N,m_2}$ on input que;
- Case 4: The label $L^{\mathcal{G}_{N,m_2}}(x)$ depends on the response res, and with a non-negligible probability over the choice of y and z, an "addition" query $(\mathsf{que}_1, \mathsf{que}_2, \mathsf{res})$ occurs when computing $A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(y), L^{\mathcal{G}_{N,m_2}}(z))$.

Now we explain our approach of building the oracle-free circuit C. We collect queries into a list, denoted as $S_{\mathsf{que\text{-}res}}$, and hardcode $S_{\mathsf{que\text{-}res}}$ into the circuit C to

[13] Note that for standard-model groups, $L(y)$ denotes the value g^y for the fixed generator g, and here y is the discrete logarithm of h with respect to g.

make sure that $C(x)$ will be able to reconstruct $L^{\mathcal{G}_{N,m_2}}(x)$ without making any query to the oracle at all:

- In case 1 (**Non-sensitive query**), same as in [ZZ23], since $L^{\mathcal{G}_{N,m_2}}(x)$ does not depend on res, when computing $L^{\mathcal{G}_{N,m_2}}(x)$ we can just replace the response with a random string without affecting the ultimate labeling. Therefore, for any query not in $S_{\mathsf{que\text{-}res}}$, we will have C respond with a uniformly random string.
- In case 2 (**Sensitive but frequent query**), same as in [ZZ23], since $L^{\mathcal{G}_{N,m_2}}(x)$ does depend on res, this query is a sensitive query for the ultimate labeling. In this case, it must be that $A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(y), L^{\mathcal{G}_{N,m_2}}(z))$ be able to extract res from the inputs, i.e., $L^{\mathcal{G}_{N,m_2}}(y)$ and $L^{\mathcal{G}_{N,m_2}}(z)$, which indicates that, with a high probability, $(\mathsf{que}, \mathsf{res}) \in Q_x \cap (Q_y \cup Q_z)$. On the other hand, x, y and z are pairwise independent, which means that $Q_x \cap Q_y$ and $Q_x \cap Q_z$ *only* contains "frequent" queries. Therefore, this query/response pair can be collected by running $L^{\mathcal{G}_{N,m_2}}(\cdot)$ on sufficiently many random inputs.
- In case 3 (**Sensitive labeling query**), same as in [ZZ23], $S_{\mathsf{que\text{-}res}}$ collects all the labeling queries that occur when running $A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(y), L^{\mathcal{G}_{N,m_2}}(z))$. We know that, with a non-negligible probability $(\mathsf{que}, \mathsf{res})$ will be amongst the queries in $S_{\mathsf{que\text{-}res}}$. By repeating several times, we have that $(\mathsf{que}, \mathsf{res}) \in S_{\mathsf{que\text{-}res}}$ with a high probability.
- In Case 4 (**Sensitive addition query**), different from [ZZ23], the addition query, i.e., $(\mathsf{que}_1, \mathsf{que}_2, \mathsf{res})$ occurs, where que_1 and que_2 are two valid group encodings of $\mathcal{G}_{N,m_2}$. Although res appears in this query, collecting this kind of query is not usually useful for our purpose. Specifically, when running $L^{\mathcal{G}_{N,m_2}}(x)$, the algorithm might make labeling queries on points $(x_1, \ldots, x_q)$, whereas $S_{\mathsf{que\text{-}res}}$ might only store query/response pairs in the form of addition, i.e., $(\mathcal{G}_{N,m_2}(y_i), \mathcal{G}_{N,m_2}(z_i), \mathcal{G}_{N,m_2}(y_i + z_i))$, without explicitly knowing either y_i or z_i. As a result, $C(x)$ might fail to reconstruct $L^{\mathcal{G}_{N,m_2}}(x)$: when running $C(x) := L^{S_{\mathsf{que\text{-}res}}}(x)$, although C knows that $\mathcal{G}_{N,m_2}(x_i)$ exists in the database $S_{\mathsf{que\text{-}res}}$, it does *not* know which tuple corresponds to the correct one.

  To resolve the problem, we need to transform this addition query into a labeling query. Observe that if the discrete logarithms of que_1 and que_2 are known, then the transformation is trivial. Exploring deeper, during the procedure of computing $A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(y), L^{\mathcal{G}_{N,m_2}}(z))$, the algorithm $A^{\mathcal{G}_{N,m_2}}$ can only extract valid group encodings in $Q_y \cup Q_z$[14]. Moreover, we have that $L^{\mathcal{G}_{N,m_2}}(y)$ and $L^{\mathcal{G}_{N,m_2}}(z)$ are independent of the responses that are in $Q_z \setminus Q_y$ and $Q_y \setminus Q_z$, respectively.

  Now, we *leverage the length tool*. Concretely, from $A^{\mathcal{G}_{N,m_2}}$'s perspective, $L^{\mathcal{G}_{N,m_2}}(y)$ is the only string that carries information of the valid group encodings $\in Q_y \setminus Q_z$. If $m_2 - m_1$ is sufficiently large, say $m_2 - m_1 \geq \omega(\log \lambda)$, where λ is the security parameter, then it is impossible for $A^{\mathcal{G}_{N,m_2}}$ to extract a valid group encoding from $Q_y \setminus Q_z$ except for a negligible probability, indicating that the valid group encodings that $A^{\mathcal{G}_{N,m_2}}$ can extract are in $Q_y \cap Q_z$. Having

[14] Other group encodings in $\mathcal{G}_{N,m_2}$ are independent of $L^{\mathcal{G}_{N,m_2}}(y)$ and $L^{\mathcal{G}_{N,m_2}}(z)$.

that x, y and z are pairwise independent, we know that queries in $Q_y \cap Q_z$ are frequent with a high probability, which can be easily captured as in case 2.

Next, we consider the value of $C(x')$ for $x' \neq x$. If we are lucky and $S_{\mathsf{que\text{-}res}}$ contains all sensitive queries of $Q_{x'}$, then $C(x') = L^{\mathcal{G}_{N,m_2}}(x') \neq L^{\mathcal{G}_{N,m_2}}(x)$, indicating that $C(x')$ rejects as desired. Otherwise, if $S_{\mathsf{que\text{-}res}}$ does not contain all the sensitive queries of $Q_{x'}$, then $S_{\mathsf{que\text{-}res}}$ would respond to the query with random value, which means that $C(x')$ computes an invalid label for x'. As explained in [ZZ23], the random response would only serve to inject further randomness into the label, and the invalid label would be unequal to $L^{\mathcal{G}_{N,m_2}}(x)$ with a high probability. Combining the above together, we build an oracle-free circuit that *only* accepts the discrete logarithm x.

The above sketch is not precise; please find the low-level details in Sect. 4.

The Hierarchy is Tight. To complement our results of the hierarchy, we next show that if $m_2 - m_1$ is *small*, then $\mathcal{G}_{N,m_1}$ and $\mathcal{G}_{N,m_2}$ are equivalent under the indifferentiability framework. To explain our idea, we illustrate the simplest case, where $m_2 - m_1 = 1$[15]. Let Trunc be the function that chops off the last bit of the input, we build an indifferentiable group in $\mathcal{G}_{N,m_2}$ as follows:

$$L^{\mathcal{G}_{N,m_2}}(x) := \mathsf{Trunc}(\mathcal{G}^{\mathsf{label}}_{N,m_2}(x));$$

$$A^{\mathcal{G}_{N,m_2}}(\mathsf{str}_0, \mathsf{str}_1) := \begin{cases} \mathsf{Trunc}(\mathcal{G}^{\mathsf{add}}_{N,m_2}(\mathsf{str}_0||b_0, \mathsf{str}_1||b_1)), & \text{if } \mathsf{str}_0||b_0 \text{ and } \mathsf{str}_1||b_1 \text{ are valid;} \\ \perp & \text{otherwise.} \end{cases}$$

For clarity, if there exist $b_0, b_1 \in \{0,1\}$ such that both $\mathsf{str}_0||b_0$ and $\mathsf{str}_1||b_1$ are valid, then the addition algorithm outputs $\mathsf{Trunc}(\mathcal{G}^{\mathsf{add}}_{N,m_2}(\mathsf{str}_0||b_0, \mathsf{str}_1||b_1))$, otherwise it aborts. Based on the fact that the group encodings of $\mathcal{G}_{N,m_2}$ are sparse, we know that for any string str, the probability that both $\mathsf{str}||0$ and $\mathsf{str}||1$ are valid is negligible, which indicates that the addition algorithm is well defined. Moreover, we prove that the construction above is indifferentiable from $\mathcal{G}_{N,m_1}$. Due to the space limit, we leave the proof in the full version of this paper [ZJW+24].

Due to the composition of indifferentiability, our results can be easily extended to the case that $m_2 - m_1 \leq \Theta(\log \lambda)$, which completes the entire picture of the hierarchy asymptotically.

1.4 Organization

In Sect. 2, we present the necessary notations, concepts, and definitions. We establish a separation between two CDH-secure groups with sufficiently large

[15] We also require that group encodings in $\mathcal{G}_{N,m_2}$ are sparse, say $m_2 - \log N \geq \omega(\log \lambda)$.

encoding length difference in Sect. 3. We then establish a hierarchy among GGMs with different encoding lengths in Sect. 4. All formal proofs can be found in the full version of this paper [ZJW+24] due to the space limitation.

2 Preliminaries

Notation. For a finite set S, we denote a random sample s from S according to the uniform distribution as $s \xleftarrow{\$} S$. We say a positive function $\mathsf{negl}(\cdot)$ is negligible, if for all positive polynomial $p(\cdot)$, there exists a constant $\lambda_0 > 0$ such that for all $\lambda > \lambda_0$, it holds that $\mathsf{negl}(\lambda) < 1/p(\lambda)$. We say a function $\rho(\cdot)$ is noticeable in λ, if the inverse $1/\rho(\lambda)$ is polynomial in λ. We write $y \xleftarrow{\$} \mathsf{Alg}(I)$ to denote variable y that is obtained by running a randomized algorithm Alg on input I (which may consist of a tuple $I := (I_1, ..., I_n)$). If Alg is deterministic, we write "$\leftarrow$" instead of "$\xleftarrow{\$}$". By $x||y$, we mean the concatenation of strings x and y.

Algorithms. Denote $\lambda \in \mathbb{N}$ as the security parameter. Here we use a non-uniform circuit to formalize the model of computation. An algorithm Alg is a collection of circuits $\{C_\lambda\}_{\lambda\in\mathbb{N}}$ with domain Dom_λ and range Ran_λ, respectively. When considering interactive algorithms $(\mathsf{Alg}_1, \ldots, \mathsf{Alg}_n)$, algorithms are treated as a sequence of circuits $C_\lambda^{(1)}, C_\lambda^{(2)}, \ldots$, where the domain of $C_\lambda^{(i)}$ is denoted as $\mathsf{Dom}_\lambda^{(i)} = \mathsf{stat}_\lambda^{(i)} \times \mathsf{input}_\lambda^{(i-1)}$, the range of $C_\lambda^{(i)}$ is denoted as $\mathsf{Ran}_\lambda^{(i)} = \mathsf{stat}_\lambda^{(i+1)} \times \mathsf{output}_\lambda^{(i)}$. Here, $\mathsf{stat}_\lambda^{(i)}$ ($\mathsf{input}_\lambda^{(i)}$, $\mathsf{output}_\lambda^{(i)}$) is the space of the state (inputs, outputs) that $C_\lambda^{(i)}$ sends to $C_\lambda^{(i+1)}$, respectively.

Games. A game is initiated by a probabilistic interactive algorithm C, called a challenger, and a predicate function $\mathsf{pf} : \{0,1\}^* \to [0,1]$. The challenger takes the security parameter as input and interacts with k communicating-restricted parties $(\mathsf{Alg}_1, \ldots, \mathsf{Alg}_k)$. We call $\mathcal{A} := (\mathsf{Alg}_1, \ldots, \mathsf{Alg}_k)$ the adversary. In the end of the game, the challenger C outputs a bit b; if $b = 1$ we say the adversary wins the game, otherwise we say the adversary loses. Let $\mathrm{Cl}(\mathcal{A})$ be a class of adversary. We say a game $(\mathsf{C}, \mathsf{pf})$ is hard with respect to $\mathrm{Cl}(\mathcal{A})$, if for any adversary $\mathcal{A} \in \mathrm{Cl}(\mathcal{A})$, we have $\Pr[\mathcal{A} \text{ wins}] \leq \mathsf{pf} + \mathsf{negl}(\lambda)$.

Cryptosystems. A cryptosystem Σ consists of a set of algorithms, which typically are non-interactive. Here, Σ is accessible via two interfaces Σ.hon and Σ.adv, where Σ.hon provides an honest interface through which the system can be accessed by all parties in a black-box manner, and Σ.adv models the adversarial access to the inner working part of Σ.

2.1 Primitives, Idealized Models, and Reduction Notions

In this work, we treat CDH-secure groups as cryptographic primitives, and explore black-box reduction between them with different lengths. First of all, we recall the definition of primitive formalized by [RTV04].

2.1.1 Cryptographic Primitives

Definition 1 (Cryptographic Primitive [RTV04]**).** *A primitive $\mathcal{P}$ is a pair $\langle \mathcal{F}_{\mathcal{P}}, \mathcal{R}_{\mathcal{P}} \rangle$, where $\mathcal{F}_{\mathcal{P}}$ is a set of functions $f : \{0,1\}^* \mapsto \{0,1\}^*$, and $\mathcal{R}_{\mathcal{P}}$ is a relation over pairs $\langle f, \mathcal{A} \rangle$ of a function $f \in \mathcal{F}_{\mathcal{P}}$ and an adversarial machine $\mathcal{A}$. (The set $\mathcal{F}_{\mathcal{P}}$ is required to contain at least one function which is computable by a* PPT *machine.)*

- Efficient implementation. *We say a function f* **implements** *$\mathcal{P}$ or is an implementation of $\mathcal{P}$ if $f \in \mathcal{F}_{\mathcal{P}}$. An* **efficient implementation** *of $\mathcal{P}$ is an implementation of $\mathcal{P}$ which is polynomial-time computable.*
- Secure implementation. *We say an adversarial machine $\mathcal{A}$* **$\mathcal{P}$-breaks** *$f \in \mathcal{F}_{\mathcal{P}}$ if $\langle f, \mathcal{A} \rangle \in \mathcal{R}_{\mathcal{P}}$. A* **secure implementation** *of $\mathcal{P}$ is an implementation of $\mathcal{P}$ such that no* PPT *adversarial machine $\mathcal{P}$-breaks f.*

We say the **primitive $\mathcal{P}$ exists** *if there is an efficient and secure implementation of $\mathcal{P}$.*

As mentioned before, we treat CDH-secure groups as a cryptographic primitive. Now we formalize this primitive by using the terms in [RTV04].

Definition 2 (CDH-Secure Groups). *A CDH-secure group $\mathcal{P}^{\mathsf{CDH}}$ consists of the following pair $\langle \mathcal{F}_{\mathcal{P}^{\mathsf{CDH}}}, \mathcal{R}_{\mathcal{P}^{\mathsf{CDH}}} \rangle$:*

1. *The set $\mathcal{F}_{\mathcal{P}^{\mathsf{CDH}}}$ for specifying syntax and capturing the correctness property. Here, the set $\mathcal{F}_{\mathcal{P}^{\mathsf{CDH}}}$ consists of functions f, where f represents the* group generation function *for generating group description of finite cycle groups. Concretely, we write $(\mathbb{G}, g, N, m) \xleftarrow{\$} f(1^\lambda)$, where $\mathbb{G}$ is a cyclic group of prime order N, g is a generator $\mathbb{G}$, and m is the length of group encoding (that is, each group element in $\mathbb{G}$ can be represented as an m-bit string).*
   We note that the correctness is guaranteed by the basic properties of the cyclic group.
2. *The relation $\mathcal{R}_{\mathcal{P}^{\mathsf{CDH}}}$ for capturing the security property.*
   For function $f \in \mathcal{F}_{\mathcal{P}^{\mathsf{CDH}}}$ and PPT *(adversarial) machine $\mathcal{A}$, we define $\langle f, \mathcal{A} \rangle \in \mathcal{R}_{\mathcal{P}^{\mathsf{CDH}}}$ if there exists a polynomial $p(\cdot)$ such that $\Pr[\mathcal{A}(\mathbb{G}, g, N, m, h_1, h_2) = g^{x_1 x_2}] > 1/p(\lambda)$ for infinitely many λ.*
   Here, $(\mathbb{G}, g, N, m) \xleftarrow{\$} f(1^\lambda)$, and $h_1, h_2 \in \mathbb{G}$ are two uniformly chosen group elements where $h_1 = g^{x_1}$, $h_2 = g^{x_2}$, and $x_1, x_2 \in \mathbb{Z}_N$.

We say CDH-secure group $\mathcal{P}^{\mathsf{CDH}}$ exists, if there exists a function $f \in \mathcal{F}_{\mathcal{P}^{\mathsf{CDH}}}$, it holds that no PPT *adversarial machine $\mathcal{A}$ such that $\langle f, \mathcal{A} \rangle \in \mathcal{R}_{\mathcal{P}^{\mathsf{CDH}}}$. Often, we make the parameters, the order N and the encoding length m, explicit, and denote the CDH-secure group as $\mathcal{P}^{\mathsf{CDH}}_{N,m}$.*

Non-interactive key exchange (NIKE) was initially studied by Diffie and Hellman in their breakthrough paper [DH76]. We now describe this primitive by using the terms in [RTV04].

Definition 3 (Non-Interactive Key Exchange). *A non-interactive key exchange protocol* $\mathcal{P}^{\mathsf{NIKE}}$ *consists of the following pair* $\langle \mathcal{F}_{\mathcal{P}^{\mathsf{NIKE}}}, \mathcal{R}_{\mathcal{P}^{\mathsf{NIKE}}} \rangle$*:*

1. *The set* $\mathcal{F}_{\mathcal{P}^{\mathsf{NIKE}}}$ *for specifying syntax and capturing the correctness property. Here, the set* $\mathcal{F}_{\mathcal{P}^{\mathsf{NIKE}}}$ *consists of functions* f*, where* $f := (\mathsf{KGen}, \mathsf{SHK})$ *represents*
   - *the* public-key message function $\mathsf{KGen} : \mathcal{SK} \mapsto \mathcal{PK}$ *for generating the public-key message based on a randomly chosen private-key, where* $\mathcal{PK}$ *and* $\mathcal{SK}$ *are public-key space and private-key space, respectively.*
   - *the* shared key generation function $\mathsf{SHK} : \mathcal{PK} \times \mathcal{SK} \mapsto \mathcal{K} \cup \{\perp\}$ *for generating the shared key, where* $\mathcal{K}$ *is shared-key space, and* $\perp$ *denotes that the computation fails.*

   Concretely, for randomly chosen $\mathsf{sk} \xleftarrow{\$} \mathcal{SK}$*, we write* $\mathsf{pk} \leftarrow \mathsf{KGen}(\mathsf{sk})$*, where* pk *is called public key. Furthermore, for randomly chosen* $\mathsf{sk}' \xleftarrow{\$} \mathcal{SK}$*, compute* $\mathsf{pk}' \leftarrow \mathsf{KGen}(\mathsf{sk}')$*. We write* $\mathsf{shk} \leftarrow \mathsf{SHK}(\mathsf{pk}', \mathsf{sk})$ *and* $\mathsf{shk}' \leftarrow \mathsf{SHK}(\mathsf{pk}, \mathsf{sk}')$*. Note that, when the shared key generation function fails, we write* $\mathsf{shk} = \perp$ *or* $\mathsf{shk}' = \perp$*.*
   We say correctness *is achieved if there exists an* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\mathsf{shk} \neq \perp \wedge \mathsf{shk}' \neq \perp \wedge \mathsf{shk} \neq \mathsf{shk}'\right] \leq \mathsf{negl}(\lambda)$$

   When $\mathsf{negl}(\lambda) = 0$*, then we say* perfect correctness *is achieved.*
2. *The relation* $\mathcal{R}_{\mathcal{P}^{\mathsf{NIKE}}}$ *for capturing the security property against key-recovery attack (KRA).*
   For function $f := (\mathsf{KGen}, \mathsf{SHK}) \in \mathcal{F}_{\mathcal{P}^{\mathsf{NIKE}}}$ *and a* PPT *(adversarial) machine* $\mathcal{A}$*, we define* $\langle f, \mathcal{A} \rangle \in \mathcal{R}_{\mathcal{P}^{\mathsf{NIKE}}}$ *if there exists a polynomial* $p(\cdot)$ *such that* $\Pr[\mathcal{A}(\mathsf{pk}, \mathsf{pk}') = \mathsf{SHK}(\mathsf{pk}', \mathsf{sk}) = \mathsf{SHK}(\mathsf{pk}, \mathsf{sk}') \neq \perp] > 1/p(\lambda)$ *for infinitely many* λ*.*
   Here, for randomly chosen $\mathsf{sk} \xleftarrow{\$} \mathcal{SK}$ *and* $\mathsf{sk}' \xleftarrow{\$} \mathcal{SK}$*, compute* $\mathsf{pk} \leftarrow \mathsf{KGen}(\mathsf{sk})$ *and* $\mathsf{pk}' \leftarrow \mathsf{KGen}(\mathsf{sk}')$*, respectively.*

We say non-interactive key exchange protocol $\mathcal{P}^{\mathsf{NIKE}}$ *exists, if there exists a function* $f \in \mathcal{F}_{\mathcal{P}^{\mathsf{NIKE}}}$*, it holds that no* PPT *adversarial machine* $\mathcal{A}$ *such that* $\langle f, \mathcal{A} \rangle \in \mathcal{R}_{\mathcal{P}^{\mathsf{NIKE}}}$*. When* $\mathcal{SK} = \mathbb{Z}_N$ *and* $\mathcal{PK} = \mathcal{K} = \{0,1\}^m$*, we make the parameters* N *and* m *explicit and denote the non-interactive key exchange protocol as* $\mathcal{P}^{\mathsf{NIKE}}_{N,m}$*.*

2.1.2 Idealized Models In this subsection, we introduce idealized models including the Random Oracle Model (ROM) [BR93], the Random Permutation Model (RPM) [RS08], and the Generic Group Model (GGM) [Sho97]. In each idealized model, all entities including the adversary $\mathcal{A}$ and the challenger $\mathcal{C}$, are provided with the access to the corresponding oracle. Below we will specify the behavior of the oracle in each idealized model.

Definition 4 (Random Oracle Model [BR93]). *Let* $\mathcal{I}_{*,S}$ *denote the set of functions* $h : \{0,1\}^* \to S$*, where* $S := \{0,1\}^n$ *for some integer* n*. The random oracle model* $\mathcal{H}$ *is an idealized model, sampling a random function* h *from* $\mathcal{I}_{*,S}$*. Every algorithm can query* x*, obtaining the corresponding value* $h(x) \in S$*.*

Definition 5 (Random Permutation Model [RS08]). *Let $\mathcal{I}_{S,S}$ denote the set of permutations $\pi : S \to S$, where $S := \{0,1\}^n$ for some integer n. The random permutation model $\mathcal{E}$ is an idealized model, sampling a random permutation π from $\mathcal{I}_{S,S}$. Every algorithm can query $x \in S$ with $\mathcal{E}$ for both π and its inverse π^{-1}, obtaining the corresponding value $\pi(x) \in S$ or $\pi^{-1}(x) \in S$.*

Definition 6 (Generic Group Model [Sho97]). *Denote by $\mathcal{I}_{\mathbb{Z}_N,S}$ the set of injections $\sigma : \mathbb{Z}_N \mapsto S$, where $S := \{0,1\}^m$. The generic group model $\mathcal{G}_{N,m}$ is an idealized model, sampling a random injection σ from $\mathcal{I}_{\mathbb{Z}_N,S}$, with functions $\mathcal{G}_{N,m}^{\mathsf{label}}$ and $\mathcal{G}_{N,m}^{\mathsf{add}}$. Concretely, for each query $x \in \mathbb{Z}_N$, the "labeling" function $\mathcal{G}_{N,m}^{\mathsf{label}}$ responds with a value $\sigma(x) \in S$. For a query (g_1, g_2), the "adding" function $\mathcal{G}_{N,m}^{\mathsf{add}}$ answers as follows: if $g_1 = \sigma(x_1)$ and $g_2 = \sigma(x_2)$ for some $x_1, x_2 \in \mathbb{Z}_N$, replying by $\sigma(x_1 + x_2)$, and replying by $\perp$ otherwise.*

2.1.3 Notions of Reductions To establish separations between primitives, in this paper, we follow two notions, *fully black-box reduction* and *relativizing reduction*, as formalized by Reingold, Trevisan, and Vadhan [RTV04].

Definition 7 (Fully Black-Box Reduction [RTV04]). *There exists a fully black-box reduction from a primitive $\mathcal{P} := \langle \mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P} \rangle$ to a primitive $\mathcal{Q} := \langle \mathcal{F}_\mathcal{Q}, \mathcal{R}_\mathcal{Q} \rangle$, if there exist* PPT *oracle machines Π and $\mathcal{B}$ such that:*

Correctness *For every implementation $f \in \mathcal{F}_\mathcal{Q}$ we have that $\Pi^f \in \mathcal{F}_\mathcal{P}$.*
Security *For every implementation $f \in \mathcal{F}_\mathcal{Q}$, if there exists a* PPT *oracle machine $\mathcal{A}$ such that $\mathcal{A}^f$ $\mathcal{P}$-breaks Π^f, then there exists a* PPT *oracle machine $\mathcal{B}$ such that $\mathcal{B}^f$ $\mathcal{Q}$-breaks f.*

In literature, a typical technique for black-box separation, say for primitives $\mathcal{P}$ and $\mathcal{Q}$, is relativizing separation, which means that there is no relativizing reduction between $\mathcal{P}$ and $\mathcal{Q}$. Reingold et al. [RTV04] indicate that fully black-box reduction implies relativizing reduction, referring to that the relativizing separation from $\mathcal{P}$ to $\mathcal{Q}$ indicates the corresponding fully black-box separation.

Definition 8 (Relativizing Reduction [RTV04]). *There exists a relativizing reduction from a primitive $\mathcal{P} := \langle \mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P} \rangle$ to a primitive $\mathcal{Q} := \langle \mathcal{F}_\mathcal{Q}, \mathcal{R}_\mathcal{Q} \rangle$, if for every oracle $\mathcal{O}$, the primitive $\mathcal{P}$ exists relative to $\mathcal{O}$ whenever $\mathcal{Q}$ exists relative to $\mathcal{O}$. A primitive $\mathcal{P}$ is said to exist relative to $\mathcal{O}$, if there exists $f \in \mathcal{F}_\mathcal{P}$ which has an efficient implementation when having access to the oracle $\mathcal{O}$ such that no* PPT *oracle machine with access to $\mathcal{O}$, can $\mathcal{P}$-break f.*

2.2 Indifferentiability

The framework of indifferentiability is proposed by Maurer, Renner, and Holenstein [MRH04], which formalizes a set of necessary and sufficient conditions for securely replacing one cryptosystem with another in an arbitrary environment. This framework is used to justify the structural soundness of various cryptographic primitives, including hash functions [CDMP05, DRS09], block

ciphers [ABD+13, CHK+16, DSSL16, GWL23], domain extenders [CDMS10], authenticated encryption with associated data [BF18], and public key cryptosystems [ZZ20]. It can also be used to study the relationship between idealized models [ZZ23]. Within the context of the indifferentiability framework, it is customary to consider that a cryptosystem either implements certain ideal objects denoted as $\mathcal{F}$, or it is a construction denoted as $C^{\mathcal{F}'}$ that relies on underlying ideal objects $\mathcal{F}'$.

Definition 9 (Indifferentiability [MRH04]). *Let Σ_1 and Σ_2 be two cryptosystems and $\mathcal{S}$ be a simulator. The indifferentiability advantage of a distinguisher $\mathcal{D}$ against (Σ_1, Σ_2) with respect to $\mathcal{S}$ is*

$$\mathrm{Adv}^{\mathsf{indif}}_{\Sigma_1,\Sigma_2,\mathcal{S},\mathcal{D}}(1^\lambda) := \Pr[\mathrm{Real}_{\Sigma_1,\mathcal{D}}] - \Pr[\mathrm{Ideal}_{\Sigma_2,\mathcal{S},\mathcal{D}}],$$

where games $\mathrm{Real}_{\Sigma_1,\mathcal{D}}$ and $\mathrm{Ideal}_{\Sigma_2,\mathcal{S},\mathcal{D}}$ are defined in Fig. 2. We say Σ_1 is indifferentiable from Σ_2, if there exists an efficient simulator $\mathcal{S}$ such that for any efficient distinguisher $\mathcal{D}$, the advantage above is negligible. Moreover, we say Σ_1 is statistically indifferentiable from Σ_2, if there exists an efficient simulator such that, for any unbounded distinguisher $\mathcal{D}$, the advantage above is negligible.

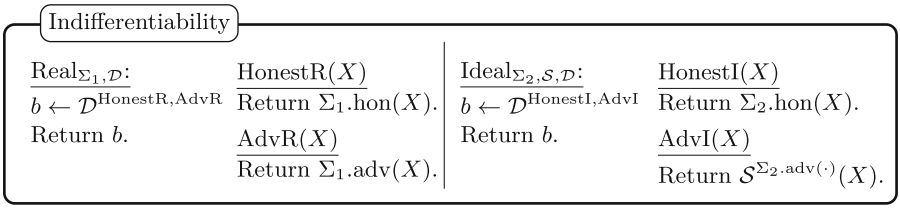

Fig. 2. Indifferentiability of Σ_1 and Σ_2, where $\mathcal{S}$ is the simulator and $\mathcal{D}$ is the adversary.

Below, we also use the notations in [BF18] and consider the definition above to two systems with interfaces as:

$$(\Sigma_1.\mathrm{hon}(X), \Sigma_1.\mathrm{adv}(x)) := (\Pi^{\mathcal{F}_1}(X), \mathcal{F}_1(x)),$$
$$(\Sigma_2.\mathrm{hon}(X), \Sigma_2.\mathrm{adv}(x)) := (\mathcal{F}_2(X), \mathcal{F}_2(x)),$$

where $\mathcal{F}_1$ and $\mathcal{F}_2$ are two ideal objects sampled from their distributions and $\Pi^{\mathcal{F}_1}$ is a construction of $\mathcal{F}_2$ by calling $\mathcal{F}_1$. Maurer, Renner, and Holenstein prove the composition theorem for the framework of indifferentiability; for simplicity, we give a game-based formalization from [RSS11].

Theorem 3 (Composition Theorem [MRH04]). *Let $\Sigma_1 := (\Pi^{\mathcal{F}_1}, \mathcal{F}_1)$ and $\Sigma_2 := (\mathcal{F}_2, \mathcal{F}_2)$ be two systems that Σ_1 is indifferentiable from Σ_2 with respect to a simulator $\mathcal{S}$, then Σ_1 is as secure as Σ_2 for any single-stage game. More concretely, let Game be a single-stage game, then for any adversary $\mathcal{A}$, there is an adversary $\mathcal{B}$ and a distinguisher $\mathcal{D}$ such that*

$$\Pr[\mathsf{Game}_{\Pi^{\mathcal{F}_1},\mathcal{A}^{\mathcal{F}_1}}] \leq \Pr[\mathsf{Game}_{\mathcal{F}_2,\mathcal{B}^{\mathcal{F}_2}}] + \mathrm{Adv}^{\mathsf{indif}}_{\Sigma_1,\Sigma_2,\mathcal{S},\mathcal{D}}.$$

The proof of Theorem 3 is straightforward; due to space limit, we skip it here. Next, we give the formal definition of the separation between two idealized models in the framework of indifferentiability against computational adversaries.

Definition 10 (Computational Indifferentiable Separation [MRH04, ZZ23]). *Let Σ_1, Σ_2 be two idealized models, we say Σ_2 is computationally indifferentiably separated from Σ_1 if for any efficient algorithm Π and any efficient simulator $\mathcal{S}$, there exists an efficient distinguisher $\mathcal{D}_{\Pi,\mathcal{S}}$ and a noticeable function ρ such that*

$$\mathrm{Adv}^{\mathsf{indif}}_{\Pi^{\Sigma_1},\Sigma_2,\mathcal{S},\mathcal{D}_{\Pi,\mathcal{S}}}(1^\lambda) := \left|\Pr[\mathrm{Real}_{\Sigma_1,\mathcal{D}_{\Pi,\mathcal{S}}}] - \Pr[\mathrm{Ideal}_{\Sigma_2,\mathcal{S},\mathcal{D}_{\Pi,\mathcal{S}}}]\right| \geq \rho(\lambda).$$

Observe that, if an idealized model Σ_2 is computationally indifferentiably separated from another idealized model Σ_1, it means that, we cannot build a scheme Π^{Σ_1} such that Π^{Σ_1} is indifferentiable from Σ_2, even under arbitrarily strong computational assumptions.

3 Separation Between Cryptographic Groups

In this section, we establish the separation between two CDH-secure groups, $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ and $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$, under the condition that both N and $(m_2 - m_1)$ are sufficiently large within the same security parameter.

Theorem 4 (Main Theorem). *Let $\lambda \in \mathbb{N}$ be the security parameter. Let N, m_1, m_2 be integers such that $N \geq 2^{\omega(\log\lambda)}, m_1 > \log N$ and $m_2 - m_1 \geq \omega(\log\lambda)$. Then there is no black-box reduction from $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$ to $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$.*

Proof. To establish the theorem, we apply the so-called two-oracle technique [HR04]. Let PSPACE be a PSPACE-complete oracle. Essentially, we set $\mathcal{O} := (\mathsf{PSPACE}, \mathcal{G}_{N,m_2})$ and prove the following:

1. $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$ exists relative to $\mathcal{O}$;
2. $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ does not exist relative to $\mathcal{O}$.

The former statement holds trivially as $\mathcal{G}_{N,m_2}$ implies $\mathcal{P}^{\mathsf{CDH}}_{N,m_2}$ in the canonical manner. Therefore, it suffices to prove the latter one.

Lemma 1. *$\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ does not exist relative to $\mathcal{O}$.*

To establish the proof, we first pin down an intermediary primitive, i.e., $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$ (within the same security parameter), defined in Sect. 2.1, and then prove that:

1. $\mathcal{P}^{\mathsf{CDH}}_{N,m_1}$ implies $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$;
2. $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$ does not exist relative to $\mathcal{O}$.

The first statement holds straightforwardly. Next, we establish our theorem by proving the following lemma.

Lemma 2. $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$ *does not exist relative to* $\mathcal{O}$.

Intuitively, to prove that $\mathcal{P}^{\mathsf{NIKE}}_{N,m_1}$ does not exist relative to $\mathcal{O}$, it is sufficient to build a PPT oracle adversary $\mathcal{A}^{\mathcal{O}}$ that breaks any construction $\Pi^{\mathcal{O}} := (\mathsf{KGen}^{\mathcal{O}}, \mathsf{SHK}^{O})$. Observe that $\mathcal{A}^{\mathcal{O}}$ has access to a PSPACE-complete oracle, which means that $\mathcal{A}^{\mathcal{O}}$ implies a computationally unbounded but query-efficient adversary that only has access to $\mathcal{G}_{N,m_2}$[16]. Therefore, it suffices to construct such an adversary $\mathcal{A}^{\mathcal{G}_{N,m_2}}$. In Fig. 3, we illustrate the description of the adversary.

We first clarify some undefined notions: Let n be a sufficiently large integer that will be specified below. By $\{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\} \xleftarrow{\text{query}} \mathsf{KGen}^{\mathcal{G}_{N,m_2}}(r_i)$, we mean that when running the algorithm $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(r_i)$, the algorithm makes queries $(\mathsf{que}_1, \ldots, \mathsf{que}_q)$ to the oracle $\mathcal{G}_{N,m_2}$ and obtains $(\mathsf{res}_1, \ldots, \mathsf{res}_q)$[17].

Next, we prove that $\mathcal{A}^{\mathcal{G}_{N,m_2}}$ outputs the valid shared key with noticeable probability. Let $S_{B\text{-}\mathsf{label}}$ be the set of the valid group elements that appear when running $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}_B)$ and $\mathsf{SHK}^{\mathcal{G}_{N,m_2}}(\mathsf{pk}_A, \mathsf{sk}_B)$; those group elements are either the responses of labeling/addition queries or the valid inputs of the addition queries. It is apparent that $|S_{B\text{-}\mathsf{label}}| \leq 6q$, due to the fact that each algorithm makes at most q queries. Now, we define:

$$S_B := \{(x,h) | h \in S_{B\text{-}\mathsf{label}},\ \mathcal{G}^{\mathsf{label}}_{N,m_2}(x) = h\}.$$

Note that, for any iteration, if the adversary successfully guesses S_B in $\hat{S}_{\mathcal{A}}$, then the shared key computed in this iteration would be valid. Specifically, in such a context, there exists an instance of the GGM that is consistent with the query views of both the adversary and the user B, and the validity of the shared key follows by the perfect correctness of $\Pi^{\mathcal{G}_{N,m_2}}$. However, without the knowledge of sk_B, $\mathcal{A}$ might not guess S_B correctly with a good probability. In fact, there are three events:

- Event 1: There exist $(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}$ and $(\mathsf{que}_B, \mathsf{res}_B) \in S_B$ such that $\mathsf{que}_A = \mathsf{que}_B$ but $\mathsf{res}_A \neq \mathsf{res}_B$.
- Event 2: There exist $(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}$ and $(\mathsf{que}_B, \mathsf{res}_B) \in S_B$ such that $\mathsf{que}_A \neq \mathsf{que}_B$ but $\mathsf{res}_A = \mathsf{res}_B$.
- Event 3: For any $(\mathsf{que}_A, \mathsf{res}_A) \in \hat{S}_{\mathcal{A}}, (\mathsf{que}_B, \mathsf{res}_B) \in S_B$, we have that if $\mathsf{que}_A = \mathsf{que}_B$ then $\mathsf{res}_A = \mathsf{res}_B$, and vice versa.

[16] Any computationally unbounded but query-efficient adversary can be simulated by a PPT oracle machine with access to a PSPACE-complete oracle, that is because what we need are specific labeling query-response tuples of GGM. These tuples can be picked by using a PSPACE-complete oracle. See [MM11] for more details.

[17] As explained above, we stress that $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}$ only makes labeling queries.

Adversary $\mathcal{A}^{\mathcal{G}_{N,m_2}}(\mathsf{pk}_A, \mathsf{pk}_B)$

$\mathcal{A}^{\mathcal{G}_{N,m_2}}(\mathsf{pk}_A, \mathsf{pk}_B)$:

$S_{\mathsf{key}} \leftarrow \emptyset$; $S_{\mathsf{que\text{-}res}} \leftarrow \emptyset$; $r_1, \ldots, r_n \xleftarrow{\$} \mathbb{Z}_N$;

Initial phase: *//collecting frequent queries*

for $i = 1$ to n:

$\{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\} \xleftarrow{\text{query}} \mathsf{KGen}^{\mathcal{G}_{N,m_2}}(r_i)$

$S_{\mathsf{que\text{-}res}} \leftarrow S_{\mathsf{que\text{-}res}} \cup \{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\}$;

for $i = 1$ to $12q + 1$: *//running $12q + 1$ iterations*

Simulation phase: *//searching a proper view*

Search a secret key $\widetilde{\mathsf{sk}}_A$ and a set of query-response tuples $\hat{S}_{\mathcal{A}}$ satisfying the following properties:

Property 1: $\hat{S}_{\mathcal{A}}$ is consistent with $S_{\mathsf{que\text{-}res}}$;

Property 2: $\mathsf{pk}_A = \mathsf{KGen}^{S_{\mathsf{que\text{-}res}} \cup \hat{S}_{\mathcal{A}}}(\widetilde{\mathsf{sk}}_A)$;

Property 3: $\hat{S}_{\mathcal{A}}$ only collects tuples of labeling queries and $|\hat{S}_{\mathcal{A}}| \leq 12q$;

Property 4: $\hat{S}_{\mathcal{A}}$ is sufficient for SHK. That is, when running the algorithm $\widetilde{\mathsf{shk}}_i \leftarrow \mathsf{SHK}^{S_{\mathsf{que\text{-}res}} \cup \hat{S}_{\mathcal{A}}}(\mathsf{pk}_B, \widetilde{\mathsf{sk}}_A)$: (1) the set $S_{\mathsf{que\text{-}res}} \cup \hat{S}_{\mathcal{A}}$ covers all labeling queries; (2) the set $S_{\mathsf{que\text{-}res}} \cup \hat{S}_{\mathcal{A}}$ is able to convert any addition query into a labeling query properly. Let $\mathsf{que} = (h_1, h_2)$ be an addition query when running the shared-key algorithm, if either h_1 or h_2 is not covered in $\hat{S}_{\mathcal{A}} \cup S_{\mathsf{que\text{-}res}}$, then responds with $\perp$; otherwise the set $S_{\mathsf{que\text{-}res}} \cup \hat{S}_{\mathcal{A}}$ must cover the following three tuples: $(x_1, h_1), (x_2, h_2)$ and $(x_1 + x_2, h_3)$, and responds with h_3.

$S_{\mathsf{key}} \leftarrow S_{\mathsf{key}} \cup \{\widetilde{\mathsf{shk}}_i\}$;

Update phase: *//updating the guessing labeling queries with valid encodings*

for each $(\mathsf{que}, \mathsf{res}) \in \hat{S}_{\mathcal{A}} \setminus S_{\mathsf{que\text{-}res}}$: $S_{\mathsf{que\text{-}res}} \leftarrow S_{\mathsf{que\text{-}res}} \cup (\mathsf{que}, \mathcal{G}^{\mathsf{label}}_{N,m_2}(\mathsf{que}))$;

Final phase: *//outputting the guessing shared key*

return the majority value in S_{key}.

Fig. 3. The description of the adversary that breaks $\Pi^{\mathcal{G}_{N,m_2}}$.

We immediately observe that event 1 occurs at most $6q$ times, because the updating phase would eliminate at least one pair in S_B. Therefore, it suffices to prove that event 2 never occurs except for negligible probability and event 3 would deduce the valid shared key with high probability. According to the description of the adversary Fig. 3, we have that in event 3, the set $\hat{S}_{\mathcal{A}} \cup S_{\mathsf{que\text{-}res}}$ responds to the labeling queries perfectly and converts the addition queries into labeling queries properly. Concretely, let $\mathsf{que} := (h_1, h_2)$ be an addition query, there are two cases: (1) $\hat{S}_{\mathcal{A}} \cup S_{\mathsf{que\text{-}res}}$ covers $(x_1, h_1), (x_2, h_2)$, and $(x_1 + x_2, h_3)$; (2) either h_1 or h_2 is not stored in $\hat{S}_{\mathcal{A}} \cup S_{\mathsf{que\text{-}res}}$. For the former case, the response is valid; for latter one, the response is invalid if and only if both h_1 and h_2 are valid group encodings. Therefore, the only bad case that prevents event 3 from deducing the valid shared key is that the adversary outputs a valid group encoding h without knowing the discrete logarithm.

Moreover, in the simulation phase, $\hat{S}_{\mathcal{A}}$ must be consistent with $S_{\mathsf{que\text{-}res}}$, which indicates that when event 2 occurs, the adversary successfully outputs a valid group encoding h without making labeling query. To bound the probability, we define that, for any $\mathsf{sk} \in \mathbb{Z}_N$:

$$Q_{\mathsf{sk}} := \left\{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\right\} \xleftarrow{\mathsf{query}} \mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}).$$

Note that the adversary only takes pk_A and pk_B as inputs, where $\mathsf{pk}_A = \mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}_A)$ and $\mathsf{pk}_B = \mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}_B)$. It is apparent that the group encoding $h \notin S_{\mathsf{que\text{-}res}}$, and we next establish our analysis by considering the following four cases:

- Case 1: (**Independent group encoding**) $h \notin Q_{\mathsf{sk}_A} \cup Q_{\mathsf{sk}_B}$
- Case 2: (**Frequent group encoding**) $h \in Q_{\mathsf{sk}_A} \cap Q_{\mathsf{sk}_B}$
- Case 3: (**Dependent but hard to extract**) $h \in Q_{\mathsf{sk}_A} \setminus Q_{\mathsf{sk}_B}$.
- Case 4: (**Dependent but hard to extract**) $h \in Q_{\mathsf{sk}_B} \setminus Q_{\mathsf{sk}_A}$.

It is apparent that, for any query-efficient adversary (might be computationally inefficient), if the probability that it outputs such an h (for all cases) is bounded, then we are done.

Case 1. We note that, h is independent of pk_A and pk_B, indicating that the probability that any adversary outputs such a h is bounded by $\frac{O(q)\cdot N}{2^{m_2}} \leq \mathsf{negl}(\lambda)$.

Case 2. We first define the frequent group encodings. Specifically, let $t := 26q^2$, we say a group encoding res is frequent if

$$\Pr[(\mathsf{que}, \mathsf{res}) \in Q_z : z \xleftarrow{\$} \mathbb{Z}_N] \geq \frac{1}{t}.$$

In such a case, we also call $(\mathsf{que}, \mathsf{res})$ as a frequent query. Note that sk_A and sk_B are uniformly sampled, therefore, for any $(\mathsf{que}, \mathsf{res}) \in Q_{\mathsf{sk}_A}$, if it is not a frequent query, then $\Pr[(\mathsf{que}, \mathsf{res}) \in Q_{\mathsf{sk}_B}] \leq \frac{1}{t}$, indicating that

$$\Pr[Q_{\mathsf{sk}_A} \cap Q_{\mathsf{sk}_B} \text{ are all frequent queries}] \geq 1 - \frac{q}{t} = 1 - \frac{1}{26q}.$$

Next, we bound the probability that $h \notin S_{\mathsf{que\text{-}res}}$ conditioned on that $Q_{\mathsf{sk}_A} \cap Q_{\mathsf{sk}_B}$ are all frequent queries. Let $n := t \cdot \lambda$, we then prove that, with a high probability, $Q_{r_1} \cup \cdots \cup Q_{r_n}$ contains all frequent queries. Essentially, there are at most $q_f := q \cdot t$ frequent queries, denoted as $\{(\mathsf{que}'_i, \mathsf{res}'_i)\}_{i\in[q_f]}$. For each $(\mathsf{que}'_i, \mathsf{res}'_i)$, we have that

$$\Pr[(\mathsf{que}'_i, \mathsf{res}'_i) \notin Q_{r_1} \cup \cdots \cup Q_{r_n}] \leq \left(1 - \frac{1}{t}\right)^n \leq e^{-\lambda},$$

which means

$$\Pr[(\mathsf{que}'_i, \mathsf{res}'_i) \in Q_{r_1} \cup \cdots \cup Q_{r_n} : \forall i \in [q_f]] \geq 1 - (q \cdot t)e^{-\lambda}.$$

Therefore,

$$\Pr[\text{Case 2}] = \Pr[h \in Q_{\mathsf{sk}_A} \cap Q_{\mathsf{sk}_B} \wedge h \notin S_{\mathsf{que\text{-}res}}] \leq \frac{1}{26q} + (q \cdot t)e^{-\lambda} \leq \frac{1}{26q} + \mathsf{negl}(\lambda).$$

Case 3. We immediately observe that pk_B is independent of h, which means that only pk_A carries the information of h. Note that the length of pk_A is m_1; in contrast, the length of h is m_2; this intuitively indicates that, over the probability of sampling the GGM instance, it is impossible to extract a valid group encoding in $Q_{\mathsf{sk}_A} \setminus (Q_{\mathsf{sk}_B} \cup S_{\mathsf{que\text{-}res}})$ except for negligible probability.

To establish the formal analysis, we strengthen the adversary $\mathcal{A}$ by providing $\mathcal{A}$ the unbounded computational power, and the following information: the tuple $(\mathsf{sk}_A, \mathsf{sk}_B, \mathsf{pk}_A, Q_{\mathsf{sk}_B}, S_{\mathsf{que\text{-}res}})$. It is easy to see that $\mathcal{A}$ itself can compute $\mathsf{pk}_A, \mathsf{pk}_B$ and $S_{\mathsf{que\text{-}res}}$, therefore it suffices to prove that

$$\Pr[\mathcal{A} \text{ outputs } h \in Q_{\mathsf{sk}_A} \setminus (Q_{\mathsf{sk}_B} \cup S_{\mathsf{que\text{-}res}})] \leq \mathsf{negl}(\lambda)$$

where the probability is over the sampling of $\mathsf{sk}_A, \mathsf{sk}_B$ and the GGM instance[18]. Observe that, pk_B is independent of h, which indicates that knowing sk_B would not increase $\mathcal{A}$'s winning probability. To further simplify the analysis, we prove a more general statement: for any secret key sk and any S (set of query-response tuples, poly-size),

$$\Pr[\mathcal{A}(\mathsf{sk}, \mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}), S) \to h : h \in Q_{\mathsf{sk}} \setminus S] \leq \mathsf{negl}(\lambda)$$

where the probability is *only* over the sampling of the GGM instance, conditioned on that the GGM instance $\mathcal{G}_{N,m_2}$ is consistent with S.

Note that, for any fixed poly-size S, the total number of the GGM instances (mapping from N to $\{0,1\}^{m_2}$) that are consistent with S is

$$(2^{m_2} - |S|) \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)).$$

Next, we introduce some notations. Note that, once the secret key sk and the GGM instance $\mathcal{G}_{N,m_2}$ are fixed, the algorithm $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk})$ is deterministic (including the queries made to $\mathcal{G}_{N,m_2}$). We here define $Q_{\mathsf{sk}\text{-}\mathcal{G}}$ as the sequence of the query-response tuples, denoted as

$$Q_{\mathsf{sk}\text{-}\mathcal{G}} := \{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\}.$$

More clearly, when running the algorithm $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk})$, the i-th query that the algorithm makes to $\mathcal{G}_{N,m_2}$ is que_i and the corresponding response is res_i. Besides, for each $(\mathsf{sk}, \mathcal{G}_{N,m_2})$, the algorithm $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk})$ outputs a public key. Next, we categorize the public keys into two types, namely the "good public keys" and the "bad public keys", with respect to the fixed secret key sk. We denote

$$T = 2^{\frac{m_2 - m_1}{2}} \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)),$$

and for any public key pk we denote S_{pk} as the set of the GGM instances such that $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}) = \mathsf{pk}$. Now, we say a public key pk (with respect to sk) is bad if $|S_{\mathsf{pk}}| \leq T$, otherwise we say the public key is good. Note that, given a

[18] The instance of GGM must be consistent with $Q_{\mathsf{sk}_B} \cup S_{\mathsf{que\text{-}res}}$.

bad public key pk (e.g., $|S_{\mathsf{pk}}| = 1$), the adversary might output a valid group encoding, thus we need to prove that, over the sampling of the GGM instance,

$$\Pr[\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}) \text{ is bad}] \leq \mathsf{negl}(\lambda).$$

Note that the space of public keys is $\{0,1\}^{m_1}$, which means that there are at most 2^{m_1} public keys. Therefore, the counting of the GGM instances that induce to a bad public key is bounded by $2^{m_1} \times T$, referring to

$$\Pr[\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}) \text{ is bad}] \leq \frac{2^{m_1} \cdot 2^{\frac{m_2-m_1}{2}}}{(2^{m_2} - |S|)} \leq \frac{1}{2^{\frac{m_2-m_1}{2}} - |S|} \leq \mathsf{negl}(\lambda).$$

Hence, it suffices to prove that, given any good public key, any adversary $\mathcal{A}$ cannot extract a valid group encoding $h \in Q_{\mathsf{sk}} \setminus S$ except for negligible probability.

For readability, we first elaborate the analysis in the case that $S = \emptyset$, where the adversary only has knowledge of $(\mathsf{sk}, \mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk}))$. Let str be any string in $\{0,1\}^{m_2}$, we denote S_{str} as the set of GGM instances such that $\mathsf{str} \in Q_{\mathsf{sk}\text{-}\mathcal{G}}$. Therefore it is sufficient to prove that, for any $\mathsf{str} \in \{0,1\}^{m_2}$, the size of S_{str} is much smaller than T (in this special case, $|S| = 0$). Specifically, by having that

$$T > 2^{\frac{m_2-m_1}{2}} \cdot (2^{m_2} - 1) \cdots (2^{m_2} - (N-1))$$

we prove that

$$|S_{\mathsf{str}}| \leq q \cdot (2^{m_2} - 1) \cdots (2^{m_2} - (N-1))$$

Note that, once the secret key sk and the GGM instance $\mathcal{G}_{N,m_2}$ are fixed, the algorithm $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk})$ is deterministic. We next illustrate an observation about $Q_{\mathsf{sk}\text{-}\mathcal{G}}$. Let $\mathcal{G}_{N,m_2}$ and $\mathcal{G}'_{N,m_2}$ be two different instances of GGM, and we denote

$$Q_{\mathsf{sk}\text{-}\mathcal{G}} := \{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\}$$
$$Q_{\mathsf{sk}\text{-}\mathcal{G}'} := \{(\mathsf{que}'_1, \mathsf{res}'_1), \ldots, (\mathsf{que}'_q, \mathsf{res}'_q)\}$$

We claim that either $Q_{\mathsf{sk}\text{-}\mathcal{G}} = Q_{\mathsf{sk}\text{-}\mathcal{G}'}$ or $\exists i \in [q]$ such that $\mathsf{res}_i \neq \mathsf{res}'_i$. In other words, it is impossible that $Q_{\mathsf{sk}\text{-}\mathcal{G}} \neq Q_{\mathsf{sk}\text{-}\mathcal{G}'}$ but $(\mathsf{res}_1, \ldots, \mathsf{res}_q) = (\mathsf{res}'_1, \ldots, \mathsf{res}'_q)$. In fact, if such an event occurs, then there exists an index $j \in [q]$ such that (1) $\forall i < j$, $(\mathsf{que}_i, \mathsf{res}_i) = (\mathsf{que}'_i, \mathsf{res}'_i)$; (2) $\mathsf{que}_j \neq \mathsf{que}'_j$, which contradicts to that $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\mathsf{sk})$ is deterministic.

This observation illustrates that $Q_{\mathsf{sk}\text{-}\mathcal{G}}$ can be represented only by $(\mathsf{res}_1, \ldots, \mathsf{res}_q)$; that is, once the sequence of the responses is fixed, then the corresponding sequence of the queries is also settled down. We denote

$$V = ((2^{m_2} - q) \cdots (2^{m_2} - (N-1)))$$

and note that for each response sequence $(\mathsf{res}_1, \ldots, \mathsf{res}_q)$, there are exactly V numbers of GGM instances that would induce it.

Next, we compute the upper bound of $|S_{\mathsf{str}}|$. If str appears in the sequence $(\mathsf{res}_1, \ldots, \mathsf{res}_q)$, then there exists an index i such that $\mathsf{res}_i = \mathsf{str}$. For the rest,

we maximize the possibility and have that the number of all possible sequences that contain str is bounded by

$$q \cdot ((2^{m_2} - 1) \cdots (2^{m_2} - (q-1))).$$

Combining the above together, we have that

$$|S_{\mathsf{str}}| \leq q \cdot (2^{m_2} - 1) \cdots (2^{m_2} - (N-1)).$$

In the following, we extend our analysis into the general case, where S is poly-size and

$$T = 2^{\frac{m_2 - m_1}{2}} \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1))$$

We immediately observe that, the upper bound above does not serve our purpose any more. The reason is that the upper bound above is calculated over all possible GGM instances, while what we need to count are the ones over the GGM instances that are consistent with S.

It is apparent that $Q_{\mathsf{sk}\text{-}\mathcal{G}}$ can be still represented by the sequence of responses when $S \neq \emptyset$. To complete the analysis, we then illustrate an additional observation about $Q_{\mathsf{sk}\text{-}\mathcal{G}}$. Let $(\mathsf{res}_1, \ldots, \mathsf{res}_q)$ and $(\mathsf{res}'_1, \ldots, \mathsf{res}'_q)$ be two different sequences. We claim it is impossible that there exists an index $j \in [q]$ such that(1) $\forall i < j$, $\mathsf{res}_i = \mathsf{res}'_i$; (2)$\mathsf{res}_j \in S$ but $\mathsf{res}'_j \notin S$[19]. More specifically, given the statement that $\forall i < j$, $\mathsf{res}_i = \mathsf{res}'_i$, it is apparent that $\mathsf{que}_j = \mathsf{que}'_j$. Moreover, by having $(\mathsf{que}_j, \mathsf{res}_j) \in S$, we claim that the response of que'_j must be res_j, because the GGM instances must be consistent with S. Based on this new observation, we next prove the upper bound by induction.

Let str be a string such that $\mathsf{str} \notin S$ (note that the adversary's goal is to output a valid group encoding without knowing the discrete logarithm), we denote $S_{\mathsf{str}\text{-}k}$ as the set of the GGM instances such that: (1) the algorithm $\mathsf{KGen}^{\mathcal{G}_{N,m_2}}(\cdot)$ makes k queries; (2) $\mathsf{str} \in Q_{\mathsf{sk}\text{-}\mathcal{G}} \setminus S$. We then prove that for any k,

$$|S_{\mathsf{str}\text{-}k}| \leq k \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)).$$

We first compute $|S_{\mathsf{str}\text{-}1}|$. Note that que_1 is always fixed, and if $\mathsf{que}_1 \in S$[20], then $|S_{\mathsf{str}\text{-}1}| = 0$ because str would never appear. On the other hand, if $\mathsf{que}_1 \notin S$, then the response must be str because str appears. Thus, the counting of the GGM instances that are consistent with $S \cup \{(\mathsf{que}_1, \mathsf{str})\}$ is

$$1 \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1))$$

Note that, the response of que_1 is str if and only if those GGM instances are sampled. Moreover, based on our second observation, we have that, either $\mathsf{res}_1 \in S$ or $\mathsf{res}_1 \notin S$. Hence,

$$|S_{\mathsf{str}\text{-}1}| \leq \max\{0, 1 \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1))\}.$$

[19] We here abuse the notation $\mathsf{res}_j \in S$ by meaning that there exists a query/response tuple in S with the response res_j.

[20] We here abuse the notation $\mathsf{que}_1 \in S$ by meaning that there exists a query/response pair in S with the query is que_1.

Next, given the assumption that

$$|S_{\mathsf{str}\text{-}i}| \leq i \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)),$$

we prove

$$|S_{\mathsf{str}\text{-}(i+1)}| \leq (i+1) \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)),$$

Again, que_1 is always fixed, and if $\mathsf{que}_1 \in S$, then $|S_{\mathsf{str}\text{-}(i+1)}|$ is bounded by $|S_{\mathsf{str}\text{-}i}|$, because the response of que_1 is always fixed by S, and str must appear in the last i queries. Thus, it suffices to prove that $|S_{\mathsf{str}\text{-}(i+1)}|$ is properly bounded when $\mathsf{que}_1 \notin S$. Next we consider two scenarios:

- Scenario 1: $\mathsf{res}_1 = \mathsf{str}$;
- Scenario 2: $\mathsf{res}_1 \neq \mathsf{str}$.

Observe that scenario 1 occurs if and only if the GGM instances that are consistent with $S \cup \{(\mathsf{que}_1, \mathsf{str})\}$ are selected. Therefore, the counting of those GGM instances is:

$$1 \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)).$$

When scenario 2 occurs, there are at most $2^{m_2} - (|S|+1)$ options for res_1. Once the response of que_1 is fixed, say $(\mathsf{que}_1, \mathsf{str}')$, we apply the induction. Specifically, we denote $S' = S \cup \{(\mathsf{que}_1, \mathsf{str}')\}$ ($|S'| = |S|+1$). Note that scenario 2 occurs means that str appears in the last i queries conditioned on that all the GGM instances are consistent with S'. Applying the assumption, we have that the counting of the GGM instances is bounded by

$$\begin{aligned}&(2^{m_2} - (|S|+1)) \cdot i \cdot (2^{m_2} - (|S'|+1)) \cdots (2^{m_2} - (N-1)) \\ =&i \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)).\end{aligned}$$

Now, we see that, if $\mathsf{que}_1 \notin S$ (combining both scenario 1 and scenario 2), then

$$|S_{\mathsf{str}\text{-}(i+1)}| \leq (i+1) \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1)).$$

Again, res_1 is either in S or not in S. We have that

$$\begin{aligned}|S_{\mathsf{str}\text{-}(i+1)}| &\leq \max\{|S_{\mathsf{str}\text{-}i}|, (i+1) \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1))\} \\ &= (i+1) \cdot (2^{m_2} - (|S|+1)) \cdots (2^{m_2} - (N-1))\end{aligned}$$

By setting $\mathsf{sk} := \mathsf{sk}_A$ and $S := S_{\mathsf{que}\text{-}\mathsf{res}}$, we have that the probability that the adversary outputs $h \in Q_{\mathsf{sk}_A} \setminus S_{\mathsf{que}\text{-}\mathsf{res}}$ is bounded by $\frac{O(q^2)}{2^{\frac{m_2 - m_1}{2}}} \leq \mathsf{negl}(\lambda)$.

Case 4. It is trivial that

$$\Pr[\text{Case 4}] = \Pr[\text{Case 3}].$$

Combining together, we have that

$$\begin{aligned}\Pr[\mathcal{A}^{\mathcal{G}_{N,m_2}} \text{ outputs the valid shared key}] &\geq 1 - (6q+1)(\frac{1}{26q} + \mathsf{negl}(\lambda)) \\ &\geq \frac{2}{3} - \mathsf{negl}(\lambda).\end{aligned}$$

4 The Hierarchy of GGMs

In this section, we establish a hierarchy among GGMs, varying in distinct lengths of group encodings and prove that the shorter GGM is strictly stronger than the longer GGM. Specifically, we show that one can construct an indifferentiable longer generic group from a shorter one plus an additional independent random oracle, but the shorter generic group model is computationally indifferentiably separated from the longer generic group (when the gap between the lengths is sufficiently large).

4.1 $\mathcal{G}_{N,m_1}$ Statistically Implies $\mathcal{G}_{N,m_2}$

In this section, we show how to build an longer indifferentiable generic group model from a shorter one plus an additional independent ROM. Here are the building blocks:

- $\mathcal{G}_{N,m_1} := (\mathcal{G}^{\mathsf{label}}_{N,m_1}, \mathcal{G}^{\mathsf{add}}_{N,m_1})$ is a generic group model that maps $\mathbb{Z}_N$ to $\{0,1\}^{m_1}$;
- $\mathcal{E} : \{0,1\}^{m_2} \to \{0,1\}^{m_2}$ is a random permutation oracle with its inverse $\mathcal{E}^{-1}$.

For simplicity, we denote $\mathcal{O}$ as the tuple $(\mathcal{G}_{N,m_1}, (\mathcal{E}, \mathcal{E}^{-1}))$. The following is the construction $\Pi^{\mathcal{O}}_{\mathsf{L\text{-}GGM}} := (L^{\mathcal{O}}_{\mathsf{L\text{-}GGM}}, A^{\mathcal{O}}_{\mathsf{L\text{-}GGM}})$, depicted in Fig. 4. Correctness easily follows, and it rests to prove the indifferentiability. Formally,

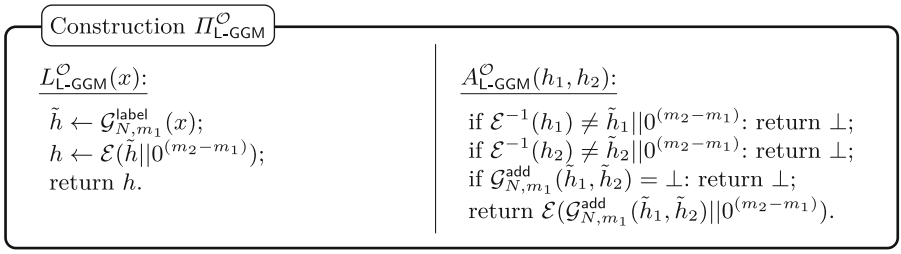

Fig. 4. The construction $\Pi^{\mathcal{O}}_{\mathsf{L\text{-}GGM}}$ in the $\mathcal{G}_{N,m_1}$ and RPM.

Theorem 5. *Let m_1, m_2 be two integers that $m_2 \geq m_1$. The scheme $\Pi^{\mathcal{O}}_{\mathsf{L\text{-}GGM}}$ in Fig. 4, with access to a generic group $\mathcal{G}_{N,m_1}$, a random permutation $\mathcal{E}$ and its inverse $\mathcal{E}^{-1}$, is indifferentiable from a generic group $\mathcal{G}_{N,m_2}$. More precisely, there exists a simulator $\mathcal{S}$ such that for all $(q_{\mathcal{G}^{\mathsf{label}}_{N,m_1}}, q_{\mathcal{G}^{\mathsf{add}}_{N,m_1}}, q_{\mathcal{E}}, q_{\mathcal{E}^{-1}})$-query distinguisher $\mathcal{D}$ with $q_{\mathcal{G}^{\mathsf{label}}_{N,m_1}} + q_{\mathcal{G}^{\mathsf{add}}_{N,m_1}} + q_{\mathcal{E}} + q_{\mathcal{E}^{-1}} \leq q$, we have*

$$\mathrm{Adv}^{\mathsf{indif}}_{\Pi^{\mathcal{O}}_{\mathsf{L\text{-}GGM}},\mathcal{G}_{N,m_2},\mathcal{S},\mathcal{D}} \leq \frac{6q^2}{N} + \frac{10q^2+4q}{2^{m_1}} + \frac{3q}{2^{\lambda}} + \frac{2q}{2^{m_1}-2q}.$$

The simulator makes at most $3q$ queries to $\mathcal{G}_{N,m_2}$.

Due to space limit, we leave the proof in the full version of this paper [ZJW+24].

4.2 $\mathcal{G}_{N,m_2}$ Does Not Computationally Imply $\mathcal{G}_{N,m_1}$

In this section, we show that the shorter GGM is computationally indifferentiably separated from the longer one. Formally,

Theorem 6. *Let λ be the security parameter. Let $\mathcal{G}_{N,m_1}$ and $\mathcal{G}_{N,m_2}$ be two generic group models. If $(m_2 - m_1) \geq \omega(\log \lambda)$, then $\mathcal{G}_{N,m_1}$ is computationally indifferentiably separated from $\mathcal{G}_{N,m_2}$.*

To prove it, we adopt the discrete logarithm identification (DLI) problem proposed by [ZZ23]. To absorb Zhang and Zhandry 's analysis into our setting, we propose the DLI problem w.r.t the shorter groups in the longer GGM. Below, we give the proof sketch of Theorem 6 and the formal proof can be found in the full version of this paper [ZJW+24].

Proof Sketch. Suppose $\Pi^{\mathcal{G}_{N,m_2}} := (L^{\mathcal{G}_{N,m_2}}, A^{\mathcal{G}_{N,m_2}})$ is indifferentiable from $\mathcal{G}_{N,m_1}$ in the longer GGM $\mathcal{G}_{N,m_2}$. The argument goes in three steps:

1. DLI w.r.t. $\Pi^{\mathcal{G}_{N,m_2}}$ is easy.
2. If $\Pi^{\mathcal{G}_{N,m_2}}$ is indifferentiable from $\mathcal{G}_{N,m_1}$ and DLI w.r.t. $\Pi^{\mathcal{G}_{N,m_2}}$ is easy, then DLI w.r.t. $\mathcal{G}_{N,m_1}$ is also easy.
3. Yet, DLI w.r.t. the generic group $\mathcal{G}_{N,m_1}$ is hard.

The above three steps draw a contradiction, so the statement "$\Pi^{\mathcal{G}_{N,m_2}}$ is indifferentiable from $\mathcal{G}_{N,m_1}$" cannot be true, completing our proof. Note that, Step 2 is already proven in [ZZ23]; and the proof of Step 3 is straightforward according to Definition 9 for indifferentiability. Due to the space limit, we skip them here. Below, we prove Step 1.

By the definition of indifferentiability, the algorithms $L^{\mathcal{G}_{N,m_2}}$ and $A^{\mathcal{G}_{N,m_2}}$ are deterministic; and they shall support group operations correctly with high probability. We stress that $L^{\mathcal{G}_{N,m_2}}$ only makes labeling queries. Let q be an integer in $\mathsf{poly}(\lambda)$. We assume that both $L^{\mathcal{G}_{N,m_2}}$ and $A^{\mathcal{G}_{N,m_2}}$ make at most q queries to $\mathcal{G}_{N,m_2}$. Next, we prove that the DLI problem w.r.t. $\Pi^{\mathcal{G}_{N,m_2}}$ is easy by constructing an efficient adversary $\mathcal{A}$ and a query-free circuit $C_{\mathsf{G\text{-}GGM}}$ in Fig. 5. (Here, G-GGM denotes the shorter group in the longer GGM.)

We first clarify some undefined notions in Fig. 5. Let n be a sufficiently large integer to be specified below. By $\{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\} \xleftarrow{\text{query}} L^{\mathcal{G}_{N,m_2}}(r_i)$, we denote that on input r_i, the algorithm $L^{\mathcal{G}_{N,m_2}}(r_i)$ makes queries $(\mathsf{que}_1, \ldots, \mathsf{que}_q)$ to $\mathcal{G}_{N,m_2}$ and gets responses of $(\mathsf{res}_1, \ldots, \mathsf{res}_q)$; and similar for the notation $\{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\} \xleftarrow{\text{query}} A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(x - z), L^{\mathcal{G}_{N,m_2}}(z))$.[21] Given an input $z \in \mathbb{Z}_N$, the query-free circuit $C_{\mathsf{G\text{-}GGM}}$ runs algorithm $L^{\mathcal{G}_{N,m_2}}(z)$ except for replacing the querying oracle by looking up the table $S_{\mathsf{que\text{-}res}}$ (and lazy sampling); we denote that as $L^{S_{\mathsf{que\text{-}res}}}$.

We argue that the query-free circuit $C_{\mathsf{G\text{-}GGM}}$ in Fig. 5 identifies x with a good probability, which means DLI w.r.t. $\Pi^{\mathcal{G}_{N,m_2}}$ is easy. Note that, we say $C_{\mathsf{G\text{-}GGM}}$

[21] Here, we abuse the notation $L^{\mathcal{G}_{N,m_2}}(x-z)$ as both the group element and the labeling operation on $x - z$.

Adversary $\mathcal{A}^{\mathcal{G}_{N,m_2}}$

$\underline{\mathcal{A}^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(x)):}$

$S_{\mathsf{que\text{-}res}} \leftarrow \emptyset$; $z, r_1, \ldots, r_n \xleftarrow{\$} \mathbb{Z}_N$;
for $i = 1$ to n: *//collecting frequent queries*
  $\{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\} \xleftarrow{\text{query}} L^{\mathcal{G}_{N,m_2}}(r_i)$;
  $S_{\mathsf{que\text{-}res}} \leftarrow S_{\mathsf{que\text{-}res}} \cup \{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\}$;
compute $L^{\mathcal{G}_{N,m_2}}(z)$, $L^{\mathcal{G}_{N,m_2}}(x - z) \leftarrow A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(x), L^{\mathcal{G}_{N,m_2}}(-z))$;
$\{(\mathsf{que}_1, \mathsf{res}_1), \ldots, (\mathsf{que}_q, \mathsf{res}_q)\} \xleftarrow{\text{query}} A^{\mathcal{G}_{N,m_2}}(L^{\mathcal{G}_{N,m_2}}(x - z), L^{\mathcal{G}_{N,m_2}}(z))$;
for $j = 1$ to q: *//collecting queries in group addition*
  if que_j is an addition query: *//converting addition queries into labeling queries*
    parse the addition query que_j to two labels h_1, h_2;
    if $\exists (x_1, h_1), (x_2, h_2) \in S_{\mathsf{que\text{-}res}}$: $S_{\mathsf{que\text{-}res}} \leftarrow S_{\mathsf{que\text{-}res}} \cup \{(x_1 + x_2, \mathsf{res}_j)\}$;
  else: $S_{\mathsf{que\text{-}res}} \leftarrow S_{\mathsf{que\text{-}res}} \cup \{(\mathsf{que}_j, \mathsf{res}_j)\}$; *//collecting labeling queries*
return $C_{\mathsf{G\text{-}GGM}}(\cdot, S_{\mathsf{que\text{-}res}}, L^{\mathcal{G}_{N,m_2}}(x))$.

$\underline{C_{\mathsf{G\text{-}GGM}}(\cdot, S_{\mathsf{que\text{-}res}}, L^{\mathcal{G}_{N,m_2}}(x)):}$

take $z \in \mathbb{Z}_N$ as input; run $\mathsf{str} \leftarrow L^{S_{\mathsf{que\text{-}res}}}(z)$;
when L makes a labeling query $\mathsf{que} = x$: *//responding to labeling queries*
  if $\exists (x, h) \in S_{\mathsf{que\text{-}res}}$: respond with h;
  else: respond with a uniformly sampled h; $S_{\mathsf{que\text{-}res}} \leftarrow S_{\mathsf{que\text{-}res}} \cup \{(x, h)\}$;
when L makes an addition query $\mathsf{que} = (h_1, h_2)$: *//responding to addition queries*
  if $\exists (x_1, h_1), (x_2, h_2) \in S_{\mathsf{que\text{-}res}}$:
    if $\exists (x_1 + x_2, h) \in S_{\mathsf{que\text{-}res}}$: respond with h;
    else: respond with a uniformly sampled h; $S_{\mathsf{que\text{-}res}} \leftarrow S_{\mathsf{que\text{-}res}} \cup \{(x_1 + x_2, h)\}$;
  else: respond with $\perp$; *//if addition query has a new labeling, then responds with* $\perp$
if $\mathsf{str} = L^{\mathcal{G}_{N,m_2}}(x)$: return 1; else: return 0.

Fig. 5. Efficient Adversary $\mathcal{A}^{\mathcal{G}_{N,m_2}}$ and query-free circuit $C_{\mathsf{G\text{-}GGM}}$ w.r.t. $\Pi^{\mathcal{G}_{N,m_2}}$.

identifies x with a good probability if it satisfies following properties. Due to the space limit, we leave the proof in the full version of this paper [ZJW+24].

- $\Pr[C_{\mathsf{G\text{-}GGM}}(x) = 1] \geq \frac{2}{3}$;
- for any noticeable function ρ: $\Pr_{x' \neq x}[C_{\mathsf{G\text{-}GGM}}(x') = 1] \leq \rho$.

Acknowledgment. Cong Zhang was supported by the National Key Research and Development Program of China (Grant No. 2023YFB3106000). This work was also supported by Ant Group through CCF-Ant Research Fund (Grant No. CCF-AFSG RF20230308). Bingsheng Zhang was supported by the National Natural Science Foundation of China (Grant No. 62072401 and No. 62232002) and Input Output (iohk.io). Hong-Sheng Zhou was supported in part by NSF grant CNS-1801470 and a VCU Research Quest grant.

References

ABD+13. Elena Andreeva, Andrey Bogdanov, Yevgeniy Dodis, Bart Mennink, and John P. Steinberger.On the indifferentiability of key-alternating ciphers.In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 531–550. Springer, Heidelberg, August 2013.

Bar20. Elaine Barker. Recommendation for key management: Part 1 – general, 2020. https://doi.org/10.6028/NIST.SP.800-57pt1r5.

BF18. Manuel Barbosa and Pooya Farshim. Indifferentiable authenticated encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 187–220. Springer, Heidelberg, August 2018.

BKSY11. Zvika Brakerski, Jonathan Katz, Gil Segev, and Arkady Yerukhimovich. Limits on the power of zero-knowledge proofs in cryptographic constructions.In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 559–578. Springer, Heidelberg, March 2011.

BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.

CDMP05. Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function.In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 430–448. Springer, Heidelberg, August 2005.

CDMS10. Jean-Sébastien Coron, Yevgeniy Dodis, Avradip Mandal, and Yannick Seurin. A domain extender for the ideal cipher. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 273–289. Springer, Heidelberg, February 2010.

CHK+16. Jean-Sébastien Coron, Thomas Holenstein, Robin Künzler, Jacques Patarin, Yannick Seurin, and Stefano Tessaro.How to build an ideal cipher: The indifferentiability of the Feistel construction.*Journal of Cryptology*, 29(1):61–114, January 2016.

CMR+23. Lily Chen, Dustin Moody, Karen Randall, Andrew Regenscheid, and Angela Robinson. Recommendations for discrete logarithm-based cryptography: Elliptic curve domain parameters, 2023. https://doi.org/10.6028/NIST.SP.800-186.

DH76. Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theory*, 22(6):644–654, 1976.

DHH+21. Nico Döttling, Dominik Hartmann, Dennis Hofheinz, Eike Kiltz, Sven Schäge, and Bogdan Ursu.On the impossibility of purely algebraic signatures.In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part III*, volume 13044 of *LNCS*, pages 317–349. Springer, Heidelberg, November 2021.

DRS09. Yevgeniy Dodis, Thomas Ristenpart, and Thomas Shrimpton. Salvaging Merkle-Damgård for practical applications. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 371–388. Springer, Heidelberg, April 2009.

DSSL16. Yevgeniy Dodis, Martijn Stam, John P. Steinberger, and Tianren Liu. Indifferentiability of confusion-diffusion networks. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 679–704. Springer, Heidelberg, May 2016.

ElG85. Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.

GMM17. Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. When does functional encryption imply obfuscation? In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 82–115. Springer, Heidelberg, November 2017.

GWL23. Chun Guo, Lei Wang, and Dongdai Lin. Impossibility of indifferentiable iterated blockciphers from 3 or less primitive calls. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part IV*, volume 14007 of *LNCS*, pages 408–439. Springer, Heidelberg, April 2023.

HKT11. Thomas Holenstein, Robin Künzler, and Stefano Tessaro. The equivalence of the random oracle model and the ideal cipher model, revisited. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 89–98, New York, NY, USA, 2011. Association for Computing Machinery.

HMQS23. Mohammad Hajiabadi, Mohammad Mahmoody, Wei Qi, and Sara Sarfaraz. Lower bounds on assumptions behind registration-based encryption. In Guy Rothblum and Hoeteck Wee, editors, *Theory of Cryptography*, pages 306–334, Cham, 2023. Springer Nature Switzerland.

HR04. Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 92–105. Springer, Heidelberg, August 2004.

IR89. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st ACM STOC*, pages 44–61. ACM Press, May 1989.

Mau05. Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Heidelberg, December 2005.

MM11. Takahiro Matsuda and Kanta Matsuura. On black-box separations among injective one-way functions. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 597–614. Springer, Heidelberg, March 2011.

MMN16. Mohammad Mahmoody, Ameer Mohammed, and Soheil Nematihaji. On the impossibility of virtual black-box obfuscation in idealized models. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 18–48. Springer, Heidelberg, January 2016.

MPZ20. Ueli Maurer, Christopher Portmann, and Jiamin Zhu. Unifying generic group models. Cryptology ePrint Archive, Report 2020/996, 2020. https://eprint.iacr.org/2020/996.

MRH04. Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology.In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, Heidelberg, February 2004.

Nec94. Vassiliy Ilyich Nechaev. Complexity of a determinate algorithm for the discrete logarithm. *Mathematical Notes*, 55(2):165–172, 1994.

PH78. Stephen Pohlig and Martin Hellman. An improved algorithm for computing logarithms over GF(p) and its cryptographic significance (Corresp.). *IEEE Transactions on Information Theory*, 24(1):106–110, 1978.

Pol78. John M Pollard. Monte Carlo methods for index computation (mod p). *Mathematics of Computation*, 32(143):918–924, 1978.

PRV12. Periklis A. Papakonstantinou, Charles W. Rackoff, and Yevgeniy Vahlis. How powerful are the DDH hard groups? Cryptology ePrint Archive, Report 2012/653, 2012. https://eprint.iacr.org/2012/653.

RS08. Phillip Rogaway and John P. Steinberger. Constructing cryptographic hash functions from fixed-key blockciphers. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 433–450. Springer, Heidelberg, August 2008.

RSS11. Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferentiability framework.In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 487–506. Springer, Heidelberg, May 2011.

RSS20. Lior Rotem, Gil Segev, and Ido Shahaf. Generic-group delay functions require hidden-order groups. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 155–180. Springer, Heidelberg, May 2020.

RTV04. Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 1–20. Springer, Heidelberg, February 2004.

SGS20. Gili Schul-Ganz and Gil Segev. Accumulators in (and beyond) generic groups: Non-trivial batch verification requires interaction.In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 77–107. Springer, Heidelberg, November 2020.

SGS21. Gili Schul-Ganz and Gil Segev. Generic-group identity-based encryption: A tight impossibility result. In *Information Theoretic Cryptography*, 2021.

Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997.

Zha22. Mark Zhandry. To label, or not to label (in generic groups).In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part III*, volume 13509 of *LNCS*, pages 66–96. Springer, Heidelberg, August 2022.

ZJW+24. Cong Zhang, Keyu Ji, Taiyu Wang, Bingsheng Zhang, Hong-Sheng Zhou, Xin Wang, and Kui Ren. On the complexity of cryptographic groups and generic group models. In *Cryptology ePrint Archive, Paper 2024/1452*, 2024. https://eprint.iacr.org/2024/1452.

ZZ18. Mark Zhandry and Cong Zhang. Impossibility of order-revealing encryption in idealized models. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 129–158. Springer, Heidelberg, November 2018.

ZZ20. Mark Zhandry and Cong Zhang. Indifferentiability for public key cryptosystems. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 63–93. Springer, Heidelberg, August 2020.

ZZ23. Cong Zhang and Mark Zhandry. The relationship between idealized models under computationally bounded adversaries. In *ASIACRYPT 2023*, 2023.