

The Multi-user Constrained PRF Security of Generalized GGM Trees for MPC and Hierarchical Wallets*

Chun Guo^{1,2,3}, Xiao Wang⁴, Xiang Xie⁵, and Yu Yu⁶

¹ School of Cyber Science and Technology, Shandong University, Qingdao, China

² Key Laboratory of Cryptologic Technology and Information Security of Ministry of Education, Shandong University, Qingdao, Shandong, 266237, China,

³ Shandong Research Institute of Industrial Technology, Jinan, Shandong, 250102, China

⁴ Northwestern University

⁵ Shanghai Qizhi Institute & PADO Labs

⁶ Shanghai Jiao Tong University

chun.guo.sc@gmail.com, wangxiao@cs.northwestern.edu, xiexiangiscas@gmail.com, yuyu@cs.sjtu.edu.cn

Abstract. Multi-user (μ) security considers large-scale attackers that, given access to a number of cryptosystem instances, attempt to compromise at least one of them. We initiate the study of μ security of the so-called GGM tree that stems from the PRG-to-PRF transformation of Goldreich, Goldwasser, and Micali, with a goal to provide references for its recently popularized use in applied cryptography. We propose a generalized model for GGM trees and analyze its μ *prefix-constrained PRF* security in the random oracle model. Our model allows to derive concrete bounds and improvements for various protocols, and we show-case on the Bitcoin-Improvement-Proposal standard Bip32 hierarchical wallets and function secret sharing (FSS) protocols. In both scenarios, we propose improvements with better performance and concrete security bounds at the same time. Compared with the state-of-the-art designs, our SHACAL3- and KECCAK- p -based Bip32 variants reduce the communication cost of MPC-based implementations by 73.3%~93.8%, while our AES-based FSS substantially improves μ security while reducing computations by 50%.

* Published at ACM Transactions on Privacy and Security (TOPS, previously known as TISSEC) in 04/2023: <https://dl.acm.org/doi/abs/10.1145/3592608>. Technical difference between this doc and the published version is minor.

Table of Contents

The Multi-user Constrained PRF Security of Generalized GGM Trees for MPC and Hierarchical Wallets	1
<i>Chun Guo, Xiao Wang, Xiang Xie, and Yu Yu</i>	
1 Introduction	2
1.1 Our Contribution	4
1.2 Related Work	7
1.3 Organization	8
2 Preliminaries	8
3 A Framework for Generalized GGM Trees	13
3.1 Intuitions	13
3.2 Our General Tree Model	15
4 Multi-user Leakage CPRF Security of GGGM Trees	16
4.1 Concrete Settings	17
4.2 Random Oracle-based Trees	18
4.3 Davies-Meyer-based Trees	26
5 Multi-user Leakage PPRF Security of GGGM Trees	30
6 Improving Hierarchical Deterministic Wallets	31
6.1 Hierarchical Deterministic Wallets	31
6.2 Formalism of Public-underivable HDW	32
6.3 Bip32 HDW and the Underlying GGGM Instance	33
6.4 μ Security Definitions for Hierarchical Deterministic Wallet ...	35
6.5 Multi-user Security of Bip32	36
6.6 Improving Prim for Bip32, and Performance of MPC Implementations	40
7 Improving Function Secret Sharing	41
A The LBC Bitcoin Cracking Project	48

1 Introduction

GGM tree. Pseudorandom functions (PRFs) and Pseudorandom generators (PRGs) are fundamental building blocks of virtually all cryptosystems. The PRG-to-PRF transformation of Goldreich, Goldwasser, and Micali, dubbed GGM tree [GGM86], was originally proposed to show the theoretical feasibility of constructing PRFs from any one-way function. Interestingly, subsequent works exhibited various “beyond-PRF” security for the GGM tree. Importantly for us, Boneh et al. [BW13,BGI14] showed a *constrained PRF* based on GGM tree with prefix predicates, meaning that it allows to delegate “constrained” keys that can only be used to evaluate the GGM PRF on a subset of inputs with certain prefixes. This can also be used as a *puncturable PRF*, the “dual” of constrained PRFs. These features have motivated extensive practical uses of the GGM: as a constrained PRF in delegatable computations [KPTZ13] and

functional signatures [BGI14], as a hierarchical access mechanism in the Bip32 hierarchical deterministic wallets (HDWs), as a puncturable PRF in Function-Secret Sharing (FSS) protocols [BGI15,BGI16]), and as part of the construction of a private puncturable PRF in Pseudorandom Correlation Generator (PCG) [BCGI18,BCG⁺19b] which supports more sophisticated applications including (RAM-based) secure multi-party computation [Ds17,SGRR19] and zero-knowledge proofs [BMRS21,DIO21].

Multi-user security. The classical (constrained/puncturable) PRF security of a function family F concerns with the behavior of $F(K, \cdot)$ under a (fixed) secret key K , which is now known as the *single-user (su) setting*. Due to the pervasive uses, a large number of GGM tree instances using the same specification and independently chosen keys become available from different entities across the Internet. For example, the Exodus Platform is based on Bip32 and have millions of active users [Exo21]. This effectively creates millions of independently keyed “Bip32 key trees” or generalized GGM trees. As another example, to compute stable marriage (as in [DEs16]), a two-party computation (2PC) protocol (like GMW) needs to evaluate 2^{38} AND gates using 2^{39} OT executions. Using the state-of-the-art PCG-based OT protocol [YWL⁺20] which need 1295 GGM trees to generate 10^7 OTs, this requires executing FSS and instantiating GGM for $\approx 2^{26}$ times. By these, once different users have conducted such computations 2^{14} times, the total number of GGM instances approaches 2^{40} .

It could be sufficient to corrupt *just one* of the numerous instances: the breach of any of the Bip32 users already incurs severe loss of funds, and the breach of any of the trees in 2PC already breaks necessary (pseudo)randomness in the protocol. This challenges security *in the multi-user (mu) setting*, in which for a parameter u representing the maximal number of *users*, there are u independently chosen secret keys K_1, \dots, K_u , and the adversary succeeds as long as it compromises at least one out of u instances $F(K_1, \cdot), \dots, F(K_u, \cdot)$ of the target F .

The mu setting was first formalized in public-key cryptography [BBM00]. Asymptotically, it is equivalent to su security: a scheme with κ bits su security trivially ensures $\kappa - \log_2 u$ bits mu security by a standard hybrid argument. But, as discussed, u can be large in practice, and the $\log_2 u$ bits degradation has incurred serious concerns [BT16,GKW⁺20] and even practical Bitcoin cracking project [LBC16] (see Appendix A). Formally exploring the relation between concrete bounds and various parameters (including u) is technically challenging, as made evident in recent works regarding public-key [BBM00,Ber15,KMP16] and symmetric-key schemes [BBT16,BT16,BHT18,GKW⁺20], and results have explicitly influenced real world designs [BT16,Ber15].

Despite the raised mu challenge, formal mu constrained or puncturable PRF security treatments of GGM remain missing, and this undermines security of applications. As discussed, trivial bounds are seldom satisfactory. Recall the example of 2PC using 128-bit keys: the trivial bound indicates 88-bit security⁷

⁷ This is tight: since each key guess falls in the 2^{40} keys with probability $2^{40}/2^{128}$, with 2^{88} computations or key guesses we can succeed to “hit” one of the 2^{40} keys with probability ≈ 1 .

and falls far below the expected 128 bits. For Bip32, recent work [DEF+21] proved 111-bit su security, by which the trivial bound indicates 91-bit mu security as long as 2^{20} Bip32 users are available.

Facing the difficulty in providing dedicated mu security proofs for each application, we seek for unified treatments with full proofs of non-trivial mu bounds, in order to provide a systematic reference for their practical deployments (which is the ultimate goal of cryptography).

1.1 Our Contribution

We provide generalized mu security definitions, a unified model of the GGM tree, and concrete mu security proofs. Improved tree schemes with non-trivial, proven mu bounds can be easily derived from our model, and we showcase on Bip32 and recent FSS protocols. Below we elaborate in detail.

Generalized GGM tree and its mu security. To unify the models in various settings, we extend both the constrained PRF (CPRF) security definition and the GGM scheme. Regarding the former, we start with the fully adaptive constrained PRF security notion of Hofheinz et al. [HKKW19], and extend it along two axes to reach a *mu leakage constrained PRF security* definition in the simulation paradigm:

- (i) *Multi-user:* Hofheinz et al.’s (su) CPRF security definition [HKKW19] allows the adversary to adaptively acquire constrained keys and function values, and requires the function value $F(K, x)$ to be pseudorandom for any x that has never been constrained. Our natural mu extension considers u secret user keys K_1, \dots, K_u , and requires the function values $F(K_i, x), F(K_j, x')$ to be pseudorandom and independent for any two involved user keys K_i, K_j , and for x unconstrained w.r.t. K_i and x' unconstrained w.r.t. K_j ;
- (ii) *Protocol-level Leakages:* In some applications such as the Bip32, the (secret) intermediate values of the F evaluations are used to derive *public* information and incur protocol-level leakages. To formalize these settings, we augment the real world with a (*context-dependent*) *leakage oracle* L outputting leakages of the corresponding intermediate values upon every query, and the ideal world with a *leakage simulator* S outputting faked leakages. The detailed leakages depend on the concrete context, and we refer to Sect. 4.1 for example.

Regarding the scheme, we propose a generalized model for GGM trees. Our model is built upon a *public* cryptographic primitive **Prim** (that can be accessed by the adversary), uses κ -bit internal secrets, and allows for multiple branches. To reflect influences of public parameters, every **Prim**-call in the tree has an additional input which we informally called “label”. Modeling **Prim** as either a fixed-input-length random oracle (FIL RO) **H** (to justify instantiating **Prim** with cryptographic hash functions) or a Davies-Meyer construction $DM^{\mathbf{E}}(L, x) = \mathbf{E}(L, x) \oplus x$ based on an ideal cipher **E** (to justify instantiating

Prim with blockciphers), we analyze the generalized tree w.r.t. our mu leakage CPRF security definition. We prove $\kappa - \log_2 C$ bits mu CPRF security, where C is a parameter depending on the probability among distinct “label” inputs of the internal **Prim** calls *across all the users*. When all “labels” are the same, security becomes (inferior) $\kappa - \log_2 D$ bits, where D , called *effective data complexity*, is the total number of **Prim**-calls internally made by the u tree instances and is related to the number of adversarial queries to the trees. When collisions among “labels” are unlikely, security becomes nearly optimal $\approx \kappa$ bits. Moreover, the leakage of certain intermediate values are indistinguishable from random “simulated leakages”, indicating that protocols can extract (a limited amount of) pseudorandom bits from the intermediate values in the tree and use them in arbitrary.

The random oracle-based trees can be instantiated with truncated KECCAK- p permutations of SHA3 [CLL19], while the results on Davies-Meyer-based trees enable instantiations using the compression function of SHA512 or the AES in Davies-Meyer mode. Note that even if \mathbf{E} is ideal, the Davies-Meyer construction $\text{DM}^{\mathbf{E}}$ cannot be modeled as a random oracle [DRST12], and we thus have to appeal for a dedicated analysis, and the proved bounds differ by a factor of 2.

While our security definition, model and provable bounds appear complicated, our analyses complete a large step of proofs for a wide range of GGM tree variants and shed lights on the influences of parameters. To derive bounds for concrete designs, designers just need to fill in the parameters and make some additional counting for the aforementioned C and D . We believe this could help characterize state-of-the-art designs and provide building blocks for the coming NIST standardization [NIS21]. We will showcase on two applications.

We remark that our bounds are proven in the ideal (function or cipher) model, and should be taken as a heuristic insurance for their practical instantiations. This theoretical caveat is shared by similar works [DFL19, ADE⁺20, DEF⁺21]. Though, the use of ideal model appears necessary to characterize how local computation (approximated by the number of ideal primitive queries) affects security in the mu setting [BT16, BHT18, GKW⁺20]. Meanwhile, as noticed in [ST16], standard model proofs fail to yield “realistic” mu security bounds for many symmetric schemes with rekeying [BHT18, GKW⁺20] (which *is* extensively used in Bip32).⁸

Puncturable PRFs We make a natural step further and consider generalized GGM as a *puncturable pseudorandom function (PPRF)* [KPTZ13, BW13, BGI16]. This is a PRF F such that given an input x and a key K , one can generate a punctured key, denoted $K\{x\}$, which allows evaluating F at every point *except for* x , and does not reveal any information about the function value $F(K, x)$. The notion was subsequently extended to allow for puncturing multiple inputs [HKW15]. This functionality is the “dual” of CPRFs. In this respect, we extend PPRFs

⁸ Though, we believe that assuming **Prim** is a weak PRF, inferior leakage security bounds can be proven. This is an interesting open question.

to the *multi-user and leakage* setting, and establish multi-user leakage PPRF security for generalized GGM using the aforementioned CPRF results.

Mu security of Bip32, and improvements. Bip32 hierarchical deterministic wallet (HDW) specifies a GGM tree-based approach to derive a collection of digital signature keys organized under an access hierarchy [Med18]. We refer to Fig. 1 for an overview and Sect. 6.1 for more details. Prior works on such HDWs typically focused on enhancing functionalities or achieving new security notions (see Related Work below). Despite the importance for practical assurance, (a) the hierarchical security of Bip32 was never formally proved; (b) the concrete bounds were never characterized even in the su setting (modulo a concurrent work [DEF⁺21]: see Related Work below).

Along a parallel though related axis regarding efficiency, the HMAC-SHA512-based Bip32 standard [Med18] consumes a huge number of AND gates and is costly when implemented in the MPC setting (for distributed key management), and improvements were mentioned by Lindell as an open problem [Lin19]. While some ideas appear obvious (e.g., using more efficient hashing instead of HMAC-SHA512), the soundness is unclear due to the lack of formalism and justification.

To address the gap, we consider mu *hierarchical unforgeability* and mu *hierarchical unlinkability* of Bip32: the former guarantee that the collaboration of several “accounts” cannot forge transactions for the other “accounts” (and are thus unable to spend their money), while the latter ensure that in the view of several “accounts”, the public signature keys derived by the other “accounts” are pseudorandom and independent. Then, using our leakage CPRF security of the RO-based trees, we prove that with u users, Bip32 achieves (roughly) $\min\{247, 256 - \log_2 u\}$ bits mu hierarchical unlinkability security and $\min\{247, 256 - \log_2 u, f(q_S)\}$ bits mu hierarchical unforgeability security, where $f(q_S)$ is the mu security of the underlying signature scheme for q_S the number of adversarial signing queries. By these, the concrete (even mu) unlinkability bounds of the whole Bip32 system is mostly close to the expectation of its designers (256 bits), and the degradation with u is limited. The concrete unforgeability bounds, however, depend on $f(q_S)$ the mu security of the signature, and this is unavoidable. This emphasizes on *carefully choosing the signature*. Besides, the relation with CPRFs provide cryptographic insights on HDWs including Bip32.

Our formal analysis opens the way to improving Bip32 tree using SHACAL3, the compression function underlying SHA512, and KECCAK- p [800, 11], a cryptographic permutation from the SHA3 family. Our proposals reduce the number of AND gates by 73.3%~93.8% compared with the standard Bip32, and are promising in the context of threshold cryptography and side-channel protections. To demonstrate, we benchmark MPC implementations of our proposals and the standard Bip32, with results in Table 1 indicating expected improvements.

Mu security of trees in FSS, and improvements. FSS is a cryptographic primitive where the client can secretly share a function f to f_1 and f_2 such that each of the function does not reveal the parameters of f . Two servers holding

f_i and x can locally evaluate $f_i(x)$ and the scheme ensures that $f(x) = f_1(x) \oplus f_2(x)$. When f is a point function, such construction can be easily used for private information retrieval. Furthermore, the splitting of function f can be performed in a two-party computation protocol and thus eliminate the need of a client.

As mentioned before, most such applications require a private puncturable PRF, which is built upon a GGM tree. In detail, in PCG, a client (which could be simulated by MPC) samples a key of a puncturable PRF, which is essentially the root of a GGM tree, and then sends the key/root to one of the party. The other party chooses a leaf (in an oblivious manner) and obtains the punctured key enabling evaluating all but that leaf. A popular way to instantiate the GGM tree is to define the length-doubling PRG as $s \mapsto \text{AES}_{fk}(s) \oplus s \parallel \text{AES}_{fk}(s \oplus 1) \oplus s \oplus 1$ (denote FBTr), where AES_{fk} is the AES using a fixed, publicly known keys fk . However, security characterizations are lacked even in the su setting, leaving gaps in both concrete bounds and design references.

We bridge both gaps. First, regarding security of the aforementioned scheme FBTr, we formally prove (roughly tight) concrete mu puncturable PRF security of $128 - \log_2 d - \log_2 u$ bits, where d is the depth of the tree. These, as discussed, suffer from the (notable) $\log_2 u$ bit degradation.

We then show how to improve. 1) To overcome the $\log_2 u$ degradation, we propose to use *random IV as the fixed AES key in every FSS protocol instance*, and this improves the mu security to $128 - \log_2 d - 2$ bits which is a good bound since d is typically small in practice. 2) We propose multi-branch generalizations that offers a promising tradeoff with the same mu security, i.e., it saves 50% computations at the expense of 50% larger puncturable keys.

While our analysis does not cover the full FSS, it indicates the trees exhibit no practical security issues. We leave the full characterization for future work.

1.2 Related Work

Regarding CPRFs/PPRFs, a series of works overcame standard model challenges (which is complementary to our goal of providing practical reference) and greatly advance the theory [FKPR14,HKW15,BLW17,HKKW19].

Regarding HDW, Das et al. [DFL19,ADE⁺20] appeared the first to formalize security of *deterministic wallets* as *unforgeability* and *unlinkability*. Concurrently to ours, they extend their treatment to Bip32 [DEF⁺21]. Das et al.’s hierarchical unforgeability definition is stronger than ours, as it requires unforgeability of signatures even if the parent chain codes are leaked. This models the settings of “hot wallet breach” and granting chain codes to auditors. Due to this, they have to resort to *rerandomizable signatures* and related-key properties of ECDSA. In contrast, our unforgeability definition assumes parent chain codes secret. Thus our bounds do not cover the “audit” use case [Med18, Use cases]. Moreover, our analysis does not distinguish “hardened” and “non-hardened” derivations (since parent chain codes are always secret). Our model mostly covers security against a subset of malicious offices in an enterprise. On the other hand, our treatments provided non-trivial multi-user security bounds. Thanks

to the fine-grained analysis using multi-collisions, our mu unlinkability bound is much better than [DEF⁺21, Theorem A.1]. In all, our analysis of Bip32 offers another complementary viewpoint. Perhaps more importantly, our main focus is in the generalized GGM tree construction and its results, with Bip32 a mere application.

Luzio et al. provided a systematic study of HDW and a new design Arcula [LFA20]. Their work emphasized more on offering a new proposal with richer functions than characterizing security of existing schemes (i.e., Bip32). In comparison, we emphasize more on schemes already used “in the wild”. For a survey of earlier works [TVR16,AGKK19,MPs19], we refer to [DFL19, Sect. 1.3].

1.3 Organization

We establish notations and models in Sect. 2. Our extensions of constrained and puncturable PRF notions are also in Sect. 2. Then, we formally define our general GGM tree model in Sect. 3, and prove the mu leakage CPRF security in Sect. 4. The mu leakage PPRF security is established in Sect. 5. We finally demonstrate applications to Bip32 and FSS in Sect. 6 and Sect. 7 respectively.

2 Preliminaries

Denote by $[i]_m$ the m -bit binary encoding of the non-negative integer i , and by \perp the empty string. Given an n -bit string x and $a \leq n$, denote by $\text{left}_a(x)$ (resp., $\text{right}_a(x)$) the a leftmost (resp., rightmost) bits of x .

To describe Bip32 wallets in Sect. 3 and 6, denote by \mathbb{G} the group of the elliptic curve in use and G its primitive element. Denote by $\mathbb{Z}_{|\mathbb{G}|}^+$ the set $\{1, \dots, |\mathbb{G}|-1\}$. The function $\text{int}(X)$ interprets a 256-bit string X as a 256-bit number. The function $\text{ser}_P(sk \cdot G) = \text{ser}_P(pk)$ serializes the coordinate $pk = (e_x, e_y)$ as a bit string in the SEC1’s compressed form, i.e., $\text{ser}_P(e_x, e_y) = [2]_8 \parallel [e_x]_{256}$ when $e_y = 0 \pmod 2$, while $\text{ser}_P(e_x, e_y) = [3]_8 \parallel [e_x]_{256}$ when $e_y = 1 \pmod 2$.

Signature schemes. Due to Bip32, we need the definition and security of digital signature schemes.

Definition 1 (Signature Scheme). A signature scheme $\text{Sig} = (\text{KGen}, \text{Sign}, \text{Vrfy})$ is a triple of algorithms. The randomized key generation algorithm KGen takes as input public parameters pp and returns a pair $(pk, sk) \in \mathcal{K}_{\text{sign}}$ of public and secret keys. The randomized signing algorithm Sign takes as input a secret key sk and a message m and returns a signature σ . The deterministic verification algorithm Vrfy takes as input a public key pk , a signature σ , and a message m . It returns 1 (accept) or 0 (reject). We require correctness, i.e., $\forall (pk, sk) \leftarrow \text{KGen}(\text{pp}) \forall m : \text{Vrfy}(pk, \text{Sign}(sk, m), m) = 1$.

We will reduce the security of the Bip32 wallet to the standard *existential unforgeability under chosen message attacks* (UFCMA) security of the signatures in the multi-user setting. For this, we adopt the formalism of Bernstein [Ber15].

Definition 2 (Multi-user UFCMA Security). A (digital) signature scheme $\text{Sig} = (\text{KGen}, \text{Sign}, \text{Vrfy})$ is (u, q_S, t, ε) -muUFCMA secure, if for any adversary \mathcal{A} making q_S queries to $\text{Sign}(sk_1, \cdot), \dots, \text{Sign}(sk_u, \cdot)$ and running in time t , it holds $\Pr[\mathcal{A}^{\text{Sign}(sk_1, \cdot), \dots, \text{Sign}(sk_u, \cdot)}(pk_1, \dots, pk_u) \text{ forges}] \leq \varepsilon$, where the event “forges” means \mathcal{A} outputs a pair (m^*, σ^*) such that $\text{Vrfy}(pk_i, m^*, \sigma^*) = 1$ for some $i \in \{1, \dots, u\}$ and \mathcal{A} never queried $\text{Sign}(sk_{i'}, m^*)$ for any i' with $sk_{i'} = sk_i$.

A naive hybrid argument establishes $\kappa - \log_2 u$ muUFCMA security from κ bit UFCMA security. For the Schnorr and BLS schemes, this multi-user security loss can be overcome in the random oracle model [KMP16, Lac18]—for example, using 256-bit secret keys, the muUFCMA security of Schnorr is of 128 bits according to [KMP16],—or by using the key-prefixing technique [Ber15, Lac18].

Constrained PRF and its Multi-user security. Our formalism of constrained PRFs (CPRFs) basically follows [BW13, KPTZ13, BGI14], which is as follows.

Definition 3 (Constrained PRF). With key space \mathcal{K} , domain \mathcal{X} , and range \mathcal{Y} , a constrained pseudorandom function for a set system $\mathbb{S} \subseteq 2^{\mathcal{X}}$ is a keyed function F with an additional constrained key space \mathcal{K}_c and four probabilistic polynomial-time algorithms $(F.\text{KGen}, F.\text{Ev}, F.\text{Co}, F.\text{SubCo}, F.\text{CEv})$:

- the key generation algorithm $F.\text{KGen}$ returns $K = (k, \text{pp}) \in \mathcal{K}$, where k is the ordinary secret master key and pp is the public parameter;
- the (ordinary) evaluation algorithm $F.\text{Ev}(K, x)$ (always) outputs $F(K, x)$ for the inputs $K \in \mathcal{K}$ and $x \in \mathcal{X}$;
- the constraining algorithm $F.\text{Co}(K, \mathcal{S})$ outputs a constrained key $K\{\mathcal{S}\} \in \mathcal{K}_c$ on input a key $K \in \mathcal{K}$ and a set $\mathcal{S} \in \mathbb{S}$;
- the constrained evaluation algorithm $F.\text{CEv}(K\{\mathcal{S}\}, x)$: on input $K\{\mathcal{S}\}$ constraining all points in \mathcal{S} and x , outputs $F(K, x)$ if $x \in \mathcal{S}$, and \perp otherwise.

We will focus on CPRFs built upon a public ideal primitive **Prim**, and write F^{Prim} to highlight. For a CPRF F^{Prim} , we follow [HKKW19] and formalize the adversarial goal as distinguishing the real world oracles $(F^{\text{Prim}}.\text{Co}, F^{\text{Prim}}.\text{Ev}, \text{Prim})$ from the ideal world oracles $(F^{\text{Prim}}.\text{Co}, R, \text{Prim})$ for a random function R . We extend it with *multiple users* and *leakages*, in the *concrete security paradigm*. Before the formal presentation, we first elaborate on the new ingredients.

Multi-user. Let u be the maximal number of users.⁹ In the mu setting, the adversarial goal becomes distinguishing the real world $(F^{\text{Prim}}.\text{Co}(K_1, \cdot), F^{\text{Prim}}.\text{Ev}(K_1, \cdot), \dots, F^{\text{Prim}}.\text{Co}(K_u, \cdot), F^{\text{Prim}}.\text{Ev}(K_u, \cdot))$ and the ideal world $(F^{\text{Prim}}.\text{Co}(K_1, \cdot), R(1, \cdot), \dots, F^{\text{Prim}}.\text{Co}(K_u, \cdot), R(u, \cdot))$ for u independent random keys $\mathbf{K} = (K_1, \dots, K_u) = ((k_1, \text{pp}_1), \dots, (k_u, \text{pp}_u))$, where $R(1, \cdot), \dots, R(u, \cdot)$ instantiate u independent random functions. To simplify notations, we will use a single oracle $\text{muEv}_{\mathbf{K}}(i, \cdot)$ for the functionality of $F^{\text{Prim}}.\text{Ev}(K_i, \cdot)$, and a single $\text{muCo}_{\mathbf{K}}(i, \cdot)$ for $F^{\text{Prim}}.\text{Co}(K_i, \cdot)$.

⁹ This is equivalent to the alternative definitions allowing adaptively adding new instances [BBT16, BT16].

Protocol-level leakages. As mentioned, to capture the cases where (secret) intermediate values become public information, we augment the real world oracles $\text{muEv}_{\mathbf{K}}$ and $\text{muCo}_{\mathbf{K}}$ with a *leakage oracle* L . Every time the distinguisher issues a query to either $\text{muEv}_{\mathbf{K}}$ or $\text{muCo}_{\mathbf{K}}$, besides the ordinary response, it obtains leakages due to the corresponding internal computations from L . To highlight, we use the notations $\text{muCo}_{\mathbf{K}}^{\mathsf{L}}$ and $\text{muEv}_{\mathbf{K}}^{\mathsf{L}}$ for these “leaky oracles”.

For the ideal world oracles $\text{muCo}_{\mathbf{K}}$ and R to produce consistent outputs, we augment them with a *leakage simulator* S and obtain the “leaky ideal oracles” $\text{muCo}_{\mathbf{K}}^{\mathsf{S}}$ and R^{S} . Every time the distinguisher issues a query, it obtains both ordinary responses and leakages produced by S . This enables formalizing indistinguishability of the leaky real and ideal worlds, with L and S being parameters. The detailed definitions of L and S depend on the concrete applications and security requirements, and we refer to Sect. 4.1 for example.

While our simulation-based definition seems quite strong, we have provided positive results (Theorems 1 and 2) with *explicit simulators*, and our simulators simply output random leakages, indicating pseudorandomness of (leaked) intermediate values. This is in contrast with simulatability of *side-channel* leakages, the achievability of which remains open [LMO+14].

Effective data complexity. Concrete security of a cryptosystem is qualified by the attack advantage regarding adversaries with data and time complexities. In the ideal model, time complexity is typically captured by T the number of queries to the ideal function \mathbf{Prim} . Data complexity shall reflect the amount of information gained from the oracles $\text{muCo}_{\mathbf{K}}^{\mathsf{L}}$ and $\text{muEv}_{\mathbf{K}}^{\mathsf{L}}$, which is a complicated function of the true number of adversarial queries. To remedy, we follow [DRST12] and introduce *effective data complexity* D , which is the total number of queries to \mathbf{Prim} internally made by $\mathsf{F}^{\mathbf{Prim}}$ during the interaction. As will be seen in Sect. 6 and 7, it is easy to count D for concrete applications. The output of every internal \mathbf{Prim} query, including the ordinary output and the corresponding leakage, will be a part of the information collected by the distinguisher. Therefore, effective data complexity does measure the amount of information gained from $\text{muCo}_{\mathbf{K}}^{\mathsf{L}}$ and $\text{muEv}_{\mathbf{K}}^{\mathsf{L}}$.

Formal definition. With the above, our formal definition is as follows.

Definition 4 (Multi-user Leakage CPRF Security). *The keyed function $\mathsf{F}^{\mathbf{Prim}}$ is a (u, T, D, ε) - (L, S) -constrained PRF, if for any distinguisher \mathcal{D} making T queries to \mathbf{Prim} and having effective data complexity D , we have*

$$\left| \Pr[\mathcal{D}^{\text{muCo}_{\mathbf{K}}^{\mathsf{L}}, \text{muEv}_{\mathbf{K}}^{\mathsf{L}}, \mathbf{Prim}} = 1] - \Pr[\mathcal{D}^{\text{muCo}_{\mathbf{K}}^{\mathsf{S}}, \mathsf{R}^{\mathsf{S}}, \mathbf{Prim}} = 1] \right| \leq \varepsilon,$$

where the probability is taken over the u user keys $\mathbf{K} = (K_1, \dots, K_u)$ with $K_i \leftarrow \text{KGen}$, over \mathcal{D} 's random tape and the ideal primitive \mathbf{Prim} , and where:

- (i) $\text{muCo}_{\mathbf{K}}^{\mathsf{L}}(i, \mathcal{S})$: for $1 \leq i \leq u$, outputs the constrained key $\mathsf{F}^{\mathbf{Prim}}.\text{Co}(K_i, \mathcal{S})$ and the corresponding information leakage $\mathsf{L}(K_i, \mathcal{S})$ for the i -th user;

- (ii) $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}(i, x)$: for $1 \leq i \leq u$, outputs $\mathbf{F}^{\mathbf{Prim}}.\text{Ev}(K_i, x)$ and the corresponding information leakage $\mathbf{L}(K_i, x)$ for the i -th user;
- (iii) $\text{muCo}_{\mathbf{K}}^{\mathbf{S}}(i, \mathcal{S})$: for $1 \leq i \leq u$, outputs the constrained key $\mathbf{F}^{\mathbf{Prim}}.\text{Co}(K_i, \mathcal{S})$ and the corresponding simulated leakage $\mathbf{S}(i, \mathcal{S})$ for the i -th user;¹⁰
- (iv) $\mathbf{R}^{\mathbf{S}}(i, x)$: for $1 \leq i \leq u$, outputs $y \xleftarrow{\mathbf{S}} \mathcal{Y}$ for every new pair of inputs (i, x) and the corresponding simulated leakage $\mathbf{S}(i, x)$.

\mathcal{D} is not allowed to make a constraining query $\text{muCo}_{\mathbf{K}}^{\mathbf{L}}(i, \mathcal{S})$ and an evaluation query $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}(i, x)$ with $x \in \mathcal{S}$, in any order, since this leads to trivial win.

The classical CPRF security notion can be recovered by eliminating the leakages \mathbf{L} and \mathbf{S} .

Puncturable PRF and its Multi-user security. Depending on its structure, a PPRF scheme may impose certain relations among the multiple punctured inputs. To capture this, we follow Definition 3 and define PPRF for set systems.

Definition 5 (Puncturable PRF). With key space \mathcal{K} , domain \mathcal{X} , and range \mathcal{Y} , a puncturable pseudorandom function for a set system $\mathbb{S} \subseteq 2^{\mathcal{X}}$ is a keyed function \mathbf{F} with an additional punctured key space \mathcal{K}_p and four probabilistic polynomial-time algorithms $(\mathbf{F}.\text{KGen}, \mathbf{F}.\text{Ev}, \mathbf{F}.\text{Pu}, \mathbf{F}.\text{PEv})$:

- the key generation algorithm $\mathbf{F}.\text{KGen}$ returns $K = (\mathbf{k}, \mathbf{pp}) \in \mathcal{K}$, where \mathbf{k} is the ordinary secret master key and \mathbf{pp} is the public parameter;
- the evaluation algorithm $\mathbf{F}.\text{Ev}(K, x)$ outputs $\mathbf{F}(K, x)$ for the inputs $K \in \mathcal{K}$ and $x \in \mathcal{X}$;
- the puncturing algorithm $\mathbf{F}.\text{Pu}(K, \mathcal{S})$ outputs a punctured key $K\{\mathcal{S}\} \in \mathcal{K}_p$ on input a key $K \in \mathcal{K}$ and a set $\mathcal{S} \in \mathbb{S}$;
- the punctured evaluation algorithm $\mathbf{F}.\text{PEv}(K\{\mathcal{S}\}, x)$: on input $K\{\mathcal{S}\}$ punctured at all points in \mathcal{S} and x , outputs $\mathbf{F}(K, x)$ if $x \notin \mathcal{S}$, and \perp otherwise.

For a secure PPRF $\mathbf{F}^{\mathbf{Prim}}$, $\mathbf{F}^{\mathbf{Prim}}.\text{Ev}(K, x)$ shall be pseudorandom for any $x \in \mathcal{S}$, even if the punctured key $K\{\mathcal{S}\}$ has been given. This is just the dual of a restricted form of CPRF. For a formal definition in the mu setting, we borrow the ingredients from our Definition 4.

Definition 6 (Multi-user Leakage PPRF Security). $\mathbf{F}^{\mathbf{Prim}}$ is a (u, T, D, ε) - (\mathbf{L}, \mathbf{S}) -puncturable PRF, if for any distinguisher \mathcal{D} making T queries to \mathbf{Prim} and having effective data complexity D , we have

$$\left| \Pr[\mathcal{D}^{\text{muPu}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{Prim}} = 1] - \Pr[\mathcal{D}^{\text{muPu}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{Prim}} = 1] \right| \leq \varepsilon,$$

where the probability is taken over the u user keys $\mathbf{K} = (K_1, \dots, K_u)$, with $K_i \leftarrow \text{KGen}$, over \mathcal{D} 's random tape and the ideal primitive \mathbf{Prim} , and where:

¹⁰ This paper focuses on the setting where \mathbf{S} does not have access to the \mathbf{Prim} oracle, which suffices for producing random leakages.

- (i) $\text{muPu}_{\mathbf{K}}^{\mathbf{L}}(i, \mathcal{S})$: for $1 \leq i \leq u$, outputs the punctured key $\text{F}^{\text{Prim}}.\text{Pu}(K_i, \mathcal{S})$ and the corresponding information leakage $\mathbf{L}(K_i, \mathcal{S})$ for the i -th user;
- (ii) $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}(i, x)$: for $1 \leq i \leq u$, outputs $\text{F}^{\text{Prim}}.\text{Ev}(K_i, x)$ and the leakage $\mathbf{L}(K_i, x)$ for the i -th user;
- (iii) $\text{muPu}_{\mathbf{K}}^{\mathbf{S}}(i, \mathcal{S})$: for $1 \leq i \leq u$, outputs the punctured key $\text{F}^{\text{Prim}}.\text{Pu}(K_i, \mathcal{S})$ and the corresponding simulated leakage $\mathbf{S}(i, \mathcal{S})$ for the i -th user;
- (iv) $\mathbf{R}^{\mathbf{S}}(i, x)$: for $1 \leq i \leq u$, outputs $y \xleftarrow{\mathbf{S}} \mathcal{Y}$ and simulated leakage $\mathbf{S}(i, x)$ for every new pair (i, x) .

The following two restrictions are imposed on the distinguisher's queries:

- (a) \mathfrak{D} is not allowed to make a puncture query $\text{muPu}_{\mathbf{K}}^{\mathbf{L}}(i, \mathcal{S})$ if a previous evaluation query $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}(i, x)$ with $x \notin \mathcal{S}$ is made;
- (b) \mathfrak{D} is not allowed to make two distinct puncturing queries $\text{muPu}_{\mathbf{K}}^{\mathbf{L}}(i, \mathcal{S})$ and $\text{muPu}_{\mathbf{K}}^{\mathbf{L}}(i', \mathcal{S}')$ with $i = i'$. Namely, each user is punctured (at most) once.

Restriction (a) is necessary to prevent trivial wins (i.e., query both $\text{muPu}_{\mathbf{K}}^{\mathbf{L}}(i, \mathcal{S})$ and $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}(i, x)$ with $x \notin \mathcal{S}$ to check if the latter matches the result computed from the former). Restriction (b) stems from the nature of PPRFs: note that as the interaction proceeds, the number of $\text{F.Ev}(K, x)$ that can be computed by \mathfrak{D} can not decrease. Thus, if \mathfrak{D} makes two puncturing queries $\text{muPu}_{\mathbf{K}}^{\mathbf{L}}(i, \mathcal{S})$ and $\text{muPu}_{\mathbf{K}}^{\mathbf{L}}(i, \mathcal{S}')$ to the same i -th user/instance, then the later query \mathcal{S}' has to be a subset of \mathcal{S} , otherwise the later puncturing request cannot succeed. But it's odd for \mathfrak{D} to make such two puncturing queries. Therefore, we simply forbid \mathfrak{D} puncturing the same user/instance twice.

H-coefficient method. We use Patarin's H-coefficient method [Pat09] to prove security of the trees, and provide a quick overview here. Our presentation borrows heavily from that of [CS14]. Fix a distinguisher \mathfrak{D} that makes at most q queries to its oracles. As in the security definition presented above, \mathfrak{D} 's aim is to distinguish between two worlds: a "real world" and an "ideal world". Assume wlog that \mathfrak{D} is deterministic. The execution of \mathfrak{D} defines a *transcript* that includes the sequence of queries and answers received from its oracles; \mathfrak{D} 's output is a deterministic function of its transcript. Thus, if $T_{\text{re}}, T_{\text{id}}$ denote the probability distributions on transcripts induced by the real and ideal worlds, respectively, then \mathfrak{D} 's distinguishing advantage is upper bounded by the statistical distance $\text{SD}(T_{\text{re}}, T_{\text{id}}) := \frac{1}{2} \sum_{\mathcal{Q}} |\Pr[T_{\text{re}} = \mathcal{Q}] - \Pr[T_{\text{id}} = \mathcal{Q}]|$, with sum taken over all possible transcripts \mathcal{Q} .

Let Θ denote the set of all transcripts that can be generated by \mathfrak{D} in either world. We look for a partition of Θ into two sets Θ_{good} and Θ_{bad} of "good" and "bad" transcripts, respectively, along with a constant $\varepsilon_1 \in [0, 1)$ such that

$$\mathcal{Q} \in \Theta_{\text{good}} \implies \frac{\Pr[T_{\text{re}} = \mathcal{Q}]}{\Pr[T_{\text{id}} = \mathcal{Q}]} \geq 1 - \varepsilon_1. \quad (1)$$

It is then possible to show (see [CS14]) that the statistical distance is upper bounded by

$$\text{SD}(T_{\text{re}}, T_{\text{id}}) \leq \varepsilon_1 + \Pr[T_{\text{id}} \in \Theta_{\text{bad}}]. \quad (2)$$

3 A Framework for Generalized GGM Trees

Before presenting our model in Sect. 3.2, we first serve intuitions in Sect. 3.1 to ease understanding.

3.1 Intuitions

Example 1. The classical GGM tree uses a length-doubling PRG $G : \{0, 1\}^n \rightarrow \{0, 1\}^{2n}$. Thus, every G invocation expands an n -bit intermediate value s into two children, and this constitutes a binary tree. An instantiation of [GKWY20] defines $G^{\text{AES}_{fk}}(s) := \text{DM}^{\text{AES}}(fk, s) \parallel \text{DM}^{\text{AES}}(fk, s \oplus [1]_{128})$ with $n = 128$, where $\text{DM}^{\text{AES}}(x, y) := \text{AES}_x(y) \oplus y$ and AES_{fk} is the AES using a fixed, publicly known keys fk (though, distinct high-level protocols may use distinct fk). The security of $G^{\text{AES}_{fk}}$ is only justifiable by assuming AES is an ideal cipher. A natural extension is to increase “parallelization degree”, i.e., mapping s to θn bits $\text{DM}^{\text{AES}}(fk, s) \parallel \dots \parallel \text{DM}^{\text{AES}}(fk, s \oplus [\theta - 1]_{128})$ for some integer $\theta \geq 3$. If we view AES as a *public primitive*, then every internal secret s is involved in $\theta \geq 2$ distinct primitive-calls.

Example 2. The Bip32 wallet defines a more sophisticated approach to generate a collection of keys organized in a *key tree* [Med18]. The tree has multiple branches. The 256-bit “chain codes” in [Med18] (the value ch_p in the dashed box in Fig. 1) essentially constitute the internal secret states of the tree. For each such state, the key tree makes $\theta \leq 2^{32}$ calls to HMAC-SHA512 (henceforth abbreviated as HMAC) to derive θ children. Every HMAC-call has additional complicated inputs, including an index j and a signature key. A half of the HMAC-output (the I values in Fig. 1) will be used to derive the child signature (private and public) key, meaning that it may not be perfectly secret anymore. The default key tree recommended in [Med18] has depth 3, as shown in Fig. 1. If we replace HMAC by a general double-input function **Prim** that accept two inputs of ν and λ bits, then the obtained key tree is depicted in Fig. 1.

Summary. Inspired by the above examples, we would like to have a unified model compatible with flexible choices as follows.

First, we build our model $\text{GGGM}^{\text{Prim}}$ on a public primitive **Prim** functioning as the sub-tree derivation function. In many scenarios the output size of AES may be insufficient (since it may limit the security bound), and “larger” primitives such as SHA512 and SHA3 may be preferred. In this respect, our results should address both the case that **Prim** is a keyed random oracle, i.e., a function that maps each (key,message) pair to an independent and uniform point [DRST12], and the case that **Prim** is the the aforementioned Davies-Meyer mode of an ideal cipher.

Second, to increase design choices, we prefer that all parameters are flexible, including: the size n of the node in the tree, the size κ of the secret which can be viewed as security parameter and “parallelization degree” θ . Moreover, when the output size of **Prim** is long, a single **Prim** invocation may give rise to multiple child nodes, and we denote by w the number of such children.

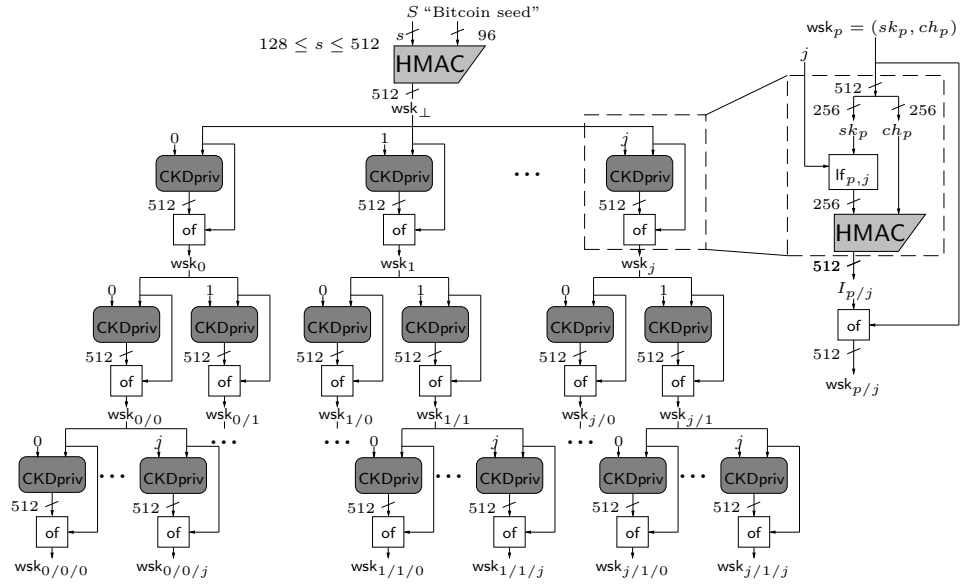


Fig. 1. The key tree in default configuration of Bip32. Due to space, we omit some nodes. S is the seed of the HDW instance, and "Bitcoin seed" is the 96-bit literal string. The values $wsk_{\perp}, wsk_0, \dots$ are wallet secret keys of Bip32 and will be formalized in Sect. 6.2 and 6.3. The dashed box in the corner shows the internal computations of the function CKDpriv (which will be formally defined in Fig. 6). The functions lf and of will be defined in Eqs. (25) and (26).

Finally, to somewhat separate distinct **Prim** invocations, one may inject “labels” into the inputs to **Prim**; one may also “disturb” the node input by e.g., xoring constants, as the aforementioned use of s and $s \oplus [1]_{128}$ in the two parallel calls [GKWY20]. We model these ideas as two input mappings sf and lf . We will elaborate on the details in the next subsection.

3.2 Our General Tree Model

By the above intuitions, our general model $\text{GGGM}^{\text{Prim}}$ is built upon a public primitive $\text{Prim} : \{0, 1\}^\nu \times \{0, 1\}^\lambda \mapsto \{0, 1\}^{wn}$ functioning as the sub-tree derivation function, and is formally described in Fig. 2. Also see Fig. 3 for depiction.

In $\text{GGGM}^{\text{Prim}}$, every node represents a string of n bits. Every “internal” node is involved in θ calls to **Prim**, and every **Prim**-call gives rise to w new n -bit node values. Hence, the tree is $w\theta$ -branched. Inspired by [Med18], we refer to positions in a tree using “paths” of the form $p = i_1/i_2/\dots/i_j$, referring to a node at depth j . A node value at “position” $p = i_1/i_2/\dots/i_j$ is denoted $\text{Nd}(p)$, and $\text{Nd}(\perp) = \text{k}$ denotes the root. By this, when the depth of the tree is d , the domain of $\text{GGGM}^{\text{Prim}}.\text{Ev}$ is $\mathcal{X} = \{i_1/i_2/\dots/i_d\}_{i_\ell \in \{0, \dots, w\theta - 1\} \text{ for } \ell = 1, \dots, d}$. We further define

- $\mathcal{P}^* := \{\perp\} \cup \{i_1/i_2/\dots/i_{d'}\}_{d' \in \{1, \dots, d-1\}, i_\ell \in \{0, \dots, w\theta - 1\} \text{ for } \ell = 1, \dots, d'}$ denoting the set of “incomplete” paths in such trees;
- $\mathcal{P} := \mathcal{X} \cup \mathcal{P}^*$ denoting the set of all valid paths in such trees;
- For any two paths $p, p' \in \mathcal{P}$, p is prefix of p' , if $p = p'$, or if there exist i'_1, \dots, i'_ℓ such that $p' = p/i'_1/\dots/i'_\ell$.

Correspondingly, the prefix set system is

$$\mathbb{S}_{\text{pre,GGGM}} = \{\mathcal{S}_{p^*, \text{GGGM}} : p^* \in \mathcal{P}^*\}, \text{ with } \mathcal{S}_{p^*, \text{GGGM}} = \{p \in \mathcal{X} : p^* \text{ is prefix of } p\}. \quad (3)$$

Since every set $\mathcal{S}_{p^*, \text{GGGM}} \in \mathbb{S}_{\text{pre,GGGM}}$ has an associated path p^* , in Fig. 2 we adopt the formalism of $\text{GGGM}^{\text{Prim}}.\text{CEv}$ taking a path p^* as the second input.

For a node value $z = \text{Nd}(p)$, $p \in \mathcal{P}^*$, and an index $j \in \{0, 1, \dots, \theta - 1\}$, the j -th derivation call of $\text{Nd}(p)$ is $\text{Prim}(\text{lf}_{p,j}(\text{left}_{n-\kappa}(z), \text{pp}), \text{sf}_{p,j}(\text{right}_\kappa(z)))$ (as shown in StepDown, Fig. 2), where the two input map functions sf and lf are such that:

- The *labeling function* $\text{lf}_{p,j}(\text{left}_{n-\kappa}(z), \text{pp})$ maps the left most $n - \kappa$ bits of a node value z and the public parameter pp to the ν -bit 1st input for **Prim**. Subsequently, we call this ν -bit value the *label* of the derivation call.
- The *seeding function* $\text{sf}_{p,j}(\text{right}_\kappa(z))$ is injective, and maps the right most κ bits of a node value z to the λ -bit 2nd input for **Prim**. In some sense, $\text{sf}_{p,j}(\text{right}_\kappa(z))$ serves as the “secret seed” of the sub-tree, and κ is *security parameter* of the tree;

The output of the **Prim**-call $I = \text{Prim}(\text{lf}_{p,j}(\text{left}_{n-\kappa}(z), \text{pp}), \text{sf}_{p,j}(\text{right}_\kappa(z)))$ does not immediately give rise to children nodes. Instead, an *outputting function* of $(z, \text{left}_n(\text{right}_{bn}(I)))$ maps the parent node $\text{Nd}(p)$ and the b -th chunk $\text{left}_n(\text{right}_{bn}(I))$ to the child node $\text{Nd}(p/b)$.

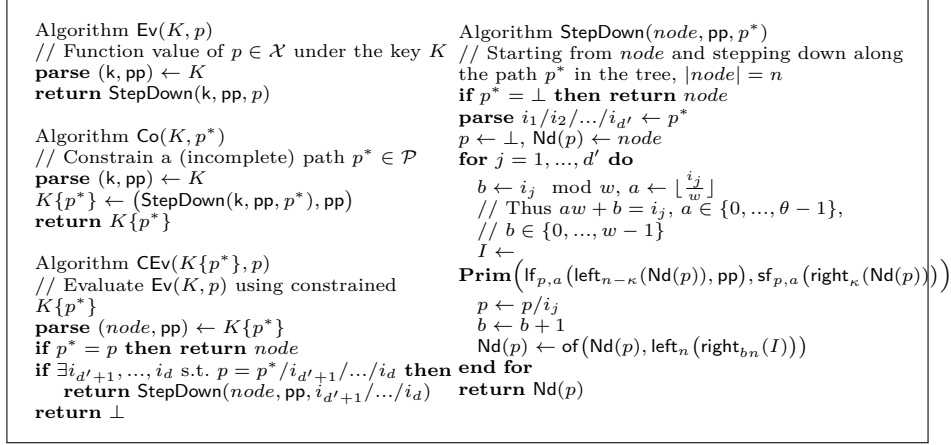


Fig. 2. Constrained PRF based on the generalized GGM tree $\text{GGGM}^{\text{Prim}}[\kappa, n, \nu, \theta, w, d, \text{sf}, \text{lf}]$. For simplicity, below we omit the suffix $[\kappa, \dots, \text{lf}]$.

For the input functions, we require that $j \neq j' \Rightarrow (\text{lf}_{p,j}(\text{left}_{n-\kappa}(z), pp), \text{sf}_{p,j}(\text{right}_{\kappa}(z))) \neq (\text{lf}_{p,j'}(\text{left}_{n-\kappa}(z), pp), \text{sf}_{p,j'}(\text{right}_{\kappa}(z)))$ for any $p \in \mathcal{P}^*$, i.e., distinct **Prim**-calls using the same node value $z \in \{0, 1\}^n$ are necessarily on distinct inputs. For the outputting function, we limit our discussion to *seeded bijections*, i.e., $\text{of}(z, \cdot)$ is a bijection on $\{0, 1\}^n$ for *any* $z \in \{0, 1\}^n$. This means of has an inverse of^{-1} such that $\text{of}^{-1}(z, \text{of}(z, I)) = I$ for any $z, I \in \{0, 1\}^n$.

The complexity of our model is worthy, and it has covered the intuitive examples. We refer to Sect. 6 and 7 for details.




Fig. 3. Derivation calls relevant to a single node $Nd(p)$ in the tree $\text{GGGM}^{\text{Prim}}$.

4 Multi-user Leakage CPRF Security of GGM Trees

For clarity, below in Sect. 4.1 we first cast our general formalism in Sect. 2 into the concrete setting of $\text{GGGM}^{\text{Prim}}$. We also elaborate on the information leakages to-be-considered in this paper, as well as useful structural properties.

```

Algorithm  $S(i_0, p^*)$ 
 $leakages \leftarrow \emptyset$ 
parse  $i_1/i_2/\dots/i_{d'} \leftarrow p^*$ 
for  $\ell = 0, \dots, d' - 1$  do
    if  $Table(i_0, p^*) = \perp$  then
         $r \xleftarrow{\mathcal{S}} \{0, 1\}^{n-\kappa}$ ,  $Table(i_0, i_1^*/i_2^*/\dots/i_\ell^*) \leftarrow r$ 
         $leakages \leftarrow leakages \cup \{(i_0/i_1^*/i_2^*/\dots/i_\ell^*, Table(i_0, i_1^*/i_2^*/\dots/i_\ell^*))\}$ 
    end for
return  $leakages$ 

```

Fig. 4. Leakage simulator S for Theorems 1, 2, and 3.

With these preparations, we consider $\text{GGGM}^{\text{Prim}}$ with Prim being a FIL RO \mathbf{H} in Sect. 4.2, and then the case of Prim being the Davies-Meyer construction $\text{DM}^{\mathbf{E}}$ for an ideal cipher \mathbf{E} in Sect. 4.3.

4.1 Concrete Settings

The oracles and leakage simulator. We assume the leakage oracle \mathbf{L} in $\text{muCo}_{\mathbf{K}}^{\mathbf{L}}$ and $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}$ leaks the leftmost bits $\text{left}_{n-\kappa}(z)$ for every intermediate node $z \in \{0, 1\}^n$ appeared during the computations. In detail, note that by our convention, a query to the evaluation oracle $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}$ is of the form (i_0, p) , $p = i_1/i_2/\dots/i_d$. In the real world, $\text{muEv}_{\mathbf{K}}$ will return $\text{Nd}(i_0, p)$ the n -bit node value at the end of the path p in the i_0 -th tree. As the corresponding leakages, we assume that $\mathbf{L}(K_{i_0}, p)$ will provide d additional values of $n - \kappa$ bits, i.e.,

$$\text{left}_{n-\kappa}(\text{Nd}(i_0, \perp)), \text{left}_{n-\kappa}(\text{Nd}(i_0, i_1)), \dots, \text{left}_{n-\kappa}(\text{Nd}(i_0, i_1/\dots/i_{d-1})).$$

To wit, after issuing a query to $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}$, \mathfrak{D} obtains $n + d(n - \kappa)$ bits information.

Similarly, a query to the constraining oracle $\text{muCo}_{\mathbf{K}}^{\mathbf{L}}$ is of the form (i_0, p^*) , $p^* = i_1^*/\dots/i_{d'}^*$ (see Fig. 2) which may be incomplete (i.e., $d' \leq d$), and $\text{muCo}_{\mathbf{K}}$ will return the n -bit node value $\text{Nd}(i_0, p^*)$. As the leakages, we assume that $\mathbf{L}(K_{i_0}, p^*)$ provides d' values $\text{left}_{n-\kappa}(\text{Nd}(i_0, \perp)), \text{left}_{n-\kappa}(\text{Nd}(i_0, i_1^*)), \dots, \text{left}_{n-\kappa}(\text{Nd}(i_0, i_1^*/\dots/i_{d'-1}^*))$, which resembles the evaluation queries. This means a single query to $\text{muCo}_{\mathbf{K}}^{\mathbf{L}}(i_0, i_1^*/\dots/i_{d'}^*)$ gives rise to $n + d'(n - \kappa)$ bits.

As in Definition 4, the ideal oracle \mathbf{R} returns an n -bit random value which is of the same size as $\text{Nd}(i_0, p)$, and we consider a simulator S that simply outputs random simulated leakages. Formally, S is described in Fig. 4. Consequently, \mathfrak{D} obtains the same amount of random bits (i.e., $n + d'(n - \kappa)$ bits) as in the real world. We stress that, while our security definition is simulation-based, our leakage simulator never “hijacks” adversarial random oracle queries. Therefore, our subsequent results only need *non-programmable random oracles*.

Query restrictions. Besides the query restriction imposed in Definition 4, we additionally assume that the distinguisher \mathfrak{D} never makes redundant queries. Clearly, this cannot decrease attack advantage. For $\text{GGGM}^{\text{Prim}}$, these indicate:

- a. \mathfrak{D} never queries both $\text{muEv}_{\mathbf{K}}^{\mathbf{L}}(i_0, p)$ and $\text{muCo}_{\mathbf{K}}^{\mathbf{L}}(i_0, p^*)$ such that p^* is prefix of p (otherwise $p \in \mathcal{S}_{p^*, \text{GGGM}}$).

- b. \mathfrak{D} never makes distinct queries $(i_0, p^*), (i_0, p^{**})$ to muCoK^L such that p^* is prefix of p^{**} (otherwise (i_0, p^{**}) is redundant);

Brother paths. Given a path $p = i_1/\dots/i_{d'-1}/i_{d'} \in \mathcal{P}$, we define a set $\mathcal{B}r(p) = \{p_1, p_2, \dots, p_w\}$ of w paths, where for $1 \leq \ell \leq w$,

$$p_\ell = i_1/\dots/i_{d'-1}/i_{d'}^{(\ell)}, \quad i_{d'}^{(\ell)} = \lfloor \frac{i_{d'}}{w} \rfloor \cdot w + \ell - 1.$$

These w nodes $\text{Nd}(i_0, p_1), \dots, \text{Nd}(i_0, p_w)$ are attributed to the same wn -bit output of a single call to **Prim** in the i_0 -th tree. We thus call them *brother paths (of p)*. This notion will be used in subsequent analyses.

Concentration of labels. In the random oracle model, distinct labels separate corresponding **Prim**-calls. To formalize, consider the following sampling process. Given D distinct pairs $\mathcal{I} = \{(i_0^{(\ell)}, p^{(\ell)}) \in \{1, \dots, u\} \times \mathcal{P}\}_{\ell=1, \dots, D}$,

1. We invoke **KGen** for u times to have u public parameters $\mathbf{P} = (\text{pp}_1, \dots, \text{pp}_u)$;
2. We follow the strategy of **S** and sample D strings of $n - \kappa$ bits with replacement and define the list $\mathcal{L} = \{(i_0^{(1)}, p^{(1)}, r^{(1)}), \dots, (i_0^{(D)}, p^{(D)}, r^{(D)})\}$.

Based on \mathcal{L} , we define a quantity for the maximal frequency of a certain label value, i.e.,

$$\mu(\mathcal{L}) := \max_{t \in \{0,1\}^\nu} \left\{ \left| \{(i_0, p, r) \in \mathcal{L} : \text{lf}_{p,j}(r, \text{pp}_{i_0}) = t\} \right| \right\}. \quad (4)$$

We denote this sampling process by $\mathbf{P} \leftarrow \text{KGen}, \mathcal{L} \leftarrow \text{S}$. Note that the quantity reflects a *property across multiple users in the ideal world*. A trivial upper bound is $\mu(\mathcal{L}) \leq D$, see Sect. 4.2 below. As will be seen in Theorems 1–3, the smaller $\mu(\mathcal{L})$, the better concrete security. A promising choice is to define $\mu(\mathcal{L})$ to be a (pseudo)random variable, as will be treated in Theorems 1-3. For example, one can choose a random initialization vector IV for every $\text{GGGM}^{\text{Prim}}$ instances/users and define $\text{lf}_{p,j}(r, \text{pp}) := \cdot \|\text{IV}$. This limits collisions between labels in distinct $\text{GGGM}^{\text{Prim}}$ instances and decreases $\mu(\mathcal{L})$. This approach will be used to improve FSS (see Sect. 7). As will be seen, a more sophisticated approach is to extract a part of the intermediate values as pseudorandom bits and “embed” them in the images of $\text{lf}_{p,j}(r, \text{pp})$. This approach is used in Bip32 key tree (see Sect. 6).

4.2 Random Oracle-based Trees

In detail, we consider defining $\text{Prim}(x, y) := \mathbf{H}(x\|y)$ for an FIL RO $\mathbf{H} : \{0, 1\}^{\nu+\omega} \mapsto \{0, 1\}^{wn}$. Obviously, this is equivalent with $\text{Prim}(x, y) := \mathbf{KH}(x, y)$ for a keyed FIL random oracle $\mathbf{KH} : \{0, 1\}^\nu \times \{0, 1\}^\lambda \mapsto \{0, 1\}^{wn}$. Our main result is as follows.

Theorem 1. *Assume using the simulator **S** defined in Fig. 4, and:*

- (i) $\mathbf{H} : \{0, 1\}^{\nu+\omega} \mapsto \{0, 1\}^{wn}$ is modeled as a random oracle, and
- (ii) \mathbf{L} leaks $\text{left}_{n-\kappa}(\text{Nd}(i_0, p^*))$ for every intermediate node $\text{Nd}(i_0, p^*)$ (see Sect. 4.1), and
- (iii) There exist quantities C and ε_μ such that

$$\Pr_{\mathbf{P} \leftarrow \text{KGen}, \mathcal{L} \leftarrow \mathbf{S}}[\mu(\mathcal{L}) > C] \leq \varepsilon_\mu. \quad (5)$$

Then, $\text{GGGM}^{\mathbf{H}}$ is a (u, T, D, ε) - (\mathbf{L}, \mathbf{S}) -constrained PRF for set system $\mathbb{S}_{\text{pre}, \text{GGGM}}$ of Eq. (3), where

$$\varepsilon = 2\varepsilon_\mu + \frac{2C \cdot (T + D)}{2^\kappa}. \quad (6)$$

Interpretation. The “leakages” and its simulator can be eliminated from Theorem 1 by setting $\kappa = n$. Therefore, the result implies the classical CPRF security of $\text{GGGM}^{\mathbf{H}}$.

If $\text{pp}_1, \dots, \text{pp}_u$ are picked according to some distribution, then $\mu(\mathcal{L})$ remains a random variable. An example is $\text{FMTr}^{\text{DM}^{\text{AES}}}$ in Sect. 7. If $\text{pp}_1, \dots, \text{pp}_u$ are not random, then $\mu(\mathcal{L})$ will be fixed by the design of the $\text{GGGM}^{\mathbf{H}}$ instance. In particular, when lf is a constant function, we could (only) appeal for the trivial upper bound $C = D$. In this case, $\varepsilon_\mu = 0$, and the obtained bound $DT/2^\kappa + D^2/2^\kappa$ indicates the $\kappa - \log_2 D$ bits security mentioned in the introduction.

The leaked intermediate values are indistinguishable with the random simulated leakages. This means protocols can extract $n - \kappa$ pseudorandom bits from every n -bit intermediate value in $\text{GGGM}^{\mathbf{H}}$ for arbitrary uses, as if these bits are “leaked” (to everyone including the adversaries). Indeed, **Bip32** does use these $n - \kappa$ bits to derive signature keys: see Fig. 1 or Sect. 6.

Proof idea. Recall that the goal is to derive a bound for

$$\left| \Pr[\mathfrak{D}^{\text{muCo}_K^{\mathbf{L}}, \text{muEv}_K^{\mathbf{L}}, \mathbf{H}} = 1] - \Pr[\mathfrak{D}^{\text{muCo}_K^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H}} = 1] \right|.$$

In the ideal world $\mathfrak{D}^{\text{muCo}_K^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H}}$, the outputs of the oracle $\text{muCo}_K^{\mathbf{S}}$ (i.e., the constrained keys) depend on the secret keys k_1, \dots, k_u . This complicates the H-coefficient based analysis. To remedy, we introduce a *random constraining oracle* $\mathfrak{S}\text{Cons}^{\mathbf{S}}$, which accepts queries of the same form (i_0, p^*) as $\text{muCo}_K^{\mathbf{S}}$ (and $\text{muCo}_K^{\mathbf{L}}$), but $\mathfrak{S}\text{Cons}$ returns true random n -bit strings as the constrained key $\text{Nd}(i_0, p^*)$ (while \mathbf{S} returns random leakages as before). This means for a query $(i_0, i_1^*/\dots/i_d^*)$, the random constraining oracle $\mathfrak{S}\text{Cons}^{\mathbf{S}}$ returns in total $n + d'(n - \kappa)$ random bits. We take $(\mathfrak{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$ as an intermediate world. In this vein, when interacting with $(\mathfrak{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$, the information gained by \mathfrak{D} is completely independent of any CPRF key, easing the analysis.

We then proceed with two steps: first, we prove indistinguishability of the real and the intermediate worlds; and then, we prove indistinguishability of the intermediate and the ideal worlds.

Indistinguishability of $(\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{H})$ and $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$. For this step, we view the real $(\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{H})$ as the real world, and the intermediate world $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$ the ideal world. We prove the following bound using the H-coefficient method.

$$\left| \Pr[\mathcal{D}^{\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{H}} = 1] - \Pr[\mathcal{D}^{\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H}} = 1] \right| \leq \varepsilon_{\mu} + \frac{C \cdot (T + D)}{2^{\kappa}}. \quad (7)$$

Transcripts Since we are in the Random Oracle Model and aim at statistical indistinguishability, we could, without loss of generality [Pat09], consider a deterministic distinguisher \mathcal{D} interacting with the three oracles as reflected in Definition 4. To summarize the information gained by \mathcal{D} in a clear form, we introduce two list $\mathcal{Q}_{\mathbf{H}}$ and \mathcal{Q}_{Nd} . The list $\mathcal{Q}_{\mathbf{H}} = \{(x^{(1)}, y^{(1)}), \dots\}$ records \mathcal{D} 's queries/answers to/from \mathbf{H} , with $(x, y) \in \mathcal{Q}_{\mathbf{H}}$ meaning $\mathbf{H}(x) = y$. The list $\mathcal{Q}_{\text{Nd}} = \{(i_0^{(1)}, p^{(1)}, z^{(1)}, b^{(1)}), \dots\}$ records the values in the tree obtained by \mathcal{D} via either the “standard” responses or the leakages. Every tuple $(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}$ is such that $p \in \mathcal{P}$, and:

- $z \in \{0, 1\}^{n-\kappa}$ when $b = 0$, meaning that z is a leaked intermediate value;
- $z \in \{0, 1\}^n$ when $b = 1$, meaning that z is either a constrained key or an output of $\text{muEv}_{\mathbf{K}}$.

To simplify our proof language, we follow [GKWY20, GKW+20] and reveal a number of internal secrets to \mathcal{D} at the end of the interaction. In detail, we will reveal the κ -bit internal seeds to \mathcal{D} , and add them to the transcript \mathcal{Q}_{Nd} . In this respect, note that in the real world $(\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{H})$, for every resulted tuple $(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}$, the corresponding node value $\text{Nd}(i_0, p)$ necessarily appeared during processing the queries. Moreover, $z = \text{Nd}(i_0, p)$ when $b = 1$, and $z = \text{left}_{n-\kappa}(\text{Nd}(i_0, p))$ when $b = 0$. We will thus reveal the “missing” κ bits $z_R = \text{right}_{\kappa}(\text{Nd}(i_0, p))$ to \mathcal{D} and add them to \mathcal{Q}_{Nd} . Since $k_{i_0} = \text{Nd}(i_0, \perp)$, this means the u user keys k_1, \dots, k_u are also completely given. In the ideal world, we reveal and add random “dummy” bits to the transcript. We also append to the transcript u public parameters $\mathbf{P} = (\text{pp}_1, \dots, \text{pp}_u)$ that are sampled according to the same distribution as the real world. By this, we obtain an extended list $\mathcal{Q}_{\text{Nd}}^{\mathbf{X}} = \{(i_0^{(1)}, p^{(1)}, z^{(1)}, b^{(1)}), \dots\}$, among which $(i_0^{(1)}, p^{(1)}), (i_0^{(2)}, p^{(2)}), \dots$ are exactly the same as those in \mathcal{Q}_{Nd} , while $z^{(1)}, z^{(2)}, \dots$ are all “full” n -bit strings regardless of the values of $b^{(1)}, b^{(2)}, \dots$. In all, we define

$$\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}, \mathcal{Q}_{\mathbf{H}}, \mathbf{P})$$

as a transcript. It is without loss of generality to provide these additional values, since the distinguisher is free to ignore them. Below we will first define the set Θ_{bad} of bad transcripts, and then analyze good transcripts.

Internal evaluation list $\mathcal{Q}_{\mathbf{H}}^{\mathbf{X}}$ Regarding $\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$, we make crucial observations as follows. For any $(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$, $p = i_1/\dots/i_{d'}$ with $d' \geq 1$,

- it holds $(i_0^\circ, p^\circ, \text{Nd}(p^\circ), \star) \in \mathcal{Q}_{\text{Nd}}^{\text{X}}$ for every prefix p° of p , since $\text{Nd}(i_0^\circ, p^\circ)$ indeed appeared during computing $\text{Nd}(i_0, p) = z$;
- it holds $(i_0, p', \text{Nd}(p'), \star) \in \mathcal{Q}_{\text{Nd}}^{\text{X}}$ for every brother path $p' \in \text{Br}(p)$ of p . To see this, note that for any path $p \in \mathcal{P}$, $(i_0, p, \text{Nd}(p), \star) \notin \mathcal{Q}_{\text{Nd}}^{\text{X}}$ only if the corresponding value $\text{Nd}(i_0, p)$ is never computed during the interaction. The only possibility is that there exists a constraining query $\text{muCo}_{\mathbf{K}}^{\text{L}}(i_0, p^*)$ such that p^* is a prefix of p , and this forbids querying (i_0, p) due to the query restrictions mentioned in Sect. 4.1. This essentially implies $\text{Nd}(i_0, p')$ is never computed for any $p' \in \text{Br}(p)$. By this, $(i_0, p, \text{Nd}(i_0, p), \star) \notin \mathcal{Q}_{\text{Nd}}^{\text{X}} \Leftrightarrow \forall p' \in \text{Br}(p) : (i_0, p', \text{Nd}(i_0, p'), \star) \notin \mathcal{Q}_{\text{Nd}}^{\text{X}}$, and thus the claim.

By these, consider any $(i_0, p, z, \star) \in \mathcal{Q}_{\text{Nd}}^{\text{X}}$ with $p = i_1/\dots/i_{d'-1}/i_{d'}$, $1 \leq d' \leq d$. Let $p^* = i_1/\dots/i_{d'-1}$, $j^* = \lfloor \frac{i_{d'}}{w} \rfloor$, and let $p_\ell = i_1/\dots/i_{d'-1}/w j^* + \ell - 1$ ($1 \leq \ell \leq w$). Then $\mathcal{Q}_{\mathbf{H}}^{\text{X}}$ includes the “internal evaluation tuple” $(i_0, p^*, j^*, x^*, y^*)$, where (recall from Sect. 3.2 that $\text{of}(z, \cdot)$ is bijective)

$$\begin{aligned} x^* &= \text{lf}_{p^*, j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0, p^*)), \text{pp}_{i_0}) \parallel \text{sf}_{p^*, j^*}(\text{right}_\kappa(\text{Nd}(i_0, p^*))), \\ y^* &= \text{of}^{-1}(\text{Nd}(i_0, p^*), \text{Nd}(i_0, p_1)) \parallel \text{of}^{-1}(\text{Nd}(i_0, p^*), \text{Nd}(i_0, p_2)) \parallel \text{of}^{-1}(\text{Nd}(i_0, p^*), \text{Nd}(i_0, p_w)). \end{aligned} \quad (8)$$

By the above two observations, all the mentioned node values can be found in $\mathcal{Q}_{\text{Nd}}^{\text{X}}$, and x^*, y^* are thus well-defined. The additional fields p^* and j^* will significantly simplify proof languages. It can be seen that the set $\mathcal{Q}_{\mathbf{H}}^{\text{X}}$ records all the internal random oracle evaluations that appeared during the real world interaction, and $|\mathcal{Q}_{\mathbf{H}}^{\text{X}}| = D$ is the aforementioned *effective data complexity* of \mathcal{Q} . Indeed, this constitutes its motivation.

Bad transcripts A transcript $\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\text{X}}, \mathcal{Q}_{\mathbf{H}})$ is *bad*, if any of the following conditions is fulfilled:

- (B-1) $\mu \geq C$.
- (B-2) There exist a tuple $(i_0^*, p^*, j^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{H}}^{\text{X}}$ and a pair $(x, y) \in \mathcal{Q}_{\mathbf{H}}$ such that $x^* = x$.
- (B-3) There are distinct tuples $(i_0^*, p^*, j^*, x^*, y^*), (i_0^{**}, p^{**}, j^{**}, x^{**}, y^{**}) \in \mathcal{Q}_{\mathbf{H}}^{\text{X}}$ such that $x^* = x^{**}$.

The 1st condition captures that the “labels” of the internal calls are too concentrated. The 2nd condition addresses the case where an internal random oracle evaluation in a $\text{muCo}_{\mathbf{K}}^{\text{L}}/\text{muEv}_{\mathbf{K}}^{\text{L}}$ query collide with an adversarial offline random oracle query, while the 3rd condition addresses the case where distinct $\text{muCo}_{\mathbf{K}}^{\text{L}}/\text{muEv}_{\mathbf{K}}^{\text{L}}$ queries issue the same internal random oracle evaluation.

First, $\Pr[(\text{B-1})] \leq \varepsilon_\mu$ immediately follows from Eq. (5), since leakages in $\mathcal{Q}^{\text{SCons}^{\text{S}}, \text{R}^{\text{S}}, \mathbf{H}}$ are purely random. For (B-2), consider each choice of $((x, y), (i_0^*, p^*, j^*, x^*, y^*)) \in \mathcal{Q}_{\mathbf{H}} \times \mathcal{Q}_{\mathbf{H}}^{\text{X}}$. By the definition of $\mathcal{Q}_{\mathbf{H}}^{\text{X}}$, $x = x^*$ means $x = \text{lf}_{p^*, j^*}(\text{left}_{n-\kappa}(z^*), \text{pp}_{i_0^*}) \parallel \text{sf}_{p^*, j^*}(\text{right}_\kappa(z^*))$ for $(i_0^*, p^*, z^*, \star) \in \mathcal{Q}_{\text{Nd}}^{\text{X}}$. Since the seed $\text{right}_\kappa(z^*)$ is uniform in $\{0, 1\}^\kappa$ in the ideal world (it is the random “dummy” value appended to \mathcal{Q}), and since $\text{sf}_{p^*, j^*}(\cdot)$ is injective, the probability to have $\text{sf}_{p^*, j^*}(\text{right}_\kappa(z^*)) = \text{right}_\lambda(x)$ is $1/2^\kappa$.

Now, the condition $x = x^*$ is fulfilled only if $\text{left}_\nu(x) = \text{lf}_{p^*,j^*}(\text{left}_{n-\kappa}(z^*), \text{pp}_{i_0^*})$. By this, for any $t \in \{0,1\}^\nu$, we define $\mathcal{Q}_{\mathbf{H}}[t] := \{s \in \{0,1\}^\lambda : (t \| s, \star) \in \mathcal{Q}_{\mathbf{H}}\}$. Then, the probability that (B-2) is fulfilled w.r.t. the above tuple $(i_0^*, p^*, j^*, x^*, y^*)$ is $|\mathcal{Q}_{\mathbf{H}}[\text{lf}_{p^*,j^*}(\text{left}_{n-\kappa}(z^*), \text{pp}_{i_0^*})]|/2^\kappa$. By the above and our definition Eq. (4),

$$\begin{aligned} \Pr[(\text{B-2}) \mid \neg(\text{B-1})] &= \sum_{(i_0^*, p^*, j^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{H}}^{\times}} \frac{|\mathcal{Q}_{\mathbf{H}}[\text{lf}_{p^*,j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{pp}_{i_0^*})]|}{2^\kappa} \\ &= \sum_{t \in \{0,1\}^\nu} \sum_{(i_0^*, p^*, j^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{H}}^{\times} : \text{lf}_{p^*,j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{pp}_{i_0^*}) = t} \frac{|\mathcal{Q}_{\mathbf{H}}[t]|}{2^\kappa} \\ &\leq \sum_{t \in \{0,1\}^\nu} \frac{C \cdot |\mathcal{Q}_{\mathbf{H}}[t]|}{2^\kappa} \quad (\text{By } \neg(\text{B-1})) \\ &\leq \frac{C \cdot T}{2^\kappa} \quad (\text{Since } \sum_{t \in \{0,1\}^\nu} |\mathcal{Q}_{\mathbf{H}}[t]| = |\mathcal{Q}_{\mathbf{H}}| = T). \end{aligned}$$

For (B-3), consider each pair $(i_0^*, p^*, j^*, x^*, y^*), (i_0^{**}, p^{**}, j^{**}, x^{**}, y^{**}) \in \mathcal{Q}_{\mathbf{H}}^{\times}$. By the definition of $\mathcal{Q}_{\mathbf{H}}^{\times}$, it holds $x^* = \text{lf}_{p^*,j^*}(\text{left}_{n-\kappa}(z^*), \text{pp}_{i_0^*}) \| \text{sf}_{p^*,j^*}(\text{right}_\kappa(z^*))$ and $x^{**} = \text{lf}_{p^{**},j^{**}}(\text{left}_{n-\kappa}(z^{**}), \text{pp}_{i_0^{**}}) \| \text{sf}_{p^{**},j^{**}}(\text{right}_\kappa(z^{**}))$ for $(i_0^*, p^*, z^*, \star), (i_0^{**}, p^{**}, z^{**}, \star) \in \mathcal{Q}_{\text{Nd}}^{\times}$ respectively. Then the condition $x^* = x^{**}$ is fulfilled only if $(i_0^*, p^*) \neq (i_0^{**}, p^{**})$, as per our restriction on sf and lf. By this, $\text{right}_\kappa(z^*)$ and $\text{right}_\kappa(z^{**})$ are uniform and independent. Moreover, $\text{sf}_{p^*,j^*}(\cdot)$ and $\text{sf}_{p^{**},j^{**}}(\cdot)$ are injective. Hence, the probability to have $\text{sf}_{p^*,j^*}(\text{right}_\kappa(z^*)) = \text{sf}_{p^{**},j^{**}}(\text{right}_\kappa(z^{**}))$ is $1/2^\kappa$. Further, $x^* = x^{**}$ requires $\text{lf}_{p^*,j^*}(\text{left}_{n-\kappa}(z^*), \text{pp}_{i_0^*}) = \text{lf}_{p^{**},j^{**}}(\text{left}_{n-\kappa}(z^{**}), \text{pp}_{i_0^{**}})$. Therefore,

$$\begin{aligned} &\Pr[(\text{B-3}) \mid \neg(\text{B-1})] \\ &= \sum_{(i_0^*, p^*, j^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{H}}^{\times}} \left(\sum_{\substack{(i_0^{**}, p^{**}, j^{**}, x^{**}, y^{**}) \neq (i_0^*, p^*, j^*, x^*, y^*) : \\ \text{lf}_{p^*,j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{pp}_{i_0^*}) = \text{lf}_{p^{**},j^{**}}(\text{left}_{n-\kappa}(\text{Nd}(i_0^{**}, p^{**})), \text{pp}_{i_0^{**}})}} \right) \frac{1}{2^\kappa} \\ &\leq \sum_{(i_0^*, p^*, j^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{H}}^{\times}} \frac{C}{2^\kappa} \quad (\text{By } \neg(\text{B-1})) \\ &\leq \frac{C \cdot D}{2^\kappa}. \end{aligned}$$

In all, a union bound yields

$$\Pr[T_{\text{id}} \in \Theta_{\text{bad}}] \leq \varepsilon_\mu + \frac{C \cdot (T + D)}{2^\kappa}. \quad (9)$$

Ratio of probabilities of good transcripts. One key insight of the H-coefficient method is that, for any $\mathcal{Q} \in \Theta_{\text{good}}$, the probability ratio $\Pr[T_{\text{re}} = \mathcal{Q}] / \Pr[T_{\text{id}} = \mathcal{Q}]$

is equal to the ratio between the probability that the real-world oracles are consistent with \mathcal{Q} and the probability that the ideal-world oracles are consistent with \mathcal{Q} . Now, for any attainable transcript $\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\text{X}}, \mathcal{Q}_{\text{H}}, \mathbf{P})$, the probability that the ideal world transcript is consistent with \mathcal{Q} is always exactly

$$\frac{1}{2^{Twn}} \times \frac{1}{2^{|\mathcal{Q}_{\text{Nd}}^{\text{X}}|n}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}] \quad (10)$$

This is so since in $\mathfrak{D}^{\text{Cons}^{\text{S}}, \text{R}^{\text{S}}, \mathbf{H}}$,

- (i) the probability that a random oracle with wn -bit outputs is consistent with the T queries in \mathcal{Q}_{H} is exactly $1/2^{Twn}$;
- (ii) the probability that the random node values equal those in $\mathcal{Q}_{\text{Nd}}^{\text{X}}$ is $1/2^{|\mathcal{Q}_{\text{Nd}}^{\text{X}}|n}$. Note that these nodes include the u secret keys k_1, \dots, k_u since $\text{Nd}(i_0, \perp) = k_{i_0}$;
- (iii) the probability that the public parameter vector \mathbf{P} is produced by KGen is $\Pr[\mathbf{P} \leftarrow \text{KGen}]$.

Bounding the distinguishing advantage of \mathfrak{D} thus reduces to bounding the probability that the real world is consistent with transcripts $\mathcal{Q} \in \Theta_{\text{good}}$.

Let $\mathbf{H} \vdash \mathcal{Q}_{\text{H}}$ denote the event that a function \mathbf{H} is consistent with the queries/answers in \mathcal{Q}_{H} , i.e., that $\mathbf{H}(x) = y$ for all $(x, y) \in \mathcal{Q}_{\text{H}}$. Since, in the real world, the information in $\mathcal{Q}_{\text{Nd}}^{\text{X}}$ is completely determined by \mathbf{H} and the user keys $\mathbf{K} = (K_1, \dots, K_u)$, we can also write $(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\text{X}}$ to denote the event that the function \mathbf{H} and keys \mathbf{K} are consistent with the queries/answers in $\mathcal{Q}_{\text{Nd}}^{\text{X}}$. For a (good) transcript $\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\text{X}}, \mathcal{Q}_{\text{H}}, \mathbf{P})$, the probability that the real world is consistent with \mathcal{Q} is exactly

$$\begin{aligned} & \Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\text{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\text{H}}] \times \Pr[\mathbf{H} \vdash \mathcal{Q}_{\text{H}}] \times \Pr[\mathbf{K} \leftarrow \text{KGen}] \\ &= \Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\text{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\text{H}}] \times \Pr[\mathbf{H} \vdash \mathcal{Q}_{\text{H}}] \times \frac{1}{2^{un}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}] \quad (11) \end{aligned}$$

(using independence of \mathbf{K} and \mathbf{H}). We have $\Pr[\mathbf{H} \vdash \mathcal{Q}_{\text{H}}] = 1/2^{Twn}$ exactly as before. The crux of the proof thus reduces to bounding $\Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\text{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\text{H}}]$. For this, note that in the real world, the list $\mathcal{Q}_{\text{H}}^{\text{X}}$ essentially summarizes *all* the random oracle queries internally issued by $\text{GGGM}^{\mathbf{H}}$ for producing the transcript $\mathcal{Q}_{\text{Nd}}^{\text{X}}$. Therefore,

$$\Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\text{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\text{H}}] = \Pr[\forall (i_0^*, p^*, j^*, x^*, y^*) \in \mathcal{Q}_{\text{H}}^{\text{X}} : \mathbf{H}(x^*) = y^* \mid \mathbf{H} \vdash \mathcal{Q}_{\text{H}}].$$

We further show that the latter probability concerns with \mathbf{H} satisfying “new” and distinct equations. Let $\mathcal{Q}_{\text{H}}^{\text{X}} = ((i_0^{(1)}, p^{(1)}, j^{(1)}, x^{(1)}, y^{(1)}), \dots, (i_0^{(D)}, p^{(D)}, j^{(D)}, x^{(D)}, y^{(D)}))$ in arbitrary order. The probability can be expressed as

$$\prod_{\ell=1}^D \Pr[\mathbf{H}(x^{(\ell)}) = y^{(\ell)} \mid \mathbf{H} \vdash \mathcal{Q}_{\text{H}} \wedge \forall \ell' < \ell : \mathbf{H}(x^{(\ell')}) = y^{(\ell')}].$$

Fix some ℓ . Since the transcript is good, there is no query of the form $(x^{(\ell)}, \star)$ in $\mathcal{Q}_{\mathbf{H}}$ (since (B-2) does not occur), nor is $\mathbf{H}(x^{(\ell)})$ determined by the fact that $\mathbf{H}(x^{(\ell')}) = y^{(\ell')}$ for all $\ell' < \ell$ (since (B-3) does not occur). Thus, we have

$$\Pr[\mathbf{H}(x^{(\ell)}) = y^{(\ell)} \mid \mathbf{H} \vdash \mathcal{Q}_{\mathbf{H}} \wedge \forall \ell' < \ell : \mathbf{H}(x^{(\ell')}) = y^{(\ell')}] = 1/2^{wn}.$$

for all i . The term $1/2^{Dwn}$ thus follows.

It remains to determine the size of $\mathcal{Q}_{\mathbf{H}}^{\mathbf{x}}$. Recall that, for every pair

$$\left(i_0, p, j, \text{lf}_{p,j}(\text{left}_{n-\kappa}(\text{Nd}(i_0, p)), \text{pp}_{i_0}) \parallel \text{sf}_{p,j}(\text{right}_{\kappa}(\text{Nd}(i_0, p))), \bar{z}_1 \parallel \dots \parallel \bar{z}_w \right) \in \mathcal{Q}_{\mathbf{H}}^{\mathbf{x}},$$

it holds $(i_0, p/jw, \text{of}(\text{Nd}(i_0, p), \bar{z}_1), \star), \dots, (i_0, p/jw + w - 1, \text{of}(\text{Nd}(i_0, p), \bar{z}_w), \star) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{x}}$. On the other hand, it can be seen that, for every set of w triples $(i_0, p_1, z_1, \star), \dots, (i_0, p_w, z_w, \star) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{x}}$ such that $p_1 = p^*/j^*w, p_1, \dots, p_w \in \mathcal{B}r(p_1)$, the n -bit values $\text{Nd}(i_0, p^*) = z^*$ and z_1, \dots, z_w correspond to a unique tuple $(i_0, p^*, j^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{H}}^{\mathbf{x}}$ with $y^* = \text{of}^{-1}(z^*, z_1) \parallel \dots \parallel \text{of}^{-1}(z^*, z_w)$. By this, there is a one-to-one correspondence between *sets of w tuples in $\mathcal{Q}_{\text{Nd}}^{\mathbf{x}}$ with associated paths being brothers* and *tuples in $\mathcal{Q}_{\mathbf{H}}^{\mathbf{x}}$* , and thus

$$|\mathcal{Q}_{\mathbf{H}}^{\mathbf{x}}| = \frac{|\{(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{x}} : p \neq \perp\}|}{w}.$$

Note that $|\mathcal{Q}_{\text{Nd}}^{\mathbf{x}}| - |\{(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{x}} : p \neq \perp\}| = u$, and the u entries are $(1, \perp, k_1, \star), \dots, (u, \perp, k_u, \star)$. By this,

$$\begin{aligned} \text{Eq. (11)} &= \left(\frac{1}{2^{wn}}\right)^D \times \frac{1}{2^{Twn}} \times \frac{1}{2^{un}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}] \\ &= \left(\frac{1}{2^n}\right)^{|\{(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{x}} : p \neq \perp\}|} \times \frac{1}{2^{Twn}} \times \frac{1}{2^{un}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}] \\ &= \frac{1}{2^{Twn}} \times \left(\frac{1}{2^n}\right)^{|\mathcal{Q}_{\text{Nd}}^{\mathbf{x}}|} \times \Pr[\mathbf{P} \leftarrow \text{KGen}]. \end{aligned}$$

So the probability that the real world is consistent with the transcript is the same as Eq. (10). This means Eq. (9), the probability of obtaining bad transcripts, constitutes the gap between the real world $(\text{muCo}_{\mathbf{K}}^{\perp}, \text{muEv}_{\mathbf{K}}^{\perp}, \mathbf{H})$ and the intermediate world $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$.

Indistinguishability of $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$ and $(\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$. For this, we view $(\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$ as the real world and $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H})$ as the ideal, and prove the following bound:

$$\left| \Pr[\mathcal{D}^{\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H}} = 1] - \Pr[\mathcal{D}^{\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{H}} = 1] \right| \leq \varepsilon_{\mu} + \frac{C \cdot (T + D)}{2^{\kappa}}. \quad (12)$$

Essentially, this step requires establishing pseudorandomness of the constrained keys, which follows the same idea as the first step with a number of changes. In detail, we employ four list $\mathcal{Q}_{\mathbf{H}}, \mathcal{Q}_{\mathbf{S}}, \mathcal{Q}_{\text{Nd}}$, and \mathcal{Q}_R to keep the information

gained by \mathcal{D} . The RO-query list $\mathcal{Q}_{\mathbf{H}}$ is just the same as before. The list $\mathcal{Q}_{\mathbf{S}} = \{(i_0^{(1)}, p^{(1)}, r^{(1)}), \dots\}$ records all the leakages returned by \mathbf{S} , where $i_0^{(\ell)} \in \{1, \dots, u\}$, $p^{(\ell)} \in \mathcal{P}^*$, and $r^{(\ell)} \in \{0, 1\}^{n-\kappa}$ indicates that at the position $p^{(\ell)}$ in the $i_0^{(1)}$ -th tree, the leakage block returned by \mathbf{S} is $r^{(\ell)}$. The list \mathcal{Q}_{Nd} is modified such that $\mathcal{Q}_{\text{Nd}} = \{(i_0^{(1)}, p^{(1)}, z^{(1)}), \dots\}$ records the queries and responses of $\mathcal{S}\text{Cons}/\mu\text{Co}_{\mathbf{K}}$. Namely, the ℓ -th tuple $(i_0^{(\ell)}, p^{(\ell)}, z^{(\ell)}) \in \mathcal{Q}_{\text{Nd}}$ indicates that \mathcal{D} made a query $(i_0^{(\ell)}, p^{(\ell)})$ to $\mathcal{S}\text{Cons}/\mu\text{Co}_{\mathbf{K}}$ and the n -bit “standard” response is $z^{(\ell)}$. Finally, $\mathcal{Q}_R = \{(i_0^{(1)}, p^{(1)}, z^{(1)}), \dots\}$ records the queries and responses of \mathbf{R} , where $(i_0^{(\ell)}, p^{(\ell)}, z^{(\ell)}) \in \mathcal{Q}_{\text{Nd}}$ indicates that \mathcal{D} made a query $(i_0^{(\ell)}, p^{(\ell)})$ to $\mathbf{R}^{\mathbf{S}}$ and the n -bit (random) response of \mathbf{R} is $z^{(\ell)}$.

We also reveal the internal values corresponding to $\mu\text{Co}_{\mathbf{K}}^{\mathbf{S}}$'s evaluations to \mathcal{D} at the end of the interaction to extend \mathcal{Q}_{Nd} , but the strategy differs from the previous step (Sect. 4.2). In detail, note that every tuple $(i_0, p, z) \in \mathcal{Q}_{\text{Nd}}$ indicates \mathcal{D} making a query (i_0, p) to $\mu\text{Co}_{\mathbf{K}}^{\mathbf{S}}$, and for every p' that is prefix of p , the corresponding node value $z' = \text{Nd}(i_0, p')$ necessarily appeared during processing the query. We thus reveal all such node values $z' = \text{Nd}(i_0, p')$ to \mathcal{D} and add the corresponding triple (i_0, p', z') to \mathcal{Q}_{Nd} . We also reveal all the u secret keys k_1, \dots, k_u and add the corresponding triple (i_0, \perp, k_{i_0}) to \mathcal{Q}_{Nd} . In the ideal world, we reveal random “dummy” n -bit blocks $z' \xleftarrow{\$} \{0, 1\}^n$ to \mathcal{D} and add (i_0, p', z') to \mathcal{Q}_{Nd} correspondingly. By this, we obtain an extended list $\mathcal{Q}_{\text{Nd}}^{\mathbf{X}} = \{(i_0^{(1)}, p^{(1)}, z^{(1)}), \dots\}$, among which $(i_0^{(1)}, p^{(1)}), (i_0^{(2)}, p^{(2)}), \dots$ represents the positions to the nodes in \mathcal{Q}_{Nd} , while $z^{(1)}, z^{(2)}, \dots$ are all “full” n -bit strings that are either “real” intermediate values (in the real world) or random “dummy” blocks (in the ideal world). Finally, we also sample u public parameters \mathbf{P} in the ideal world. In all, we define

$$\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}, \mathcal{Q}_R, \mathcal{Q}_{\mathbf{S}}, \mathcal{Q}_{\mathbf{H}}, \mathbf{P})$$

as a transcript. Clearly, $|\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}| \leq D$. The real world probability $\Pr[T_{\text{re}} = \mathcal{Q}]$ is then written as

$$\Pr[\mathbf{H} \vdash \mathcal{Q}_{\mathbf{H}}] \times \Pr[\mathbf{S} \vdash \mathcal{Q}_{\mathbf{S}}] \times \Pr[\mathbf{R} \vdash \mathcal{Q}_R] \times \Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\mathbf{H}}] \times \Pr[\mathbf{P} \leftarrow \text{KGen}], \quad (13)$$

where $\mathbf{S} \vdash \mathcal{Q}_{\mathbf{S}}$ denotes the event that the random leakages returned by \mathbf{S} are consistent with those in $\mathcal{Q}_{\mathbf{S}}$, $\mathbf{R} \vdash \mathcal{Q}_R$ denotes $\mathbf{R}(i_0, p) = z$ for every $(i_0, p, z) \in \mathcal{Q}_R$, and $(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$ denotes the event that values generated by the oracle $\mu\text{Co}_{\mathbf{K}}$ (using \mathbf{H}) are consistent with the records in $\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$. The above expansion is possible since \mathbf{S} , \mathbf{R} , and (\mathbf{H}, \mathbf{K}) are independent. Similarly,

$$\begin{aligned} \Pr[T_{\text{id}} = \mathcal{Q}] &= \Pr[\mathbf{H} \vdash \mathcal{Q}_{\mathbf{H}}] \times \Pr[\mathbf{S} \vdash \mathcal{Q}_{\mathbf{S}}] \times \Pr[\mathbf{R} \vdash \mathcal{Q}_R] \\ &\quad \times \Pr[(\mathcal{S}\text{Cons}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}}] \times \Pr[\mathbf{P} \leftarrow \text{KGen}], \end{aligned} \quad (14)$$

since $(\mathcal{S}\text{Cons}, \mathbf{K})$ is independent from \mathbf{H} . Gathering Eqs. (13) and (14) yields

$$\frac{\Pr[T_{\text{re}} = \mathcal{Q}]}{\Pr[T_{\text{id}} = \mathcal{Q}]} = \frac{\Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\mathbf{H}}]}{\Pr[(\mathcal{S}\text{Cons}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}}]}, \quad (15)$$

and the problem thus reduces to bounding $\Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\mathbf{H}}]$. Subsequent analyses thus simply follow the same line as the 1st step. The definition of bad transcripts is the same as the first step, resulting in $\Pr[T_{\text{id}} \in \Theta_{\text{bad}}] \leq \varepsilon_{\mu} + \frac{C \cdot (T+D)}{2^{\kappa}}$. For a good transcript $\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}, \mathcal{Q}_R, \mathcal{Q}_S, \mathcal{Q}_{\mathbf{H}}, \mathbf{P})$, following the same line as the previous step, it can be shown that

$$\Pr[(\mathbf{H}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}} \mid \mathbf{H} \vdash \mathcal{Q}_{\mathbf{H}}] = \left(\frac{1}{2^n}\right)^{\left| \{(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{X}} : p \neq \perp\} \right|}.$$

The above established Eq. (12). Gathering Eqs. (7) and (12) yields Eq. (6) and completes the proof.

4.3 Davies-Meyer-based Trees

Using a block cipher $\mathbf{E} : \{0, 1\}^{\nu} \times \{0, 1\}^{wn} \mapsto \{0, 1\}^{wn}$, the Davies-Meyer-based derivation function is defined by $\mathbf{Prim}(x, y) := \text{DM}^{\mathbf{E}}(x, y) = \mathbf{E}(x, y) \oplus y$ (therefore, $\lambda = wn$). This setting is interesting when we are to instantiate the tree using (fixed-key) block ciphers in their Davies-Meyer modes.

Security of this model crucially relies on the feeding forward in Davies-Meyer, i.e., feeding the κ -bit secret $\text{right}_{\kappa}(z)$ forward. However, a bad outputting function may cancel these bits. In this respect, we additionally require that *the outputting function of only relies on the leftmost $n - \kappa$ bits of its seed*, i.e., $\text{of}(z, I) = \text{of}(\text{left}_{n-\kappa}(z), I)$ for any inputs $z, I \in \{0, 1\}^n$. Our formal results regarding $\text{GGGM}^{\text{DM}^{\mathbf{E}}}$ will emphasize on this additional restriction.

Theorem 2. *Assume using the simulator \mathcal{S} defined in Fig. 4, and:*

- (i) $\mathbf{E} : \{0, 1\}^{\nu} \times \{0, 1\}^{wn} \mapsto \{0, 1\}^{wn}$ is modeled as an ideal cipher;
- (ii) $\text{of} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is such that $\text{of}(z, I) = \text{of}(\text{left}_{n-\kappa}(z), I)$ for any $z, I \in \{0, 1\}^n$;
- (iii) \mathbf{L} leaks $\text{left}_{n-\kappa}(\text{Nd}(i_0, p^*))$ for every intermediate node $\text{Nd}(i_0, p^*)$;
- (iv) There exist quantities C and ε_{μ} such that

$$\Pr_{\mathbf{P} \leftarrow \text{KGen}, \mathcal{L} \leftarrow \mathcal{S}}[\mu(\mathcal{L}) > C] \leq \varepsilon_{\mu}. \quad (16)$$

Then, $\text{GGGM}^{\text{DM}^{\mathbf{E}}}$ is a (u, T, D, ε) - $(\mathbf{L}, \mathcal{S})$ -constrained PRF for the set system $\mathbb{S}_{\text{pre}, \text{GGGM}}$ of Eq. (3), where

$$\varepsilon = 2\varepsilon_{\mu} + \frac{4C \cdot (T + D)}{2^{\kappa}}. \quad (17)$$

We also rely on the random intermediate system $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{E})$ for “relay”.

Indistinguishability of $(\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{E})$ and $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{E})$ View $(\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{E})$ as the real world, and $(\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{E})$ the ideal world. We prove the following bound.

$$\left| \Pr[\mathcal{D}^{\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{E}} = 1] - \Pr[\mathcal{D}^{\mathcal{S}\text{Cons}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{E}} = 1] \right| \leq \varepsilon_{\mu} + \frac{2C \cdot (T + D)}{2^{\kappa}}. \quad (18)$$

The setting is similar to that studied in Sect. 4.2, except that the random oracle \mathbf{H} is replaced with an ideal cipher $\mathbf{E} : \{0, 1\}^{\nu} \times \{0, 1\}^{wn} \mapsto \{0, 1\}^{wn}$ that can be queried in both forward and backward directions, with $wn = \lambda$. By this, the transcript of \mathcal{D} 's interaction consists of \mathcal{Q}_{Nd} and $\mathcal{Q}_{\mathbf{E}}$, where $\mathcal{Q}_{\mathbf{E}} = \{(L_1, x_1, y_1), \dots\}$ records \mathcal{D} 's queries/answers to/from \mathbf{E} (with $(L, x, y) \in \mathcal{Q}_{\mathbf{E}}$ meaning $\mathbf{E}(L, x) = y$), while the transcript $\mathcal{Q}_{\text{Nd}} = \{(i_0^{(1)}, p^{(1)}, z^{(1)}, b^{(1)}), \dots\}$, $i_0^{(\ell)} \in \{1, \dots, u\}$, $p^{(\ell)} \in \mathcal{P}$, $z^{(\ell)} \in \{0, 1\}^n \cup \{0, 1\}^{n-\kappa}$, $b^{(\ell)} \in \{0, 1\}$, is just similar to Sect. 4.2. We also append the κ bit internal secrets to \mathcal{Q}_{Nd} to have the extended list $\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$ as in Sect. 4.2, and concentrate on $\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}, \mathcal{Q}_{\mathbf{E}}, \mathbf{P})$ with $\mathcal{Q}_{\text{Nd}}^{\mathbf{X}} = \{(i_0^{(1)}, p^{(1)}, z^{(1)}, b^{(1)}), \dots\}$, $i_0^{(\ell)} \in \{1, \dots, u\}$, $p^{(\ell)} \in \mathcal{P}$, $z^{(\ell)} \in \{0, 1\}^n$, $b^{(\ell)} \in \{0, 1\}$. Denote by $\mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}$ the event that a block cipher \mathbf{E} is consistent with the queries/answers in $\mathcal{Q}_{\mathbf{E}}$, i.e., that $\mathbf{E}(L, x) = y$ for all $(L, x, y) \in \mathcal{Q}_{\mathbf{E}}$. Then the probability of $\mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}$ for an ideal cipher \mathbf{E} (with wn -bit blocks and ν -bit keys) is exactly $(\prod_{L \in \{0, 1\}^{\nu}} (2^{wn})_{T_L})^{-1}$, where for integers $1 \leq b \leq a$, we set $(a)_b = a \cdot (a-1) \cdots (a-b+1)$ with $(a)_0 = 1$ by convention, and T_L is the number of tuples of the form (L, \star, \star) in $\mathcal{Q}_{\mathbf{E}}$ (thus $\sum_{L \in \{0, 1\}^{\nu}} T_L = T$). Therefore, for any attainable transcript $\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}, \mathcal{Q}_{\mathbf{E}}, \mathbf{P})$, the probability that the ideal world is consistent with \mathcal{Q} is

$$\frac{1}{\prod_{L \in \{0, 1\}^{\nu}} (2^{wn})_{T_L}} \times \frac{1}{2^{|\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}|n}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}]. \quad (19)$$

For the real world, we also write $(\mathbf{E}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$ to denote the event that the cipher \mathbf{E} and keys \mathbf{K} are consistent with the values in $\mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$. Similarly to Sect. 4.2, the real world probability is exactly

$$\Pr[(\mathbf{E}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}} \mid \mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}] \times \frac{1}{\prod_{L \in \{0, 1\}^{\nu}} (2^{wn})_{T_L}} \times \frac{1}{2^{un}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}], \quad (20)$$

and the problem also reduces to bounding $\Pr[(\mathbf{E}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}} \mid \mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}]$.

Internal evaluation list $\mathcal{Q}_{\mathbf{E}}^{\mathbf{X}}$. As in Sect. 4.2, we will show that, conditioned on $\mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}$, the event $(\mathbf{E}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$ is equivalent to \mathbf{E} satisfying a series of new and distinct equations. The conditions for bad transcripts are essentially defined to ensure these equations. We start by explicitly constructing the list $\mathcal{Q}_{\mathbf{E}}^{\mathbf{X}}$ of such equations. For this, consider any $(i_0, p, z, \star) \in \mathcal{Q}_{\text{Nd}}^{\mathbf{X}}$ with $p = p^*/i_{a'} \neq \perp$. Let $j^* = \lfloor \frac{i_{a'}}{w} \rfloor$ and $p_{\ell} = p^*/(wj^* + \ell - 1)$ ($1 \leq \ell \leq w$). Then the extended list $\mathcal{Q}_{\mathbf{E}}^{\mathbf{X}}$ includes

an “internal evaluation tuple” $(i_0, p^*, j^*, L^*, x^*, y^*)$, where $(z^* = \text{Nd}(i_0, p^*))$

$$\begin{aligned} L^* &:= \text{lf}_{p^*, j^*}(\text{left}_{n-\kappa}(z^*), \text{pp}_{i_0}), & x^* &:= \text{sf}_{p^*, j^*}(\text{right}_\kappa(z^*)) \\ y^* &:= (\text{of}^{-1}(z^*, \text{Nd}(i_0, p_1)) \parallel \dots \parallel \text{of}^{-1}(z^*, \text{Nd}(i_0, p_w))) \oplus \text{sf}_{p^*, j^*}(\text{right}_\kappa(z^*)) \end{aligned} \quad (21)$$

Bad transcripts. An extended transcript $\mathcal{Q} = (\mathcal{Q}_{\text{Nd}}^x, \mathcal{Q}_{\mathbf{E}}, \mathbf{P})$ is *bad*, if any of the following conditions is fulfilled:

- (B-1) $\mu \geq C$.
- (B-2) There exist a pair of tuples $((L, x, y), (i_0^*, p^*, j^*, L^*, x^*, y^*)) \in \mathcal{Q}_{\mathbf{E}} \times \mathcal{Q}_{\mathbf{E}}^x$ such that $(L, x) = (L^*, x^*)$, or $(L, y) = (L^*, y^*)$;
- (B-3) There exist distinct $(i_0^*, p^*, j^*, L^*, x^*, y^*), (i_0^{**}, p^{**}, j^{**}, L^{**}, x^{**}, y^{**}) \in \mathcal{Q}_{\mathbf{E}}^x$ with $(L^*, x^*) = (L^{**}, x^{**})$ or $(L^*, y^*) = (L^{**}, y^{**})$.

The bound $\Pr[(\text{B-1})] \leq \varepsilon_\mu$ also follows from Eq. (16) straightforwardly. For (B-2), consider each choice of $((L, x, y), (i_0^*, p^*, j^*, L^*, x^*, y^*)) \in \mathcal{Q}_{\mathbf{E}} \times \mathcal{Q}_{\mathbf{E}}^x$. By the fact that $\text{right}_\kappa(z^*)$ is uniform in $\{0, 1\}^\kappa$ for $(i_0^*, p^*, z^*, \star) \in \mathcal{Q}_{\text{Nd}}^x$, and that $\text{sf}_{p^*, j^*}(\star)$ is injective, the probability to have $x = x^*$ is $1/2^\kappa$.

The other condition $y = y^*$ translates into

$$y = (\text{of}^{-1}(r^*, \text{Nd}(i_0^*, p_1^*)) \parallel \dots \parallel \text{of}^{-1}(r^*, \text{Nd}(i_0^*, p_w^*))) \oplus \text{sf}_{p^*, j^*}(\text{right}_\kappa(\text{Nd}(i_0^*, p^*))),$$

where $p_\ell^* = p^*/(wj^* + \ell - 1)$ ($1 \leq \ell \leq w$) and $r^* = \text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*))$. In the ideal world, $\text{right}_\kappa(\text{Nd}(i_0^*, p^*))$ is uniform, and is independent from $\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{Nd}(i_0^*, p_1), \dots, \text{Nd}(i_0^*, p^*)$ (since $\text{right}_\kappa(\text{Nd}(i_0^*, p^*))$ is the random “dummy” value sampled at the end of the ideal world interaction). By this and by the injectivity of $\text{sf}_{p^*, j^*}(\star)$, the probability to have $y = y^*$ is $1/2^\kappa$.

Finally, note that $L = \text{lf}_{p^*, j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{pp}_{i_0^*})$ is necessary for both $(L, x) = (L^*, x^*)$ and $(L, y) = (L^*, y^*)$. By this, for any $L \in \{0, 1\}^\nu$, we define

$$\mathcal{Q}_{\mathbf{E}}^+[L] := \{x \in \{0, 1\}^{wn} : (L, x, \star) \in \mathcal{Q}_{\mathbf{E}}\}, \quad \mathcal{Q}_{\mathbf{E}}^-[L] := \{y \in \{0, 1\}^{wn} : (L, \star, y) \in \mathcal{Q}_{\mathbf{E}}\}.$$

Then, the condition (B-2) is equivalent with $\exists (i_0^*, p^*, j^*, L^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{E}}^x : x^* \in \mathcal{Q}_{\mathbf{E}}^+[L^*] \vee y^* \in \mathcal{Q}_{\mathbf{E}}^-[L^*]$, the probability of which is

$$\begin{aligned} \Pr[(\text{B-2}) \mid \neg(\text{B-1})] &\leq \sum_{(i_0^*, p^*, j^*, L^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{E}}^x} \frac{2|\mathcal{Q}_{\mathbf{E}}^+[L^*]|}{2^\kappa} \\ &\leq \sum_{L \in \{0, 1\}^\nu} \sum_{(i_0^*, p^*, j^*, L^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{E}}^x : \text{lf}_{p^*, j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{pp}_{i_0^*}) = L} \frac{2|\mathcal{Q}_{\mathbf{E}}^+[L]|}{2^\kappa} \\ &\leq \frac{2C \cdot T}{2^\kappa}. \end{aligned}$$

The last inequality follows from $\sum_{L \in \{0, 1\}^\nu} |\mathcal{Q}_{\mathbf{E}}^+[L]| = \sum_{L \in \{0, 1\}^\nu} |\mathcal{Q}_{\mathbf{E}}^-[L]| = T$.

For (B-3), consider each pair $(i_0^*, p^*, j^*, L^*, x^*, y^*), (i_0^{**}, p^{**}, j^{**}, L^{**}, x^{**}, y^{**}) \in \mathcal{Q}_{\mathbf{E}}^{\times}$. Due to our restriction on sf and lf , if $(i_0^*, p^*) = (i_0^{**}, p^{**})$ then $(j^* \neq j^{**}$ and) $x^* \neq x^{**}$. Thus we could focus on the case $p^* \neq p^{**}$, meaning that the probability to have $\text{sf}_{p^*, j^*}(\text{right}_{\kappa}(\text{Nd}(i_0^*, p^*))) = \text{sf}_{p^{**}, j^{**}}(\text{right}_{\kappa}(\text{Nd}(i_0^{**}, p^{**})))$ is $1/2^{\kappa}$ since both $\text{Nd}(i_0^*, p^*)$ and $\text{Nd}(i_0^{**}, p^{**})$ are independent and uniform. On the other side, the equality $y^* = y^{**}$ translates into

$$\begin{aligned} & (\text{of}^{-1}(r^*, \text{Nd}(i_0^*, p_1^*)) \parallel \dots \parallel \text{of}^{-1}(r^*, \text{Nd}(i_0^*, p_w^*))) \oplus \text{sf}_{p^*, j^*}(\text{right}_{\kappa}(\text{Nd}(i_0^*, p^*))), \\ & = (\text{of}^{-1}(r^{**}, \text{Nd}(i_0^{**}, p_1^{**})) \parallel \dots \parallel \text{of}^{-1}(r^{**}, \text{Nd}(i_0^{**}, p_w^{**}))) \oplus \text{sf}_{p^{**}, j^{**}}(\text{right}_{\kappa}(\text{Nd}(i_0^{**}, p^{**}))), \end{aligned}$$

where $p_{\ell}^* = p^*/(wj^* + \ell - 1)$ ($1 \leq \ell \leq w$), $r^* = \text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*))$, $p_{\ell}^{**} = p^{**}/(wj^{**} + \ell - 1)$ ($1 \leq \ell \leq w$) and $r^{**} = \text{left}_{n-\kappa}(\text{Nd}(i_0^{**}, p^{**}))$. In the ideal world, $\text{right}_{\kappa}(\text{Nd}(i_0^*, p^*))$ and $\text{right}_{\kappa}(\text{Nd}(i_0^{**}, p^{**}))$ are uniform and independent. By the injectivity of sf , the probability to have $y^* = y^{**}$ is $1/2^{\kappa}$.

Last, $(L^*, x^*) = (L^{**}, x^{**})$ and $(L^*, y^*) = (L^{**}, y^{**})$ hold only if $\text{lf}_{p^*, j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{pp}_{i_0^*}) = \text{lf}_{p^{**}, j^{**}}(\text{left}_{n-\kappa}(\text{Nd}(i_0^{**}, p^{**})), \text{pp}_{i_0^{**}})$. Meanwhile, by our definition Eq. (4) and by $\neg(\text{B-1})$, for each pair (p^*, j^*) , the number of pairs (p^{**}, j^{**}) satisfying $\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)) = \text{left}_{n-\kappa}(\text{Nd}(i_0^{**}, p^{**}))$ cannot exceed C . By these, the probability $\Pr[(\text{B-3}) \mid \neg(\text{B-1})]$ is bounded by

$$\begin{aligned} & \sum_{(i_0^*, p^*, j^*, L^*, x^*, y^*) \in \mathcal{Q}_{\mathbf{E}}^{\times}} \left(\sum_{\substack{(i_0^{**}, p^{**}, j^{**}, L^{**}, x^{**}, y^{**}) \neq (i_0^*, p^*, j^*, L^*, x^*, y^*) : \\ \text{lf}_{p^*, j^*}(\text{left}_{n-\kappa}(\text{Nd}(i_0^*, p^*)), \text{pp}_{i_0^*}) = \text{lf}_{p^{**}, j^{**}}(\text{left}_{n-\kappa}(\text{Nd}(i_0^{**}, p^{**})), \text{pp}_{i_0^{**}})}} \right) \frac{2}{2^{\kappa}} \\ & \leq \frac{2C \cdot |\mathcal{Q}_{\text{Nd}}^{\times}|}{2^{\kappa}}. \end{aligned}$$

In all, a union bound yields

$$\Pr[T_{\text{id}} \in \Theta_{\text{bad}}] \leq \varepsilon_{\mu} + \frac{2C \cdot (T + D)}{2^{\kappa}}. \quad (22)$$

Ratio of Probabilities of Good Transcripts. Fix a good transcript \mathcal{Q} . The idea resembles Sect. 4.2, concentrating on analyzing $\Pr[(\mathbf{E}, \mathbf{K}) \vdash \mathcal{Q}_{\text{Nd}}^{\times} \mid \mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}] = \Pr[\mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}^{\times} \mid \mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}}]$. Let

$$\mathcal{Q}_{\mathbf{E}}^{\times} = ((i_0^{(1)}, p^{(1)}, j^{(1)}, L^{(1)}, x^{(1)}, y^{(1)}), \dots, (i_0^{(D)}, p^{(D)}, j^{(D)}, L^{(D)}, x^{(D)}, y^{(D)}))$$

in arbitrary order, then the latter probability can be expressed as

$$\prod_{\ell=1}^D \Pr[\mathbf{E}(L^{(\ell)}, x^{(\ell)}) = y^{(\ell)} \mid \mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}} \wedge \forall \ell' < \ell : \mathbf{E}(L^{(\ell')}, x^{(\ell')}) = y^{(\ell')}].$$

Fix some ℓ . Since the transcript is good, there is no query of the form $(L^{(\ell)}, x^{(\ell)}, \star)$ in $\mathcal{Q}_{\mathbf{E}}$ (since (B-2) does not occur), nor is $\mathbf{E}(L^{(\ell)}, x^{(\ell)})$ determined by the fact that $\mathbf{E}(L^{(\ell')}, x^{(\ell')}) = y^{(\ell')}$ for all $\ell' < \ell$ (since (B-3) does not occur). Similarly by symmetry, there is no query of the form $(L^{(\ell)}, \star, y^{(\ell)})$ in $\mathcal{Q}_{\mathbf{E}}$ (since (B-2) does not occur), nor is


```

Algorithm Pu( $K, \mathcal{S}$ ) // Puncture a set  $\mathcal{S} \subseteq \mathcal{X}$ 
parse ( $k, \text{pp}$ )  $\leftarrow K$ 
 $list \leftarrow \emptyset$ 
Let  $\mathcal{P}_c \subseteq \mathcal{P}$  be such that  $p \notin \mathcal{S}$  if and only if there exists  $p^* \in \mathcal{P}_c$  and  $p^*$  is prefix of  $p$ 
for  $p^* \in \mathcal{P}_c$  do
   $ckey_{p^*} \leftarrow \text{StepDown}(k, \text{pp}, p^*)$ ,  $list \leftarrow list \cup (p^*, ckey_{p^*})$ 
end for
return  $K\{\mathcal{S}\} \leftarrow (\text{pp}, list)$ 

```

Fig. 5. The Pu algorithm corresponding to $\text{GGGM}^{\text{Prim}}[\kappa, n, \nu, \theta, w, d, \text{sf}, \text{lf}]$. $\mathbf{E}^{-1}(L^{(\ell)}, y^{(\ell)})$ determined by the fact that $\mathbf{E}(L^{(\ell')}, y^{(\ell')}) = x^{(\ell')}$ for all $\ell' < \ell$ (since (B-3) does not occur). Thus, we have

$$\Pr[\mathbf{E}(L^{(\ell)}, x^{(\ell)}) = y^{(\ell)} \mid \mathbf{E} \vdash \mathcal{Q}_{\mathbf{E}} \wedge \forall \ell' < \ell : \mathbf{E}(L^{(\ell')}, x^{(\ell')}) = y^{(\ell')}] \geq 1/2^{wn}.$$

for all ℓ . Finally, it also holds $|\mathcal{Q}_{\mathbf{E}}^{\times}| = |\{(i_0, p, z, b) \in \mathcal{Q}_{\text{Nd}}^{\times} : p \neq \perp\}|/w$, which resembles Sect. 4.2. By the above,

$$\begin{aligned} \text{Eq. (20)} &\geq \left(\frac{1}{2^{wn}}\right)^D \times \frac{1}{\prod_{L \in \{0,1\}^{\nu}} (2^{wn})_{T_L}} \times \frac{1}{2^{un}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}] \\ &= \left(\frac{1}{2^n}\right)^{|\mathcal{Q}_{\text{Nd}}^{\times}|} \times \frac{1}{\prod_{L \in \{0,1\}^{\nu}} (2^{wn})_{T_L}} \times \Pr[\mathbf{P} \leftarrow \text{KGen}], \end{aligned}$$

meaning that the probability that the real world is consistent with the transcript is at least Eq. (19). This completes the proof.

Indistinguishability of $(\mathcal{S}\text{Cons}^{\mathcal{S}}, \mathbf{R}^{\mathcal{S}}, \mathbf{E})$ and $(\text{muCo}_{\mathbf{K}}^{\mathcal{S}}, \mathbf{R}^{\mathcal{S}}, \mathbf{E})$ For this step, we view $(\text{muCo}_{\mathbf{K}}^{\mathcal{S}}, \mathbf{R}^{\mathcal{S}}, \mathbf{E})$ as the real world and $(\mathcal{S}\text{Cons}^{\mathcal{S}}, \mathbf{R}^{\mathcal{S}}, \mathbf{E})$ as the ideal. The core step is to establish pseudorandomness of the constrained keys, which combines the ideas of Sect. 4.2 and 4.3. In all, the bound remains.

$$\left| \Pr[\mathcal{D}^{\mathcal{S}\text{Cons}^{\mathcal{S}}, \mathbf{R}^{\mathcal{S}}, \mathbf{E}} = 1] - \Pr[\mathcal{D}^{\text{muCo}_{\mathbf{K}}^{\mathcal{S}}, \mathbf{R}^{\mathcal{S}}, \mathbf{E}} = 1] \right| \leq \varepsilon_{\mu} + \frac{2C \cdot (T + D)}{2^{\kappa}}. \quad (23)$$

Gathering Eqs. (18) and (23) yields Eq. (17) and completes the proof.

5 Multi-user Leakage PPRF Security of GGGM Trees

While it was believed obvious, we first formally describe the puncturing algorithm Pu of $\text{GGGM}^{\text{Prim}}$ in Fig. 5 for completeness. This specification would also be used in subsequent proof of Theorem 3.

The leakage assumption and simulator \mathbf{S} are roughly the same as Sect. 4.1 and Fig. 4. With these, the leakage PPRF security of GGGM is as follows.

Theorem 3. *Assume using the simulator \mathbf{S} defined in Fig. 4, and that there exist quantities C and ε_{μ} such that $\Pr_{\mathbf{P} \leftarrow \text{KGen}, \mathcal{L} \leftarrow \mathbf{S}}[\mu(\mathcal{L}) > C] \leq \varepsilon_{\mu}$. Then:*

- The tree $\text{GGGM}^{\mathbf{H}}$ built upon a random oracle $\mathbf{H} : \{0,1\}^{\nu+\omega} \mapsto \{0,1\}^{wn}$ is a (u, T, D, ε) - (\mathbf{L}, \mathbf{S}) -puncturable PPRF, where $\varepsilon = 2\varepsilon_{\mu} + \frac{2C \cdot (T+D)}{2^{\kappa}}$;

- With the additional restriction that $\text{of} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is such that $\text{of}(z, I) = \text{of}(\text{left}_{n-\kappa}(z), I)$ for any $z, I \in \{0, 1\}^n$, the tree $\text{GGGM}^{\text{DM}^E}$ built upon an ideal cipher $\mathbf{E} : \{0, 1\}^\nu \times \{0, 1\}^{wn} \mapsto \{0, 1\}^{wn}$ is a (u, T, D, ε) - (\mathbf{L}, \mathbf{S}) -puncturable PPRF, where $\varepsilon = 2\varepsilon_\mu + \frac{4C \cdot (T+D)}{2^\kappa}$.

Proof. A distinguisher \mathfrak{D}_1 against the PPRF security of $\text{GGGM}^{\text{Prim}}$ can be transformed to a distinguisher \mathfrak{D}_2 against the CPRF security of $\text{GGGM}^{\text{Prim}}$. Indeed, \mathfrak{D}_2 runs \mathfrak{D}_1 and simulates the muPu_K^L oracle for \mathfrak{D}_1 using its muCo_K^L oracle, as shown in Fig. 5 (the algorithm $\text{Pu}(K, \mathcal{S})$ calls $\text{StepDown}(K, p^*)$, but as shown in Fig. 2, this internal procedure functions the same as $\text{Co}(K, p^*)$). It is easy to see \mathfrak{D}_2 's effective data complexity D and query complexity T are the same as \mathfrak{D}_1 . The bounds thus follow Theorems 1 and 2.

6 Improving Hierarchical Deterministic Wallets

We first review hierarchical deterministic wallet: its basic ideas in Sect. 6.1, its formalism in Sect. 6.2, and specification of Bip32 in Sect. 6.3. Then, in Sect. 6.4, we introduce mu security definitions for HDWs; in Sect. 6.5, we establish mu security for Bip32 using Theorem 1. Finally, we present improvements in Sect. 6.6.

6.1 Hierarchical Deterministic Wallets

Briefly, a wallet in blockchain consists of a pair of secret and public keys for a digital signature scheme. To transfer assets, the user signs *transactions* (i.e., messages) with its secret key, with the digital identity named “address” of the receiver (e.g., the receiver’s public key) embedded in the transactions. A deterministic wallet derives a sequence of *session key pairs* from a single master key and uses distinct public session keys as multiple identities for receiving, in order to achieve anonymity and limit the damage of (session) key exposure.

A hierarchical deterministic wallet (HDW) makes one step further, derives a collection of signing keys from a master and organizes them under an access hierarchy, where each element represents a group of users and each user has its associated keys. Users staying higher in the hierarchy must be able to derive the keys of users on lower levels and to further sign transactions on their behalf. Users on lower levels, however, should not be able to escalate their privileges along the hierarchy, not even when colluding with others. This *hierarchical access control* exactly fits into the manager(s)-departments architecture of large-scale enterprises. It also eases *wallet delegation* and *auditing*, which turns out to be crucial for e-commerce and Decentralized Finance.

Clearly, secret signing keys are central in wallets. To achieve decentralized key managements, a promising approach is to use *threshold signatures* [Des88, Lin19]. In a threshold signature, the secret signing key is divided into shares held by multiple parties. A threshold number (i.e., a subset) of these parties can follow the protocol (typically the TSS protocol [AHS20]) to collaboratively sign cryptocurrency transaction. The protocol is fully decentralized, and the collapse of fewer parties/shares will not incur loss of funding. See [AHS20] for a survey.

Decentralized key managements for HDW, however, require *dividing the master key into shares* and *computing the signing key shares from the master key shares without combining any one*. Thus, multiple parties holding the master key shares have to

evaluate the hierarchical key derivations in MPC. Probably, the best approach is to garble the key derivation functions. The bottleneck is communication overhead, which is mainly determined by the number of AND gates in key derivations. The default Bip32 configuration consumes 4 HMAC-SHA512 executions with nearly 1 million AND gates to derive a concrete signing key from the wallet seed (see Fig. 1). This incurs a heavy communication cost and prohibits the use, especially for mobile users. While some ideas to remedy appear obvious, the soundness is unclear due to the lack of formalism and justification.

6.2 Formalism of Public-underivable HDW

Luzio et al. introduced a model for HDW [LFA20]. Their formalism particularly emphasizes *public-derivability*, i.e., the ability of *generating all public keys of the wallet without relying on any secret information*. As they discussed [LFA20, Sect. 7], public-derivability somewhat contradicts the privacy notion (*transaction*) *unlinkability*. However, Bip32 [Med18, Security] indeed *insisted on unlinkability rather than public-derivability* (see Sect. 6.4 for elaboration). Bip32 tries to achieve unlinkability-like notion by using *secret* “chain codes”, which essentially disables public-derivability. Auditing was intended to be carried out by granting certain secret “chain codes” to the auditor. With these considerations, we henceforth refer to Luzio et al.’s formalism [LFA20, Sect. 5] as *public-derivable HDW*, and resort to a natural relaxation (*pub-underivable*) *HDW*.¹¹

Despite disabling public-derivability, our formalism of HDW still follows [LFA20]. Concretely, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a directed acyclic graph (DAG) representing the access hierarchy of the HDW, where:

- (i) $\mathcal{V} = \{v_{\perp}, v_{p_1}, v_{p_2}, v_{p_3} \dots\}$ is the set of vertexes, where v_{\perp} has indegree 0 (meaning that v_{\perp} has the highest privilege in the hierarchy defined by \mathcal{G}).
- (ii) \mathcal{E} is the set of edges, and elements in \mathcal{E} are of the form $(v_p \xrightarrow{i_j} v_{p/i_j})$ indicating an edge from v_p to v_{p/i_j} (of lower privilege) with label i_j . To ease understanding and highlight the concrete hierarchy in \mathcal{G} , we use the path notation system of Sect. 3.2 for subscripts. This means a vertex $v_{i_1/i_2/\dots/i_d}$ is in \mathcal{V} *if and only if* the d edges $(v_{\perp} \xrightarrow{i_1} v_{i_1}), (v_{i_1} \xrightarrow{i_2} v_{i_1/i_2}), \dots, (v_{i_1/i_2/\dots/i_{d-1}} \xrightarrow{i_d} v_{i_1/i_2/\dots/i_d})$ are all in \mathcal{E} .

We define the set of descendants $\text{Desc}(\mathcal{G}, v_p)$ of node v_p to be the set of nodes $v_{p'}$ such that there exists a direct path from v_p to $v_{p'}$ in \mathcal{G} . By our notations, this means there exists $p^* \in \mathcal{P}^*$ such that $p' = p/p^*$.

With these, a (*public-underivable*) *hierarchical deterministic wallet* $\text{Wal} = (\text{Setup}, \text{DPub}, \text{DPriv}, \text{Sign}, \text{Vrfy})$, defined over seed space \mathcal{S} and message space \mathcal{M} , is defined in the following way:

- **Setup**(\mathcal{G}, \mathcal{S}): The deterministic setup algorithm (a.k.a. *master key generation* due to [DFL19]) takes as input an initial seed $S \in \mathcal{S}$ and a DAG \mathcal{G} that has a unique node with indegree 0, and returns the keys $(\text{wsk}_{\perp}, \text{wpk}_{\perp})$ of the node $v_{\perp} \in \mathcal{V}$ with indegree 0 (with the highest privileges). $(\text{wsk}_{\perp}, \text{wpk}_{\perp})$ are also viewed as the master wallet secret and public keys of the HDW.
- **DPub**($\mathcal{G}, \text{wpk}_p, v_p, v_{p'}$): The *delegated deterministic public derivation* algorithm takes as input the wallet public key wpk_p associated to node $v_p \in \mathcal{V}$ and a target node $v_{p'} \in \text{Desc}(\mathcal{G}, v_p)$, and outputs the wallet public key $\text{wpk}_{p'}$ of $v_{p'}$. This functionality fits into the “audit” use case of [Med18, Use cases].

¹¹ Concurrently to us, Das et al. [DEF⁺21] also proposed a model for HDW.

- $\text{DPriv}(\mathcal{G}, \text{wsk}_p, v_p, v_{p'})$: The deterministic private derivation algorithm takes as input the wallet secret key wsk_p associated to node $v_p \in \mathcal{V}$ and a target node $v_{p'} \in \text{Desc}(\mathcal{G}, v_p)$, and outputs the wallet secret key $\text{wsk}_{p'}$ of $v_{p'}$.
- $\text{WSign}(\text{wsk}_p, m)$: The randomized signing algorithm takes as input a message $m \in \mathcal{M}$ and a wallet secret key wsk_p , and outputs a signature σ .
- $\text{WVrfy}(pk, m, \sigma)$: The deterministic verification algorithm takes as input a signature public key pk , a message m , and a signature σ . It outputs 1 (accept) or 0 (reject).

A hierarchical deterministic wallet is *correct*, if any user v_p (that holds the wallet secret key wsk_p) can derive the private and public keys $\text{wsk}_{p'}$ and $\text{wpk}_{p'}$ of any of its descendants $v_{p'} \in \text{Desc}(\mathcal{G}, v_p)$ and create a valid signature on behalf of $v_{p'}$ (i.e., that passes the verification process against the public key $\text{wpk}_{p'}$ obtained through public key derivation). We omit the formalism.

6.3 Bip32 HDW and the Underlying GGM Instance

The Bitcoin Improvement Proposal Bip32 [Med18] uses a double-input function $\mathbf{Prim} : \{0, 1\}^{512} \times \{0, 1\}^{256} \mapsto \{0, 1\}^{512}$. In Fig. 6, we provide a description of Bip32 following the formalism in Sect. 6.2. Briefly speaking,

- During $\text{Setup}(\mathcal{G}, S)$, the wallet first invokes \mathbf{Prim} with the seed S to derive master keys $(ch_\perp, sk_\perp, pk_\perp)$ that are intended to be held by the enterprise, where (sk_\perp, pk_\perp) are the signature secret and public key and mch is an additional secret called *chain code*. Essentially, the chain code functions as secret seeds of pseudo-random primitives. As shown in Fig. 1, $\text{wsk}_\perp = (sk_\perp, ch_\perp)$ and $\text{wpk}_\perp = (pk_\perp, ch_\perp)$ are viewed as the extended private and public keys of the enterprise.
- Given the extended secret key $\text{wsk}_p = (sk_p, ch_p)$ of a node v_p , the CKDpriv algorithm derives the extended secret key $\text{wsk}_{p/i} = (sk_{p/i}, ch_{p/i})$ of a child node $v_{p/i}$ via invoking \mathbf{Prim} once and calculating a modular addition.
- Given the extended public key $\text{wpk}_p = (pk_p, ch_p)$ of a node v_p , the CKDpub algorithm derives the extended public key $\text{wpk}_{p/i} = (pk_{p/i}, ch_{p/i})$ of a child node $v_{p/i}$, thanks to the homomorphism property between the secret key and public key space.

Note that by the specification, a node v_p specified by \mathcal{G} may still have $\text{wsk}_p = \perp$, although the probability is extremely low. A default configuration for an enterprise with several offices is recommended in [Med18]. The default key tree is essentially a GGM tree with depth $d = 3$, as illustrated in Fig. 1.

As discussed in Sect. 3.1, the key tree of Bip32 is actually a major motivation of our model. Actually, the definition and analysis of Bip32 can be based on an instance of GGM that (roughly) takes the 512-bit strings $z_p = [sk_p]_{256} \| ch_p$ as nodes (with chain code ch_p functioning as the secrets). In detail, consider the GGM instance with parameters as follows.

- Node size $n = 512$, security parameter $\kappa = 256$;
- $\theta \leq 2^{32}$, and $w = 1$, i.e., output size $wn = n = 512$;
- Depth d equals the length of the longest (directed) path in \mathbf{G} . In this vein, a (directed) path $v_\perp \xrightarrow{i_1} v_{i_1} \xrightarrow{i_2} \dots \xrightarrow{i_{d'}} v_{i_1/\dots/i_{d'}}$ in \mathcal{G} identifies a path $p^* = i_1/\dots/i_{d'}$ in this tree;
- The root node $\text{Nd}(\perp) = \mathbf{Prim}(S, \text{“Bitcoin seed”})$;

<pre> Algorithm Setup(\mathcal{G}, S) // $128 \leq S \leq 512$ $I_{\perp} \leftarrow \text{HMAC}(S, \text{"Bitcoin seed"})$ $sk_{\perp} \leftarrow \text{int}(\text{left}_{256}(I_{\perp})) \bmod \mathbb{G}$ if $sk_{\perp} = 0$ or $sk_{\perp} \geq \mathbb{G}$ then return \perp $pk_{\perp} \leftarrow sk_{\perp} \cdot G, ch_{\perp} \leftarrow \text{right}_{256}(I_{\perp})$ $wsk_{\perp} \leftarrow (sk_{\perp}, ch_{\perp}), wpk_{\perp} \leftarrow (pk_{\perp}, ch_{\perp})$ Algorithm DPub($\mathcal{G}, wpk_p, v_p, v_{p'}$) if $v_{p'} \notin \text{Desc}(\mathcal{G}, v_p)$ then return \perp parse $p/i_1/i_2/\dots/i_{d'}$ $\leftarrow p'$ $wpk \leftarrow wpk_p, p^* \leftarrow p$ for $j = 1, \dots, d'$ do $wpk \leftarrow \text{CKDpub}(wpk, i_j)$ $p^* \leftarrow p^*/i_j$ endfor return wpk Algorithm DPriv($\mathcal{G}, wsk_p, v_p, v_{p'}$) if $v_{p'} \notin \text{Desc}(\mathcal{G}, v_p)$ then return \perp parse $p/i_1/i_2/\dots/i_{d'}$ $\leftarrow p'$ $wsk \leftarrow wsk_p, p^* \leftarrow p$ for $j = 1, \dots, d'$ do $wsk \leftarrow \text{CKDpriv}(wsk, i_j)$ $p^* \leftarrow p^*/i_j$ endfor return wsk Algorithm WVrfy(pk, m, σ) return $\text{Sig.Vrfy}(pk, m, \sigma)$ </pre>	<pre> Algorithm WSign(wsk, m) $(sk, ch) \leftarrow wsk$ return $\text{Sig.Sign}(sk, m)$ Algorithm CKDpriv(wsk, i) $(sk, ch) \leftarrow wsk$ if $i \geq 2^{31}$ then // "hardened derivation" $I \leftarrow \text{HMAC}([0]_8 \parallel [sk]_{256} \parallel [i]_{32}, ch)$ else // $0 \leq i < 2^{31}$, normal derivation $pk \leftarrow sk \cdot G$ $I \leftarrow \text{HMAC}(\text{ser}_P(pk) \parallel [i]_{32}, ch)$ endif $\Delta \leftarrow \text{int}(\text{left}_{256}(I)), ch' \leftarrow \text{right}_{256}(I)$ if $\Delta \geq \mathbb{G}$ then return \perp $sk' \leftarrow (sk + \Delta) \bmod \mathbb{G}$ if $sk' = 0$ then return \perp return (sk', ch') Algorithm CKDpub(wpk, i) $(pk, ch) \leftarrow wpk$ if $i \geq 2^{31}$ then // invalid "hardened" return \perp else // $0 \leq i < 2^{31}$, normal derivation $I \leftarrow \text{HMAC}(\text{ser}_P(pk) \parallel [i]_{32}, ch)$ endif $\Delta \leftarrow \text{int}(\text{left}_{256}(I)), ch' \leftarrow \text{right}_{256}(I)$ if $\Delta \geq \mathbb{G}$ then return \perp $pk' \leftarrow pk + \Delta \cdot G$ // ECC addition if pk' is the point at infinity then return \perp return (pk', ch') </pre>
--	--

Fig. 6. Specification of Bip32 HDW following the formalism of Sect. 6.2.

- For a node value $z \in \{0, 1\}^{512}$ and $(p, j) \in (\{\perp\} \cup \{i_1/i_2/\dots/i_{d'}\}_{d' \in \{1, \dots, d-1\}}) \times \{0, \dots, 2^{32}-1\}$, let $r = \text{left}_{256}(z)$, then the seeding, labeling and outputting functions are as follows.

$$\text{sf}_{p,j}(\text{right}_{256}(z)) = \text{right}_{256}(z) \quad (24)$$

$$\text{lf}_{p,j}(r, \text{pp}) = \begin{cases} [0]_8 \parallel r \parallel [j]_{32} & \text{if } 2^{31} \leq j < 2^{32} \\ \text{ser}_P(\text{int}(r) \cdot G) \parallel [j]_{32} & \text{if } 0 \leq j < 2^{31} \text{ and } \text{int}(r) \in \mathbb{Z}_{|\mathbb{G}|}^+ \\ [0]_8 \parallel r \parallel [j]_{32} & \text{if } 0 \leq j < 2^{31} \text{ and } \text{int}(r) \notin \mathbb{Z}_{|\mathbb{G}|}^+ \end{cases} \quad (25)$$

$$\text{of}(z, I) = \begin{cases} I & \text{if } \text{int}(\text{left}_{256}(I)) \geq |\mathbb{G}| \\ [\text{int}(r) + \text{int}(\text{left}_{256}(I)) \bmod |\mathbb{G}|]_{256} \parallel \text{right}_{256}(I) & \text{if } 0 \leq \text{int}(\text{left}_{256}(I)) < |\mathbb{G}| \end{cases} \quad (26)$$

The above GGM instance is *not exactly* the same as a Bip32 key tree: it may contain nodes $z \in \{0, 1\}^{512}$ such that $\text{int}(\text{left}_{256}(z)) \notin \mathbb{Z}_{|\mathbb{G}|}^+$ cannot be interpreted as a “correct” signature private key (and lf computes differently from Bip32 key tree), while the Bip32 key tree discards such “unparsable” values. But the gap is limited, as will be reflected in the subsequent reduction in Sect. 6.5. In all, the algorithm $\text{DPriv}(\mathcal{G}, wsk_p, v_p, v_{p'})$ in Fig. 6 can be redefined based on the Co algorithm of the above GGM instance. In particular, when the “staring” node v_{\perp} is the root node, the key can be derived via querying the constraining oracle of F . This will be the core idea of our reduction in Sect. 6.5.

<pre> main Game^{muHEUF}_{WalPrim, \mathfrak{A}}(\mathbf{G}) for $i_0 = 1, \dots, u$ do $S_{i_0} \xleftarrow{\\$} \mathcal{S}$ $(\text{wsk}_{\perp}^{(i_0)}, \text{wpk}_{\perp}^{(i_0)}) \leftarrow \text{Setup}(\mathcal{G}_{i_0}, S_{i_0})$ $\text{mpk} \leftarrow (\text{wpk}_{\perp}^{(1)}, \dots, \text{wpk}_{\perp}^{(u)})$ $\mathcal{Q}_{\text{Crupt}} \leftarrow \emptyset, \mathcal{Q}_{\text{WSign}} \leftarrow \emptyset$ $(i_0, m^*, \sigma^*, v_p^{(i_0)}) \leftarrow$ $\mathfrak{A}_{\text{Crupt, PKReq, muWSign, Prim}}(\text{mpk})$ if $(i_0, v_p^{(i)}) \in \mathcal{Q}_{\text{Crupt}}$ or $(i_0, m^*, v_p^{(i)}) \in \mathcal{Q}_{\text{WSign}}$ then return 0 $pk_p^{(i_0)} \leftarrow \text{PKReq}(i_0, v_p^{(i_0)})$ if $\text{WVrfy}(pk_p^{(i_0)}, m^*, \sigma^*) = 0$ then return 0 return 1 Oracle muWSign($i_0, m, v_p^{(i_0)}$) $\text{wsk}_p^{(i_0)} \leftarrow \text{DPriv}(\mathcal{G}_{i_0}, \text{wsk}_{\perp}^{(i_0)}, v_{\perp}^{(i_0)}, v_p^{(i_0)})$ $\sigma \leftarrow \text{WSign}(\text{wsk}_p^{(i_0)}, m)$ $\mathcal{Q}_{\text{WSign}} \leftarrow \mathcal{Q}_{\text{WSign}} \cup \{(i_0, m, v_p^{(i_0)})\}$ return σ </pre>	<pre> Oracle PKReq($i_0, v_p^{(i_0)}$) // This oracle returns the public signature // key $pk_p^{(i_0)}$ of the node $v_p^{(i_0)}$ in the i_0-th // HDW instance to \mathfrak{A}, mimicking the // publicity of such keys $\text{wpk}_p^{(i_0)} \leftarrow \text{DPub}(\mathcal{G}_{i_0}, \text{wpk}_{\perp}^{(i_0)}, v_{\perp}^{(i_0)}, v_p^{(i_0)})$ Recover $pk_p^{(i_0)}$ from $\text{wpk}_p^{(i_0)}$ return $pk_p^{(i_0)}$ Oracle Crupt($i_0, v_p^{(i_0)}$) if $\exists v_{p'}^{(i_0)} \in \mathcal{Q}_{\text{Reqd}} : v_{p'}^{(i_0)} = v_p^{(i_0)}$ or $v_{p'}^{(i_0)} \in \text{Desc}(\mathcal{G}_{i_0}, v_p^{(i_0)})$ then return $\text{wsk}_p^{(i_0)} \leftarrow \text{DPriv}(\mathcal{G}_{i_0}, \text{wsk}_{\perp}^{(i_0)}, v_{\perp}^{(i_0)}, \text{wsk}_p^{(i_0)})$ $\mathcal{Q}_{\text{Crupt}} \leftarrow \mathcal{Q}_{\text{Crupt}} \cup \{v_p^{(i_0)}\} \cup \text{Desc}(\mathcal{G}_{i_0}, v_p^{(i_0)})$ return $\text{wsk}_p^{(i_0)}$ </pre>
---	---

Fig. 7. Multi-user HEUF security game $\text{Game}_{\text{WalPrim}, \mathfrak{A}}^{\text{muHEUF}}(\mathbf{G})$. The invoked Setup, DPub, DPriv, WSign, and WVrfy belongs to the wallet Wal^{Prim} .

6.4 mu Security Definitions for Hierarchical Deterministic Wallet

We follow Luzio et al.'s *hierarchical unforgeability* notion [LFA20], which allows an attacker to corrupt an arbitrary number of users (and their descendants) in the hierarchy, and challenges the attacker to forge a signature on behalf of an uncorrupted node. The ability of corruption is formalized by the oracle Crupt in Fig. 7, which also models multiple nodes colluding and sharing their secrets. Below we extend it into the mu setting. Also we consider the case wallet is built upon a public ideal function **Prim**, and write Wal^{Prim} to highlight. Another parameter that helps characterizing adversarial power is *the maximal number of allowed sessions*. Concretely, let $\mathbf{G} = (\mathcal{G}_1, \dots, \mathcal{G}_u)$ be a sequence of DAGs that defines the configurations for the u wallet users, and let $\mathcal{G}_{i_0} = (\mathcal{V}_{i_0}, \mathcal{E}_{i_0})$ for $i_0 = 1, \dots, u$. Then, the maximal number of allowed sessions is defined as $V(\mathbf{G}) := \sum_{i_0=1}^u |\mathcal{V}_{i_0}|$. I.e., the DAGs limit the total number of derived keys in the system.

Definition 7 (Multi-user HEUF Security). *An HDW scheme Wal^{Prim} is $(u, q_C, q_D, q_S, T, D, t, \varepsilon)$ -multi-user hierarchically existentially unforgeable, if for every sequence of u DAGs $\mathbf{G} = (\mathcal{G}_1, \dots, \mathcal{G}_u)$ such that $V(\mathbf{G}) \leq D$ and any adversary \mathfrak{A} making q_C queries to the Crupt oracle, q_D queries to the PKReq oracle, q_S queries to the signing oracle muWSign, and T queries to the ideal primitive **Prim**, it holds*

$$\Pr[\text{Game}_{\text{WalPrim}, \mathfrak{A}}^{\text{muHEUF}}(\mathbf{G}) = 1] \leq \varepsilon,$$

where the experiment $\text{Game}_{\text{WalPrim}, \mathfrak{A}}^{\text{muHEUF}}(\mathbf{G})$ is defined in Fig. 7.

Luzio et al. mentioned *hierarchical unlinkability* without a detailed formalism [LFA20]. This notion intends to capture that there is no “non-trivial” relation between distinct child public keys of the same node. We formalize this idea in the mu setting, allowing an attacker to corrupt an arbitrary number of users/sessions in the hierarchy, and

<pre> main Game^{muHULk}_{WalPrim, \mathfrak{A}, b}(\mathbf{G}) Initializes all entries of kTable to \perp for $i_0 = 1, \dots, u$ do $S_{i_0} \xleftarrow{\\$} \mathcal{S}$ $(wsk_{\perp}^{(i_0)}, wpk_{\perp}^{(i_0)}) \leftarrow \text{Setup}(\mathcal{G}_{i_0}, S_{i_0})$ endfor $\text{mpk} \leftarrow (wpk_{\perp}^{(1)}, \dots, wpk_{\perp}^{(u)})$ $\mathcal{Q}_{\text{Crupt}} \leftarrow \emptyset, \mathcal{Q}_{\text{Reqd}} \leftarrow \emptyset$ return $\mathfrak{A}^{\text{Crupt, PKReq, muWSign, Prim}}(\text{mpk})$ Oracle muWSign($i_0, m, v_p^{(i_0)}$) if $b = 0$ then if $\text{kTable}[i_0, p] = \perp$ then $(sk_p^{(i_0)}, pk_p^{(i_0)}) \xleftarrow{\\$} \mathcal{K}_{\text{sign}}$ $\text{kTable}[i_0, p] \leftarrow (sk_p^{(i_0)}, pk_p^{(i_0)})$ endif $(sk_p^{(i_0)}, pk_p^{(i_0)}) \leftarrow \text{kTable}[i_0, p]$ else // $b = 1$ $wsk_p^{(i_0)} \leftarrow \text{DPriv}(\mathcal{G}_{i_0}, wsk_{\perp}^{(i_0)}, v_{\perp}^{(i_0)}, v_p^{(i_0)})$ Recover $sk_p^{(i_0)}$ from $wsk_p^{(i_0)}$ endif $\sigma \leftarrow \text{Sig.Sig}(sk_p^{(i_0)}, m)$ return σ </pre>	<pre> Oracle Crupt($i_0, v_p^{(i_0)}$) if $\exists v_{p'}^{(i_0)} \in \mathcal{Q}_{\text{Reqd}} : v_{p'}^{(i_0)} = v_p^{(i_0)}$ or $v_{p'}^{(i_0)} \in \text{Desc}(\mathcal{G}_{i_0}, v_p^{(i_0)})$ then return $wsk_p^{(i_0)} \leftarrow \text{DPriv}(\mathcal{G}_{i_0}, wsk_{\perp}^{(i_0)}, v_{\perp}^{(i_0)}, wsk_p^{(i_0)})$ $\mathcal{Q}_{\text{Crupt}} \leftarrow \mathcal{Q}_{\text{Crupt}} \cup \{v_p^{(i_0)}\} \cup \text{Desc}(\mathcal{G}_{i_0}, v_p^{(i_0)})$ return $wsk_p^{(i_0)}$ Oracle PKReq($i_0, v_p^{(i_0)}$) if $v_p^{(i_0)} \in \mathcal{Q}_{\text{Crupt}}$ then return \perp if $b = 0$ then if $\text{kTable}[i_0, p] = \perp$ then $(sk_p^{(i_0)}, pk_p^{(i_0)}) \xleftarrow{\\$} \mathcal{K}_{\text{sign}}$ endif $(sk_p^{(i_0)}, pk_p^{(i_0)}) \leftarrow \text{kTable}[i_0, p]$ else // $b = 1$ $wpk_p^{(i_0)} \leftarrow \text{DPub}(\mathcal{G}_{i_0}, wpk_{\perp}^{(i_0)}, v_{\perp}^{(i_0)}, v_p^{(i_0)})$ Recover $pk_p^{(i_0)}$ from $wpk_p^{(i_0)}$ endif $\mathcal{Q}_{\text{Reqd}} \leftarrow \mathcal{Q}_{\text{Reqd}} \cup \{v_p^{(i_0)}\}$ return $pk_p^{(i_0)}$ </pre>
--	---

Fig. 8. Multi-user muHULk security game $\text{Game}_{\text{WalPrim}, \mathfrak{A}, b}^{\text{muHULk}}(\mathbf{G})$.

challenging it to distinguish public signature keys of uncorrupted nodes from random points. This is formalized by the oracle PKReq in Fig. 8, which returns the true requested public signature key when $b = 1$, and random public key when $b = 0$. These two definitions essentially capture desired security in the intended use case “Per-office balances” in [Med18].

Definition 8 (Multi-user HULk Security). *An HDW scheme Wal^{Prim} is $(u, q_C, q_D, q_S, T, D, t, \varepsilon)$ -multi-user hierarchically unlinkable, if for every sequence of u DAGs $\mathbf{G} = (\mathcal{G}_1, \dots, \mathcal{G}_u)$ such that $V(\mathbf{G}) \leq D$ and any adversary \mathfrak{A} making q_C queries to the Crupt oracle, q_D queries to the oracle PKReq, q_S queries to the signing oracle muWSign, and T queries to the ideal primitive Prim, it holds*

$$\left| \Pr[\text{Game}_{\text{WalPrim}, \mathfrak{A}, 0}^{\text{muHULk}}(\mathbf{G}) = 1] - \Pr[\text{Game}_{\text{WalPrim}, \mathfrak{A}, 1}^{\text{muHULk}}(\mathbf{G}) = 1] \right| \leq \varepsilon,$$

where the experiment $\text{Game}_{\text{WalPrim}, \mathfrak{A}, b}^{\text{muHULk}}(\mathbf{G})$ is defined in Fig. 8.

6.5 Multi-user Security of Bip32

For simplicity and clearness, we assume using parameters close to the 256-bit elliptic curve secp256k1 recommended in [Med18]. In the standard Bip32, the double-input function Prim is instantiated by HMAC, i.e., $\text{Prim}(x, y) := \text{HMAC}(y, x)$ (using y as the key of HMAC). We note that the first input x is not always secret in Bip32. By Dodis et al. [DRST12, Theorem 4.4], with such parameters, HMAC instantiates a keyed FIL random oracle. Therefore, this subsection considers $\text{Prim}(x, y) = \text{KH}(x, y)$ for a keyed RO $\text{KH} : \{0, 1\}^{512} \times \{0, 1\}^{256} \mapsto \{0, 1\}^{512}$, and our results are formally stated as follows.

Theorem 4. *The Bip32^{KH} HDW built upon a keyed random oracle $\mathbf{KH} : \{0, 1\}^{512} \times \{0, 1\}^{256} \mapsto \{0, 1\}^{512}$ is $(u, q_C, q_D, q_S, T, D, t, \varepsilon)$ -multi-user hierarchically existentially unforgeable, where (assuming $D \leq |\mathbb{G}|/2$)*

$$\varepsilon \leq \frac{2u(T+D)}{2^s} + \frac{u^2}{2^s} + \frac{D}{|\mathbb{G}|} + \frac{2^9 \cdot (T+D)}{2^{256}} + D \times \left(1 - \frac{|\mathbb{G}| - 1}{2^{256}}\right). \quad (27)$$

The Bip32^{KH} HDW built upon a keyed random oracle $\mathbf{KH} : \{0, 1\}^{512} \times \{0, 1\}^{256} \mapsto \{0, 1\}^{512}$ and a $(q_S, q_S, O(t+D), \varepsilon_{\text{muUFCMA}})$ -mu unforgeable digital signature scheme is $(u, q_C, q_D, q_S, T, D, t, \varepsilon)$ -multi-user hierarchically existentially unforgeable, where (assuming $D \leq |\mathbb{G}|/2$)

$$\varepsilon \leq \frac{u(T+D)}{2^s} + \frac{u^2}{2^s} + \frac{D}{|\mathbb{G}|} + \frac{2^9 \cdot (T+D)}{2^{256}} + D \times \left(1 - \frac{|\mathbb{G}| - 1}{2^{256}}\right) + \varepsilon_{\text{muUFCMA}}. \quad (28)$$

When the 256-bit elliptic curve domain parameters recommended in [Res10, Sect. 2.4] is used, it can be checked that $D \times \left(1 - \frac{|\mathbb{G}| - 1}{2^{256}}\right) \leq D/2^{127}$ (as claimed in [Med18]) and $|\mathbb{G}|/2 \approx 2^{255}$.

As mentioned, with the parameters in Bip32, HMAC instantiates a keyed FIL random oracle $\mathbf{KH} : \{0, 1\}^{512} \times \{0, 1\}^{256} \mapsto \{0, 1\}^{512}$ with 256-bit security [DRST12, Theorem 4.4]. Thus, this is not the bottleneck. Otherwise, Eq. (27) indicates computational security of $\min\{247, s - \log_2 u\}$ bits, as long as the total number D of active Bip32 sessions across the world does not exceed $\min\{2^s/u, |\mathbb{G}|/2, 2^{127}\}$. Whereas Eq. (28) indicate computational security of $\min\{247, s - \log_2 u, f(q_S)\}$ bits, as long as the signature scheme delivers $f(q_S)$ bit muUFCMA security (Definition 2). Importantly, this means the amount of computations needed for attacks does *not* decrease as the number D of “active sessions” increases.

Consider the case of $s \geq 256$. In the su setting, i.e., $u = 1$, the amount of computations needed to break unlinkability with a notable success probability is 2^{247} . The $\log_2 u$ bit loss in the multi-user setting is inevitable, as it matches the effort for guessing one out of s -bit seeds S_1, \dots, S_u . Fortunately, u reflects the number of “enterprises” using Bip32 excluding their sub-accounts, and would not be too large.

On the other hand, the mu security loss due to the signature scheme has to be seriously addressed, as it straightforwardly determines the concrete unforgeability security *even in the su setting*. Any signature with sufficient mu security could be employed: see Sect. 2. For example, using 256-bit secret keys, the mu security of Schnorr is of 128 bits [KMP16], giving rise to $\min\{247, s - \log_2 u, 128\}$ bits computational security. We are not aware of ECDSA variants with good mu bounds, which is a natural open problem.

In the remaining of this subsection, we devote to prove the unlinkability and unforgeability claims in turn.

Unlinkability. Consider any adversary \mathfrak{A}_1 against the muHULk security of Bip32^{KH}. By construction, the security game $G_1 = \text{Game}_{\text{Wal}^{\text{KH}}, \mathfrak{A}_1, 0}^{\text{muHULk}}(\mathbb{G})$ first invokes $\text{Setup}(\mathcal{G}_{i_0}, S_{i_0})$ for $i_0 = 1, \dots, u$, which derives the u master keys or root nodes $\text{Nd}(1, \perp) = \text{wsk}_{\perp}^{(1)} = \mathbf{KH}(S_1, \text{“Bitcoin seed”}), \dots, \text{Nd}(u, \perp) = \text{wsk}_{\perp}^{(u)} = \mathbf{KH}(S_u, \text{“Bitcoin seed”})$ for the u instances. Similarly for the game $G_1^* = \text{Game}_{\text{Wal}^{\text{KH}}, \mathfrak{A}_1, 1}^{\text{muHULk}}(\mathbb{G})$. We first replace these nodes by uniform and independent strings. After this, we construct an adversary \mathfrak{A}_2 that simulates the muHULk security game using the CPRF oracles in front of \mathfrak{A}_1 . The advantage of \mathfrak{A}_1 is then bounded by Theorem 1. We describe the two steps in the subsequent two paragraphs in turn.

The initial derivation We start by replacing the u 512-bit root nodes $\text{Nd}(1, \perp), \dots, \text{Nd}(u, \perp)$ by u 512-bit independent and uniform strings, and obtain an intermediate game \mathbf{G}_2 . The gap between \mathbf{G}_1 and \mathbf{G}_2 is at most $u(T+D)/2^s + u^2/2^{s+1}$, i.e., the mu PRF security of $\mathbf{KH}(S, \text{“Bitcoin seed”})$ (viewing S as the secret key):

- (i) The u master key derivation calls $\text{Nd}(i_0, \perp) = \mathbf{KH}(S_{i_0}, \text{“Bitcoin seed”})$, $i_0 = 1, \dots, u$, are fresh and produce u random strings $\text{Nd}(1, \perp), \dots, \text{Nd}(u, \perp)$, as long as none of the u arguments ($S_1, \text{“Bitcoin seed”}$), ..., ($S_u, \text{“Bitcoin seed”}$) collides with the other queries to \mathbf{KH} in this system. Here “other queries” include queries made by both \mathfrak{A}_1 and the subsequent derivations in the Bip32 instances. The number of the former type is at most T , while the number of the latter type does not exceed $V(\mathbf{G}) \leq D$. Therefore, the probability of this type of collision event is at most $u(T+D)/2^s$;
- (ii) The u outputs $\text{Nd}(1, \perp), \dots, \text{Nd}(u, \perp)$ are independent if and only if the u seeds S_1, \dots, S_u are collision free. This probability is at most $\binom{u}{2} \times \frac{1}{2^s} \leq u^2/2^{s+1}$.

Similarly, we obtain an intermediate game \mathbf{G}_2^* from $\mathbf{G}_1^* = \text{Game}_{\text{WalKH}, \mathfrak{A}_1, 1}^{\text{muHULK}}(\mathbf{G})$ via replacing the u 512-bit root nodes by random, with a gap at most $u(T+D)/2^s + u^2/2^{s+1}$.

\mathbf{G}_2 versus \mathbf{G}_2^* using Theorem 1 For the remaining argument, \mathfrak{A}_2 runs \mathfrak{A}_1 and simulates the oracles as follows, and outputs \mathfrak{A}_1 's decision bit at the end.

- Upon \mathfrak{A}_1 querying $\text{Crupt}(i_0, v_{i_1/\dots/i_{d'}}^{(i_0)})$, \mathfrak{A}_2 queries $\text{muCo}_K^{\perp}(i_0, i_1/\dots/i_{d'})$ to have the constrained key $\text{Nd}(i_0, i_1/\dots/i_{d'})$ and the leakages $\text{left}_{256}(\text{Nd}(i_0, \perp)), \dots, \text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'}))$. \mathfrak{A}_2 returns \perp to \mathfrak{A}_1 , if any of the $d'+1$ 256-bit integers $\text{int}(\text{left}_{256}(\text{Nd}(i_0, \perp))), \dots, \text{int}(\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'})))$ is not in $\mathbb{Z}_{|\mathcal{G}|}^+$ (recall from Sect. 6.3 that this is consistent with the actual Bip32 specification). Otherwise, \mathfrak{A}_2 sets $sk_{i_1/\dots/i_{d'}}^{(i_0)} \leftarrow \text{int}(\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'})))$, $ch_{i_1/\dots/i_{d'}}^{(i_0)} \leftarrow \text{right}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'}))$, and passes $\text{wsk}_{i_1/\dots/i_{d'}}^{(i_0)} = (sk_{i_1/\dots/i_{d'}}^{(i_0)}, ch_{i_1/\dots/i_{d'}}^{(i_0)})$ to \mathfrak{A}_1 .
- Upon \mathfrak{A}_1 querying $\text{muWSign}(i_0, m, v_{i_1/\dots/i_{d'}}^{(i_0)})$, \mathfrak{A}_2 queries $\text{muCo}_K^{\perp}(i_0, i_1/\dots/i_{d'})$ for the constrained key $\text{Nd}(i_0, i_1/\dots/i_{d'})$ and the leakages $\text{left}_{256}(\text{Nd}(i_0, \perp)), \dots, \text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'}))$. If any of these $d'+1$ integers $\text{int}(\text{left}_{256}(\text{Nd}(i_0, \perp))), \dots, \text{int}(\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'})))$ is not in $\mathbb{Z}_{|\mathcal{G}|}^+$ then \mathfrak{A}_2 returns \perp to \mathfrak{A}_1 . Otherwise, \mathfrak{A}_2 computes the chain code $ch_{i_1/\dots/i_{d'}}^{(i_0)} \leftarrow \text{right}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'}))$ and the secret signing key $sk_{i_1/\dots/i_{d'}}^{(i_0)} \leftarrow \text{int}(\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'})))$ and returns $\text{WSign}((sk_{i_1/\dots/i_{d'}}^{(i_0)}, ch_{i_1/\dots/i_{d'}}^{(i_0)}), m)$ to \mathfrak{A}_1 .
- Upon \mathfrak{A}_1 querying $\text{PKReq}(i_0, v_{i_1/\dots/i_{d'}}^{(i_0)})$, \mathfrak{A}_2 pinpoints the longest path in \mathcal{G}_{i_0} that contains the node $v_j^{(i_0)}$. Formally, \mathfrak{A}_2 pinpoints $v_0^{(i_0)} \xrightarrow{i_1} v_{i_1}^{(i_0)} \xrightarrow{i_2} \dots \xrightarrow{i_{d'}} v_{i_1/\dots/i_{d'}}^{(i_0)}$ such that the outdegree of $v_{i_1/\dots/i_{d'}}^{(i_0)}$ is zero. \mathfrak{A}_1 then queries $\text{muEv}_K^{\perp}(i_0, i_1/\dots/i_{d'})$ ¹² to have $\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'}))$ and the leakages $\text{left}_{256}(\text{Nd}(i_0, \perp)), \dots, \text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'}))$. Again, if any of the $d'+1$ integers $\text{int}(\text{left}_{256}(\text{Nd}(i_0, \perp))), \dots, \text{int}(\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'})))$ is not in $\mathbb{Z}_{|\mathcal{G}|}^+$, then \mathfrak{A}_2 returns \perp to \mathfrak{A}_1 . Otherwise, \mathfrak{A}_2 computes $sk_{i_1/\dots/i_{d'}}^{(i_0)} \leftarrow \text{int}(\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'})))$ and $pk_{i_1/\dots/i_{d'}}^{(i_0)} \leftarrow sk_{i_1/\dots/i_{d'}}^{(i_0)} \cdot G$ and returns $pk_{i_1/\dots/i_{d'}}^{(i_0)}$ to \mathfrak{A}_1 .
- Upon \mathfrak{A}_1 querying \mathbf{KH} , \mathfrak{A}_2 simply relays the query and response.

¹² \mathfrak{A}_2 cannot simply query $\text{muCo}_K^{\perp}(i_0, i_1/\dots/i_{d'})$, otherwise it cannot react to \mathfrak{A}_1 querying $\text{PKReq}(i_0, v_{i_1/\dots/i_{d'}}^{(i_0)})$ later.

When \mathfrak{A}_2 is interacting with the real world $(\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{KH})$, \mathfrak{A}_1 is playing with the game \mathbf{G}_2 : the actions are exactly the same, despite the gap between the GGM instance and the Bip32 key tree.

On the other hand, when \mathfrak{A}_2 is interacting with the ideal $(\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{KH})$, a public key $pk_{i_1/\dots/i_{d'}}$ returned by \mathfrak{A}_2 due to \mathfrak{A}_1 querying $\text{PKReq}(i_0, v_{i_1/\dots/i_{d'}}^{(i_0)})$ query is computed via $pk_{i_1/\dots/i_{d'}}^{(i_0)} \leftarrow \text{int}(r_{i_1/\dots/i_{d'}}^{(i_0)}) \cdot G$, where $r_{i_1/\dots/i_{d'}}^{(i_0)} \xleftarrow{\mathbf{S}} \{0, 1\}^{256}$ is a block of simulated leakage (it corresponds to $\text{left}_{256}(\text{Nd}(i_0, i_1/\dots/i_{d'}))$ in the real world). This resembles the game \mathbf{G}_2^* , except that \mathfrak{A}_2 may respond \mathfrak{A}_1 with \perp . Denote this event by “ $\mathfrak{A}_2^{\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{KH}}$ returns \perp ”. Then, as long as this event does not occur, the computed key $pk_{i_1/\dots/i_{d'}}^{(i_0)}$ is uniformly distributed in \mathbb{G} . Moreover, the key computed due to any other query $\text{PKReq}(i'_0, v_{i'_1/\dots/i'_{d''}}^{(i'_0)})$ has $pk_{i'_1/\dots/i'_{d''}}^{(i'_0)} \leftarrow \text{int}(r_{i'_1/\dots/i'_{d''}}^{(i'_0)}) \cdot G$ and $r_{i'_1/\dots/i'_{d''}}^{(i'_0)}$ is independent from $r_{i_1/\dots/i_{d'}}^{(i_0)}$. Therefore, as long as the event “ $\mathfrak{A}_2^{\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{KH}}$ returns \perp ” does not occur, \mathfrak{A}_2 perfectly emulates the game \mathbf{G}_2^* in front of \mathfrak{A}_1 , i.e.,

$$\left| \Pr[\mathbf{G}_2^* = 1] - \Pr[\mathfrak{A}_2^{\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{KH}} = 1] \right| \leq \Pr[\mathfrak{A}_2^{\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{KH}} \text{ returns } \perp] \leq D \times \left(1 - \frac{|\mathbb{G}| - 1}{2^{256}}\right),$$

and thus

$$\left| \Pr[\mathbf{G}_2^* = 1] - \Pr[\mathbf{G}_2 = 1] \right| \leq \left| \Pr[\mathfrak{A}_2^{\text{muCo}_{\mathbf{K}}^{\mathbf{L}}, \text{muEv}_{\mathbf{K}}^{\mathbf{L}}, \mathbf{KH}} = 1] - \Pr[\mathfrak{A}_2^{\text{muCo}_{\mathbf{K}}^{\mathbf{S}}, \mathbf{R}^{\mathbf{S}}, \mathbf{KH}} = 1] \right| + D \times \left(1 - \frac{|\mathbb{G}| - 1}{2^{256}}\right).$$

Concrete bounds It remains to calculate the concrete bounds. Regardless of the adversarial strategy, it can be seen the effective data complexity cannot exceed $V(\mathbf{G}) \leq D$. For the quantity C , using the uniformness of the private/public keys (in the ideal world) and a multi-collision argument (we defer the details to the next paragraph), when $C = 256 = 2^8$ and $D \leq |\mathbb{G}|/2$ (so that $2D/|\mathbb{G}| \leq 1$), it can be shown that $\Pr[\mu(\mathcal{I}) \geq C] \leq D/2|\mathbb{G}|$. Then, using Theorem 1, we obtain

$$\left| \Pr[\mathbf{G}_2^* = 1] - \Pr[\mathbf{G}_2 = 1] \right| \leq \frac{D}{|\mathbb{G}|} + \frac{2^9 \cdot (T + D)}{2^{256}} + D \times \left(1 - \frac{|\mathbb{G}| - 1}{2^{256}}\right). \quad (29)$$

Gathering this with the gaps between \mathbf{G}_1 , \mathbf{G}_2 , \mathbf{G}_1^* and \mathbf{G}_2^* (which is $2 \times (u(T + D)/2^s + u^2/2^{s+1})$) yields Eq. (27).

Proof of $\Pr[\mu(\mathcal{I}) \geq C] \leq D/2|\mathbb{G}|$ For the quantity $\mu(\mathcal{L})$, we will rely on the uniformness of the private/public keys. In detail, consider any $(i_0, p, j) \neq (i'_0, p', j')$. Then,

- Case 1: $0 \leq j < 2^{31}$, whereas $2^{31} \leq j' < 2^{32}$. Then it is impossible to have $\text{lf}_{p,j}(\text{left}_{256}(\text{Nd}(i_0, p)), \text{pp}_{i_0}) = \text{lf}_{p',j'}(\text{left}_{256}(\text{Nd}(i'_0, p')), \text{pp}_{i'_0})$, since the former has leftmost byte $[0]_8$ while the latter has $[2]_8$ or $[3]_8$;
- Case 2: $2^{31} \leq j < 2^{32}$, $0 \leq j' < 2^{31}$. Then $\text{lf}_{p,j}(\text{left}_{256}(\text{Nd}(i_0, p)), \text{pp}_{i_0}) \neq \text{lf}_{p',j'}(\text{left}_{256}(\text{Nd}(i'_0, p')), \text{pp}_{i'_0})$ which resembles Case 1;
- Case 3: $0 \leq j, j' < 2^{31}$. Then $\text{lf}_{p,j}(\text{left}_{256}(\text{Nd}(i_0, p)), \text{pp}_{i_0}) \neq \text{lf}_{p',j'}(\text{left}_{256}(\text{Nd}(i'_0, p')), \text{pp}_{i'_0})$ holds with probability at most $\Pr[sk, sk' \xleftarrow{\mathbf{S}} \mathbb{Z}_{|\mathbb{G}|} : sk = sk'] = 1/|\mathbb{G}|$;
- Case 4: $2^{31} \leq j, j' < 2^{32}$. Then $\text{lf}_{p,j}(\text{left}_{256}(\text{Nd}(i_0, p)), \text{pp}_{i_0}) \neq \text{lf}_{p',j'}(\text{left}_{256}(\text{Nd}(i'_0, p')), \text{pp}_{i'_0})$ holds with probability at most $\Pr[pk, pk' \xleftarrow{\mathbf{S}} \mathbb{G} : \text{ser}_P(pk) = \text{ser}_P(pk')]$. The map $\text{ser}_P : \mathbb{G} \mapsto \{0, 1\}^{264}$ is bijective, and thus the probability is $1/|\mathbb{G}|$.

By the above, in any case, the probability to have $\text{lf}_{p,j}(\text{left}_{256}(\text{Nd}(i_0, p)), \text{pp}_{i_0}) = \text{lf}_{p',j'}(\text{left}_{256}(\text{Nd}(i'_0, p')), \text{pp}_{i'_0})$ does not exceed $1/|\mathbb{G}|$. Thus, for any integer $C \geq 2$,

$$\begin{aligned} \Pr[\mu(\mathcal{L}) \geq C] &= \Pr\left[\exists (i_0^{(1)}, p^{(1)}, j^{(1)}), \dots, (i_0^{(C)}, p^{(C)}, j^{(C)}) : \text{lf}_{p^{(1)},j^{(1)}}(\text{left}_{256}(\text{Nd}(i_0^{(1)}, p^{(1)})), \text{pp}_{i_0^{(1)}}) \right. \\ &\quad \left. = \dots = \text{lf}_{p^{(C)},j^{(C)}}(\text{left}_{256}(\text{Nd}(i_0^{(C)}, p^{(C)})), \text{pp}_{i_0^{(C)}})\right] \\ &\leq \binom{D}{C} \cdot \left(\frac{1}{|\mathbb{G}|}\right)^{C-1}. \end{aligned}$$

For $C = 256$, using $|\mathbb{G}| < 2^{256}$ in our context and as long as $D \leq |\mathbb{G}|/2$ (so that $2D/|\mathbb{G}| \leq 1$), it further holds

$$\Pr[\mu(\mathcal{L}) \geq C] \leq \frac{|\mathbb{G}|}{C!} \left(\frac{D}{|\mathbb{G}|}\right)^C \leq \frac{1}{C!} \left(\frac{2D}{|\mathbb{G}|}\right)^{256} \leq \frac{1}{256!} \frac{2D}{|\mathbb{G}|} \leq \frac{D}{2|\mathbb{G}|}.$$

Unforgeability. Let $\mathbf{G}_1 = \text{Game}_{\text{Wal}^{\text{muHEUF}}, \mathfrak{A}_1}(\mathbf{G})$ be the real hierarchical unforgeability game. We modify \mathbf{G}_1 by replacing the oracles PKReq and muWSign with the oracles PKReq and muWSign with $b = 0$ in Fig. 8. Denote by \mathbf{G}_2 the obtained modified game. It is easy to see that, the gap between \mathbf{G}_1 and \mathbf{G}_2 is the already established unlinkability bound, i.e.,

$$\left| \Pr[\mathbf{G}_2 = 1] - \Pr[\mathbf{G}_1 = 1] \right| \leq \frac{2u(T+D)}{2^s} + \frac{u^2}{2^s} + \frac{D}{|\mathbb{G}|} + \frac{2^9 \cdot (T+D)}{2^{256}} + D \times \left(1 - \frac{|\mathbb{G}| - 1}{2^{256}}\right).$$

Signing keys generated in the game \mathbf{G}_2 are independent and uniformly distributed. Thus, the forgery probability $\Pr[\mathbf{G}_2 = 1]$ is bounded by the *multi-user existential unforgeability security* of the signature in use (see Definition 2). Informally, consider an adversary \mathfrak{A}_2 having access to q_S signing oracles of the signature instantiated with q_S independent secret keys, and it makes q_S signing queries and runs in time $O(t_{\mathfrak{A}_1})$ and succeeds as long as it forges for *any* of the q_S keys. If the signature is $(q_S, q_S, O(t+D), \varepsilon_{\text{muUFCMA}})$ -mu unforgeable, then the success probability of \mathfrak{A}_2 is bounded by $\varepsilon_{\text{muUFCMA}}$. These establish the bound in Eq. (28).

6.6 Improving Prim for Bip32, and Performance of MPC Implementations

Since the security of Bip32 (variants) simply follow from our results on GGGM, sound improvements addressing Lindell’s question [Lin19] become clear. In detail, the initial seed S is of $128 \leq s \leq 512$ bits. The string “Bitcoin seed” is of 96 bits. However, as the security of the entire system cannot exceed 256 bits due to the limitation $\kappa = 256$ —and in fact, practical uses typically adopt $s = 256$,—we focus on the case $128 \leq s \leq 256$, and pad the seed S to $S \parallel [0]_{296-s}$, i.e., of 296 bits. With these considerations, we describe our Bip32 variants using a general function $\mathbf{Prim} : \{0, 1\}^{296} \times \{0, 1\}^{256} \rightarrow \{0, 1\}^{256}$.

- (i) The HMAC-call in $\text{Setup}(\mathcal{G}, S)$ is replaced by $I_{\perp} \leftarrow \mathbf{Prim}(S, \text{“Bitcoin seed”})$;
- (ii) The HMAC-calls in $\text{CKDpriv}(\text{wsk}, i)$ are replaced by $I \leftarrow \mathbf{Prim}([0]_s \parallel [sk]_{256} \parallel [i]_{32}, ch)$ and $I \leftarrow \mathbf{Prim}(\text{ser}_P(pk) \parallel [i]_{32}, ch)$ correspondingly;
- (iii) The HMAC-call in $\text{CKDpub}(\text{wpk}, i)$ is replaced by $I \leftarrow \mathbf{Prim}(\text{ser}_P(pk) \parallel [i]_{32}, ch)$.

We then propose two instantiations of **Prim** with less AND gates as follows.

- **SHACAL instantiation** $\text{Bip32}^{\text{SHACAL3}}$: define the function $\text{Prim}(x, y) := \text{SHACAL3}(x, y \parallel [0]_{256}) \oplus (y \parallel [0]_{256})$, i.e., using a single call to $\text{SHACAL3} : \{0, 1\}^{1024} \times \{0, 1\}^{512} \mapsto \{0, 1\}^{512}$, the block cipher underlying SHA512. The reliability is essentially the same as $\text{Bip32}^{\text{HMAC}}$, since the two wallets rely on the same assumption (i.e., the security of SHACAL3).
- **Fast instantiation** $\text{Bip32}^{\text{kp800}}$: define $\text{Prim}(x, y) := \text{Trunc}_{288}(\text{KECCAK-}p[800, 11](x \parallel y \parallel [0]_{248}))$, i.e., using a call to the permutation $\text{KECCAK-}p[800, 11]$, which is a member of the $\text{KECCAK-}p$ permutation family [BDP⁺17] with number of rounds halved (this follows the SHA3 designers, and its still has security margin of 6 rounds [BDP⁺17]). To shorten notations, we abbreviate $\text{KECCAK-}p[800, 11]$ as **kp800**.

The security analysis of $\text{Bip32}^{\text{kp800}}$ just follows Bip32^{KH} in Sect. 6.5, since truncating $\text{KECCAK-}p[800, 11]$ also yields an FIL RO with 248-bit security [CLL19, Theorem 1]. By these, $\text{Bip32}^{\text{kp800}}$ achieves $\min\{247, s - \log_2 u\}$ bits unlinkability and $\min\{247, s - \log_2 u, f(q_S)\}$ unforgeability (as long as we model the permutation **kp800** as a *public random 800-bit permutation*). On the other hand, analysis of $\text{Bip32}^{\text{SHACAL3}}$ follows Bip32^{KH} , except that the terms in Eq. (29) are replaced with terms from Theorem 2. According to Theorems 1 and 2, the concrete bounds differ by only a factor of 2. Thus, $\text{Bip32}^{\text{SHACAL3}}$ achieves $\min\{246, s - \log_2 u\}$ bits unlinkability and $\min\{246, s - \log_2 u, f(q_S)\}$ unforgeability. Concrete mu security of the two improved instantiations are thus comparable with the original **Bip32** standard.

We benchmark the two-party protocols based on garbled circuit which securely compute various instantiations of **Bip32**. Our implementations focus on the (most widely deployed) *default configuration* of the *key tree* [Med18] (see Fig. 1). In particular, we consider the process of deriving the shares of a session key $sk_{i_1/i_2/i_3}$ from the shares of the seed S . As discussed and as shown in Fig. 1, this consists of 4 **Prim** executions (in the MPC manner). We use EMP-Toolkit [WMK16] as the backend of our implementations, and choose the state-of-the-art garbling schemes with semi-honest security. Our experiments are performed in a virtual machine with an Intel(R) Core(TM) i5-1038NG7 CPU at 2.0GHZ with localhost communication, and the performance is reported in Table 1. By the results, $\text{Bip32}^{\text{SHACAL3}}$ achieves much better performance while retaining the same reliability as $\text{Bip32}^{\text{HMAC}}$, while $\text{Bip32}^{\text{kp800}}$ achieves the best performance with a moderate security margin.

Table 1. Performance of Bip32 instantiations. # AND presents the number of involved AND gates, and roughly match theoretical results on 4 **Prim** executions. They deliver similar mu security: see Sect. 6.5 and 6.6.

Scheme	# AND	Time (ms)	Comm. (KB)	Unlink. Sec. (bit)	Note
$\text{Bip32}^{\text{HMAC}}$	944628	87	29790	$\min\{247, s - \log_2 u\}$	Standard
$\text{Bip32}^{\text{SHACAL3}}$	245172	33	7937	$\min\{246, s - \log_2 u\}$	Reliability
$\text{Bip32}^{\text{kp800}}$	49908	19	1835	$\min\{247, s - \log_2 u\}$	High perform.

7 Improving Function Secret Sharing

The state-of-art. Function secret sharing of point functions [GI14] crucially relies on GGM trees functioning as PPRFs. The closest provably secure construction [GKWY20]

instantiates the classical GGM with the length-doubling PRG $G^{\text{AES}_{fk}}(s) := \text{AES}_{fk}(s) \oplus s \parallel \text{AES}_{fk}(s \oplus [1]_{128}) \oplus s \oplus [1]_{128}$, where AES_{fk} is the AES using a fixed, publicly-known key fk (as discussed in Sect. 3.1). To characterize its concrete mu security, we model AES as an ideal cipher $\mathbf{E} : \{0, 1\}^{wn} \times \{0, 1\}^{wn} \mapsto \{0, 1\}^{wn}$. Then, this state-of-the-art $\text{FBTr}^{\text{DM}^{\text{AES}}}$ (i.e., *FSS's Binary Tree*), is an instance of $\text{GGGM}^{\text{DM}^{\text{E}}}$ given in Sect. 4.3, with $\text{Prim}(x, y) := \text{AES}_x(y) \oplus y$ and parameters as follows.

- Node size equals the security parameter $n = \kappa = 128$. Note that in our model, this means *there is no “leakage” at all*, and we are essentially in the *classical puncturable PRF setting*;
- $\theta = 2$, and $w = 1$, i.e., branch number $w\theta = 2$, and output size $wn = n = 128$;
- $\text{sf}_{p,j}(\text{right}_\kappa(z)) = \text{sf}_{p,j}(z) = z \oplus [j]_{128}$ (which is indeed injective);
- $\text{pp} = fk$ which is a fixed constant, and $\text{lf}_{p,j}(\perp, \text{pp}) = fk$, i.e., labels don't depend on p, j at all. The restriction $j \neq j' \Rightarrow (\text{lf}_{p,j}(\perp, \text{pp}), \text{sf}_{p,j}(z)) \neq (\text{lf}_{p,j'}(\perp, \text{pp}), \text{sf}_{p,j'}(z))$ holds since $\text{sf}_{p,j}(z) = z \oplus [j]_{128} \neq z \oplus [j']_{128} = \text{sf}_{p,j'}(z)$;
- For any $z, I \in \{0, 1\}^{128}$, $\text{of}(z, I) = I$. This satisfies restriction (ii), i.e., $\text{of}(z, I) = \text{of}(\text{left}_{128-128}(z), I) = I$.
- Depth $d = \log_2 N_{\text{leaf}}$, where N_{leaf} is the desired number of leaves.

FSS uses a $\text{FBTr}^{\text{DM}^{\text{AES}}}$ instance as a *puncturable PRF (PPRF)*: see Definition 5. In detail, assume that the MPC system executes u FSS instances using independent keys K_1, \dots, K_u . For $i_0 = 1, \dots, u$, the i_0 -th FSS instance generates a *punctuated key* $K_{i_0}\{p_{i_0}\}$ for a single punctured point $p_{i_0} \in \{0, 1\}^d$, i.e., an entity holding the punctured key $K_{i_0}\{p_{i_0}\}$ is able to evaluate exactly $|\mathcal{X}| - 1$ functions values, *except for* $\text{FBTr}^{\text{DM}^{\text{AES}}}. \text{Ev}(K_{i_0}, p_{i_0})$. The i_0 -th punctured key $K_{i_0}\{p_{i_0}\}$ can be written in $128d$ bits. The effective data complexity is $2ud$, since (it is easy to count) the number of AES calls in each tree is exactly $2d$.

It remains to determine the quantity C . Unfortunately, the label $\text{lf}_{p,j}(\perp, \text{pp}) = fk$ is a constant, and we can only expect $C = 2ud$ and $\varepsilon_\mu = \Pr[\mu(\mathcal{L}) > C] = 0$. Injecting all the above parameters into Theorem 2, we conclude that $\text{FBTr}^{\text{DM}^{\text{AES}}}$ is a (u, T, D, ε) -puncturable PRF with effective data complexity $D = 2ud$ and

$$\varepsilon \leq \frac{8udT + 16(ud)^2}{2^{128}}. \quad (30)$$

As a concrete example, consider the 2PC discussed in the introduction: assuming $N_{2\text{pc}} = 2^{14}$ 2PC protocol instances and each protocol instantiating $N_{\text{tr}} = 2^{26}$ FSS or $\text{FBTr}^{\text{DM}^{\text{AES}}}$ tree instances with depth $d \leq 16$. Then, $u = N_{2\text{pc}} \times N_{\text{tr}} = 2^{40}$ and $8ud = 2^{47}$, indicating security up to $T \approx 2^{128}/2^{47} = 2^{81}$ AES “queries” or computations (i.e., a theoretical degradation up to 36.7%). Such computations are hardly infeasible.

Improving mu security and flexibility. We first address the mu security degradation, which mainly requires to overcome the factor $u = N_{2\text{pc}} \cdot N_{\text{tr}}$ due to the numerous instances. Following Sect. 6.6, we define the public parameter $\text{pp} := \text{IV}$ for $\text{IV} \xleftarrow{\$} \{0, 1\}^{128}$ that is *picked at uniform during the setup of every FSS instance*, and propose to use AES_{IV} instead of AES_{fk} . Then, labels $\text{lf}_{*,*}(\perp, \text{pp}_{i_0}) = \text{IV}_{i_0}$ and $\text{lf}_{*,*}(\perp, \text{pp}'_{i_0}) = \text{IV}'_{i_0}$ in distinct trees collide only if $\text{IV}_{i_0} = \text{IV}'_{i_0}$. Thus, for any threshold $t \geq 2$, if there is no t -collision $\text{IV}_{i_0}^{(1)} = \dots = \text{IV}_{i_0}^{(t)}$ among the u trees, then the number of tree instances with their IVs equaling a certain value is at most $t - 1$. As all IVs are uniform and independent, the probability to have a t -collision is $\binom{u}{t} / 2^{128(t-1)} \leq \frac{u^t}{t! \cdot 2^{128(t-1)}}$.

Then, observing that increasing branches trades punctured key size for computations and improves flexibility, we propose to “naturally” increase the “parallelization degree” of $\text{FBTr}^{\text{DMAES}}$ to $\theta \geq 3$, i.e., using large-expansion PRG $G(s) = \text{AES}_{\text{IV}}(s) \oplus s \parallel \text{AES}_{\text{IV}}(s \oplus [1]_{128}) \oplus s \oplus [1]_{128} \parallel \dots \parallel \text{AES}_{\text{IV}}(s \oplus [\theta - 1]_{128}) \oplus s \oplus [\theta - 1]_{128}$. We denote this by $\text{FMTr}^{\text{DMAES}}$, meaning *FSS’s Multi-branch Tree*.

- The depth decreases to $d = \log_{\theta} N_{\text{leaf}}$;
- The punctured key for a single tree has $128(\theta - 1)\log_{\theta} N_{\text{leaf}} + 128$ bits (including the 128-bit $\text{pp} = \text{IV}$);
- The effective data complexity $D = u\theta\log_{\theta} N_{\text{leaf}}$.

By the above, for $C = (t - 1)\theta\log_{\theta} N_{\text{leaf}}$, $\Pr[\mu(\mathcal{L}) > C] \leq \frac{u^t}{t! \cdot 2^{128(t-1)}}$, meaning that $\text{FMTr}^{\text{DMAES}}$ is a (u, T, D, ε) -puncturable PRF, with

$$\varepsilon \leq \frac{2u^t}{t! \cdot 2^{128(t-1)}} + \frac{4(t-1)\theta T \cdot \log_{\theta} N_{\text{leaf}} + 4(t-1)u(\theta\log_{\theta} N_{\text{leaf}})^2}{2^{128}}. \quad (31)$$

Eq. (31) holds for *any* threshold t . Since the parameters u, T, N_{leaf} are incomparable, it is difficult to conclude on optimizing t . But certain choices already yield satisfactory bounds. E.g., with $t = 5$, Eq. (31) becomes

$$\varepsilon \leq \frac{u^5}{2^{517}} + \frac{16\theta T \cdot \log_{\theta} N_{\text{leaf}} + 16u(\theta\log_{\theta} N_{\text{leaf}})^2}{2^{128}}, \quad (32)$$

indicating security up to $T \approx 2^{112}$ computations and $u \approx 2^{103}$ users.

More concretely, when $\theta = 2$, Eq. (32) indicates security up to $\approx 2^{128}/2^9 = 2^{119}$ computations (which is nearly optimal for $\kappa = 128$) and running the FSS protocol $u \leq 2^{103}$ times. The additional computation compared to the state-of-art FSS protocol is the mere generation of a 128-bit (pseudo)random IV, which is negligible. The punctured key is $128 \cdot \log_2 N_{\text{leaf}} + 128$ bits. Setting θ to 4, Eq. (32) indicates the same security, whereas the computation cost for each input is reduced by 50% due to the halved depth at the expense of a 50% blow-up in the punctured key size. We list the parameters in Table 2 for clarity.

Table 2. Comparison with prior works. The execution time is benchmarked on an AWS machine of type `m5.large`. N_{leaf} represents the number of leaves in the generated FSS trees. For clearness, the mu security column demonstrates numerical results for the specific case $N_{\text{leaf}} = 2^{16}$ and $u = N_{2\text{pc}} \times N_{\text{tr}} = 2^{40}$ trees.

Scheme	# Prim	Cost of Prim (ns)	Total time (ns)	Seed size (Bytes)	mu sec. (bits)
State-of-art $\text{FBTr}^{\text{DMAES}}$	N_{leaf}	3.5	$3.5N_{\text{leaf}}$	$16 \log_2 N_{\text{leaf}}$	81
$\text{FMTr}^{\text{DMAES}}$ with $\theta = 2$	N_{leaf}	3.5	$3.5N_{\text{leaf}}$	$16 \log_2 N_{\text{leaf}}$	119
$\text{FMTr}^{\text{DMAES}}$ with $\theta = 4$	$N_{\text{leaf}}/3$	7	$2.3N_{\text{leaf}}$	$24 \log_2 N_{\text{leaf}}$	119

In Table 2, we compare the performance of our protocol against the state-of-the-art. $\text{FMTr}^{\text{DMAES}}$ with $\theta = 2$ has the same overhead as $\text{FBTr}^{\text{DMAES}}$, while $\text{FMTr}^{\text{DMAES}}$ with $\theta = 4$ incurs a 50% increase in seed size (and communication) but $1.5\times$ faster computation. In major applications such as MPC based on preprocessed correlation [BCG⁺19a], the communication overhead due to the GGM punctuated key is less than 10%. Thus, using $\text{FMTr}^{\text{DMAES}}$ with $\theta = 4$ accelerates computations much with only insignificant communication overheads.

Acknowledgments

We thank the anonymous reviewers for their insightful suggestions. Chun Guo was supported by the National Key Research and Development Program of China under Grant No. 2022YFA1004900, the National Natural Science Foundation of China (Grant No. 62002202) and the Taishan Scholars Program (for Young Scientists) of Shandong. Xiao Wang was supported by NSF awards #2016240 and #2236819. Yu Yu was supported by the National Natural Science Foundation of China (Grant Nos. 62125204 and 92270201), the National Key Research and Development Program of China (Grant No. 2018YFA0704701), and the Major Program of Guangdong Basic and Applied Research (Grant No. 2019B030302008). Yu Yu also acknowledges the support from the XPLOER PRIZE.

References

- ADE⁺20. Nabil Alkeilani Alkadri, Poulami Das, Andreas Erwig, Sebastian Faust, Juliane Krämer, Siavash Riahi, and Patrick Struck. Deterministic wallets in a quantum world. In *ACM Conf. on Computer and Communications Security (CCS) 2020*, pages 1017–1031. ACM Press, 2020.
- AGKK19. Myrto Arapinis, Andriana Gkaniatsou, Dimitris Karakostas, and Aggelos Kiayias. A formal treatment of hardware wallets. In *Financial Cryptography and Data Security (FC)*, LNCS, pages 426–445. Springer, 2019.
- AHS20. Jean-Philippe Aumasson, Adrian Hamelink, and Omer Shlomovits. A survey of ECDSA threshold signing. *Cryptology ePrint Archive*, Report 2020/1390, 2020. <https://eprint.iacr.org/2020/1390>.
- BBM00. Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In *Advances in Cryptology—Eurocrypt 2000*, volume 1807 of *LNCS*, pages 259–274. Springer, 2000.
- BBT16. Mihir Bellare, Daniel J. Bernstein, and Stefano Tessaro. Hash-function based PRFs: AMAC and its multi-user security. In *Advances in Cryptology—Eurocrypt 2016, Part I*, volume 9665 of *LNCS*, pages 566–595. Springer, 2016.
- BCG⁺19a. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Peter Rindal, and Peter Scholl. Efficient two-round OT extension and silent non-interactive secure computation. In *ACM Conf. on Computer and Communications Security (CCS) 2019*, pages 291–308. ACM Press, 2019.
- BCG⁺19b. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In *Advances in Cryptology—Crypto 2019, Part III*, volume 11694 of *LNCS*, pages 489–518. Springer, 2019.
- BCGI18. Elette Boyle, Geoffroy Couteau, Niv Gilboa, and Yuval Ishai. Compressing vector OLE. In *ACM Conf. on Computer and Communications Security (CCS) 2018*, pages 896–912. ACM Press, 2018.
- BDP⁺17. Guido Bertoni, Joan Daemen, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. Team Keccak, 2017.
- Ber15. Daniel J. Bernstein. Multi-user Schnorr security, revisited. *Cryptology ePrint Archive*, Report 2015/996, 2015. <https://eprint.iacr.org/2015/996>.

- BGI14. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *Intl. Conference on Theory and Practice of Public Key Cryptography*, LNCS, pages 501–519. Springer, 2014.
- BGI15. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In *Advances in Cryptology—Eurocrypt 2015, Part II*, volume 9057 of LNCS, pages 337–367. Springer, 2015.
- BGI16. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In *ACM Conf. on Computer and Communications Security (CCS) 2016*, pages 1292–1303. ACM Press, 2016.
- BHT18. Priyanka Bose, Viet Tung Hoang, and Stefano Tessaro. Revisiting AES-GCM-SIV: Multi-user security, faster key derivation, and better bounds. In *Advances in Cryptology—Eurocrypt 2018, Part I*, volume 10820 of LNCS, pages 468–499. Springer, 2018.
- BLW17. Dan Boneh, Kevin Lewi, and David J. Wu. Constraining pseudorandom functions privately. In *Intl. Conference on Theory and Practice of Public Key Cryptography 2017, Part II*, LNCS, pages 494–524. Springer, 2017.
- BMRS21. Carsten Baum, Alex J. Malozemoff, Marc B. Rosen, and Peter Scholl. Mac’n’cheese: Zero-knowledge proofs for boolean and arithmetic circuits with nested disjunctions. In *Advances in Cryptology—Crypto 2021, Part IV*, LNCS, pages 92–122. Springer, 2021.
- BT16. Mihir Bellare and Björn Tackmann. The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In *Advances in Cryptology—Crypto 2016, Part I*, volume 9814 of LNCS, pages 247–276. Springer, 2016.
- BW13. Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In *Advances in Cryptology—Asiacrypt 2013, Part II*, volume 8270 of LNCS, pages 280–300. Springer, 2013.
- CBM15. Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *IEEE Symp. Security and Privacy 2015*, pages 321–338. IEEE, 2015.
- CDMP05. Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In *Advances in Cryptology—Crypto 2005*, volume 3621 of LNCS, pages 430–448. Springer, 2005.
- CK16. Aloni Cohen and Saleet Klein. The GGM function family is a weakly one-way family of functions. In *9th Theory of Cryptography Conference—TCC 2016*, LNCS, pages 84–107. Springer, 2016.
- CLL19. Wonseok Choi, ByeongHak Lee, and Jooyoung Lee. Indifferentiability of truncated random permutations. In *Advances in Cryptology—Asiacrypt 2019, Part I*, LNCS, pages 175–195. Springer, 2019.
- CS14. Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In *Advances in Cryptology—Eurocrypt 2014*, volume 8441 of LNCS, pages 327–350. Springer, 2014.
- CT21. Yu Long Chen and Stefano Tessaro. Better security-efficiency trade-offs in permutation-based two-party computation. In *Advances in Cryptology - ASIACRYPT 2021*, pages 275–304, 2021.
- DEF⁺21. Poulami Das, Andreas Erwig, Sebastian Faust, Julian Loss, and Siavash Riahi. The Exact Security of BIP32 Wallets. In *ACM Conf. on Computer and Communications Security (CCS) 2021*. ACM Press, 2021.
- Des88. Yvo Desmedt. Society and group oriented cryptography: A new concept. In *Advances in Cryptology—Crypto 1987*, LNCS, pages 120–127. Springer, 1988.

- DEs16. Jack Doerner, David Evans, and abhi shelat. Secure stable matching at scale. In *ACM Conf. on Computer and Communications Security (CCS) 2016*, pages 1602–1613. ACM Press, 2016.
- DFL19. Poulami Das, Sebastian Faust, and Julian Loss. A formal treatment of deterministic wallets. In *ACM Conf. on Computer and Communications Security (CCS) 2019*, pages 651–668. ACM Press, 2019.
- DIO21. Samuel Dittmer, Yuval Ishai, and Rafail Ostrovsky. Line-Point Zero Knowledge and Its Applications. In *2nd Conference on Information-Theoretic Cryptography (ITC 2021)*. Schloss Dagstuhl, 2021.
- DRST12. Yevgeniy Dodis, Thomas Ristenpart, John P. Steinberger, and Stefano Tessaro. To hash or not to hash again? (In)differentiability results for H^2 and HMAC. In *Advances in Cryptology—Crypto 2012*, volume 7417 of *LNCS*, pages 348–366. Springer, 2012.
- Ds17. Jack Doerner and abhi shelat. Scaling ORAM for secure computation. In *ACM Conf. on Computer and Communications Security (CCS) 2017*, pages 523–535. ACM Press, 2017.
- Exo21. Exodus. Preliminary Offering Circular dated February 26, 2021. https://www.sec.gov/Archives/edgar/data/1821534/000114036121006439/nt10013846x8_1a.htm, 2021.
- FKPR14. Georg Fuchsbauer, Momchil Konstantinov, Krzysztof Pietrzak, and Vanishree Rao. Adaptive security of constrained PRFs. In *Advances in Cryptology—Asiacrypt 2014, Part II*, volume 8874 of *LNCS*, pages 82–101. Springer, 2014.
- GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- GI14. Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In *Advances in Cryptology—Eurocrypt 2014*, volume 8441 of *LNCS*, pages 640–658. Springer, 2014.
- GKW⁺20. Chun Guo, Jonathan Katz, Xiao Wang, Chenkai Weng, and Yu Yu. Better concrete security for half-gates garbling (in the multi-instance setting). In *Advances in Cryptology—Crypto 2020, Part II*, *LNCS*, pages 793–822. Springer, 2020.
- GKWY20. Chun Guo, Jonathan Katz, Xiao Wang, and Yu Yu. Efficient and secure multiparty computation from fixed-key block ciphers. In *IEEE Symp. Security and Privacy 2020*, pages 825–841. IEEE, 2020.
- HKKW19. Dennis Hofheinz, Akshay Kamath, Venkata Koppula, and Brent Waters. Adaptively secure constrained pseudorandom functions. In *Financial Cryptography and Data Security (FC)*, *LNCS*, pages 357–376. Springer, 2019.
- HKW15. Susan Hohenberger, Venkata Koppula, and Brent Waters. Adaptively secure puncturable pseudorandom functions in the standard model. In *Advances in Cryptology—Asiacrypt 2015, Part I*, volume 9452 of *LNCS*, pages 79–102. Springer, 2015.
- KMP16. Eike Kiltz, Daniel Masny, and Jiaxin Pan. Optimal security proofs for signatures from identification schemes. In *Advances in Cryptology—Crypto 2016, Part II*, volume 9815 of *LNCS*, pages 33–61. Springer, 2016.
- KPTZ13. Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In *ACM Conf. on Computer and Communications Security (CCS) 2013*, pages 669–684. ACM Press, 2013.

- Lac18. Marie-Sarah Lacharité. Security of BLS and BGLS signatures in a multi-user setting. *Cryptogr. Commun.*, 10(1):41–58, 2018.
- LBC16. LBC. The LBC Server. <https://lbc.cryptoguru.org/>, 2016.
- LFA20. Adriano Di Luzio, Danilo Francati, and Giuseppe Ateniese. Arcula: A secure hierarchical deterministic wallet for multi-asset blockchains. In *CANS 20 International Conference on Cryptology and Network Security*, LNCS, pages 323–343. Springer, 2020.
- Lin19. Yehuda Lindell. A Full CryptoCurrency Custody Solution Based on MPC and Threshold ECDSA. Real World Crypto 2019, 2019. <https://rwc.iacr.org/2019/slides/Multiparty-ECDSA-RWC2019.pdf>.
- LMO⁺14. Jake Longo, Daniel P. Martin, Elisabeth Oswald, Daniel Page, Martijn Stam, and Michael Tunstall. Simulatable leakage: Analysis, pitfalls, and new constructions. In *Advances in Cryptology—Asiacrypt 2014, Part I*, volume 8873 of LNCS, pages 223–242. Springer, 2014.
- Med18. Mediawiki. BIP32 Specification. <https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki>, 2018.
- Men17. Bart Mennink. Insuperability of the standard versus ideal model gap for tweakable blockcipher security. In *Advances in Cryptology—Crypto 2017, Part II*, volume 10402 of LNCS, pages 708–732. Springer, 2017.
- MPs19. Antonio Marcedone, Rafael Pass, and abhi shelat. Minimizing trust in hardware wallets with two factor signatures. In *Financial Cryptography and Data Security (FC)*, LNCS, pages 407–425. Springer, 2019.
- NIS21. NIST. Multi-Party Threshold Cryptography | CSRC. <https://csrc.nist.gov/projects/threshold-cryptography>, 2021.
- Pat09. Jacques Patarin. The “coefficients H” technique (invited talk). In *Annual International Workshop on Selected Areas in Cryptography (SAC) 2008*, volume 5381 of LNCS, pages 328–345. Springer, 2009.
- Res10. Certicom Research. Sec 2: Recommended elliptic curve domain parameters. Standards for Efficient Cryptography, 2010. <https://www.secg.org/sec2-v2.pdf>.
- SGRR19. Phillipp Schoppmann, Adrià Gascón, Leonie Reichert, and Mariana Raykova. Distributed vector-OLE: Improved constructions and implementation. In *ACM Conf. on Computer and Communications Security (CCS) 2019*, pages 1055–1072. ACM Press, 2019.
- SPY13. François-Xavier Standaert, Olivier Pereira, and Yu Yu. Leakage-resilient symmetric cryptography under empirically verifiable assumptions. In *Advances in Cryptology—Crypto 2013, Part I*, volume 8042 of LNCS, pages 335–352. Springer, 2013.
- ST16. Thomas Shrimpton and R. Seth Terashima. Salvaging weak security bounds for blockcipher-based constructions. In *Advances in Cryptology—Asiacrypt 2016, Part I*, LNCS, pages 429–454. Springer, 2016.
- TVR16. Mathieu Turuani, Thomas Voegtlin, and Michaël Rusinowitch. Automated verification of electrum wallet. In *FC 2016 Workshops*, volume 9604 of LNCS, pages 27–42, February 2016.
- WMK16. Xiao Wang, Alex J. Malozemoff, and Jonathan Katz. EMP-toolkit: Efficient MultiParty computation toolkit. <https://github.com/emp-toolkit>, 2016.
- WYG⁺17. Frank Wang, Catherine Yun, Shafi Goldwasser, Vinod Vaikuntanathan, and Matei Zaharia. Splinter: Practical Private Queries on Public Data. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 299–313. USENIX Association, 2017.

- WYKW21. Chenkai Weng, Kang Yang, Jonathan Katz, and Xiao Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In *IEEE Symp. Security and Privacy 2021*. IEEE, 2021.
- YWL⁺20. Kang Yang, Chenkai Weng, Xiao Lan, Jiang Zhang, and Xiao Wang. Ferret: Fast extension for correlated OT with small communication. In *ACM Conf. on Computer and Communications Security (CCS) 2020*, pages 1607–1626. ACM Press, 2020.

A The LBC Bitcoin Cracking Project

Briefly speaking, given a pair of signature keys (sk_1, pk_1) , a Bitcoin user could sign with sk_1 to spend funds associated with the public “address” $addr_1 = \text{RIPEMD160}(\text{SHA256}(pk_1))$, i.e., a 160-bit hash digest of the public key pk_1 .¹³ But this means the map from signature keys to “addresses” is not bijective, and an adversary holding (sk', pk') with $\text{RIPEMD160}(\text{SHA256}(pk')) = addr_1$ can also spend funding associated with the address $addr_1$. Since $|addr_1| = 160$, the success probability to guess such a key pair (sk', pk') with T computations is nearly $T/2^{160}$ (this is much higher than $T/2^{256}$ of straightforwardly guessing the “original” secret key sk_1). While $T/2^{160}$ remains small, with u targeted Bitcoin users $(sk_1, pk_1), \dots, (sk_u, pk_u)$, the probability to successfully guess (sk', pk') such that $\text{RIPEMD160}(\text{SHA256}(pk')) = addr_i = \text{RIPEMD160}(\text{SHA256}(pk_i))$ for some i increases to $\frac{uT}{2^{160}}$. It can be seen this is a variant of multi-user secret (signing) key recovery attack, and the concrete security is of $160 - \log_2 u$ bits which degrades significantly with the number of targeted Bitcoin users. The LBC Bitcoin cracking project instantiates this idea in a distributed manner to gather computation power from multiple participants, and it recovered more than a dozen secret (signing) keys in 2016 and 2017 [LBC16].

It is tempting to ask how our treatments capture this attack. To clarify, note that our security definitions in Sect. 6.2 implicitly assume that *the cryptocurrency address of an account is its public signature key, rather than the hash digest of the key*. This ensures a bijective mapping between cryptocurrency addresses and signature key pairs (sk, pk) , and the success probability to guess a key pair (sk, pk) corresponding to a certain address with T computations is $T/2^{256}$. In this case, the success probability of the above multi-user attack decreases to (much smaller) $uT/2^{256}$.

¹³ This slightly deviates from the choice mentioned in Sect. 6.1, i.e., simply using pk_1 as the address.